NetIQ® Cloud Manager

Administrator Reference

June 2016

Legal Notice

For information about legal notices, trademarks, disclaimers, warranties, export and other use restrictions, U.S. Government rights, patent policy, and FIPS compliance, see https://www.netiq.com/company/legal/.

Copyright © 2016 NetIQ Corporation. All Rights Reserved.

Contents

About This Book and the Library 11

Part I Orchestration Server Components 13

1 Interface Layout of the Orchestration Server Console 15

2 The Orchestration Server and the Server Admin Objects 17

2.1 Orchestration Server Object . . . 17
2.1.1 The Orchestration Server Info/Configuration Page . . . 18
2.1.2 Orchestration Server Authentication Page . . . 26
2.1.3 Orchestration Server Policies Page . . . 30
2.1.4 Orchestration Server Constraints/Facts Page . . . 30
2.2 Server Admin Object . . . 30

3 The Job Object 31

3.1 Job Groups . . . 31
3.2 The Job Info/Groups Page . . . 32
3.2.1 Info . . . 32
3.2.2 Groups . . . 41
3.3 The Job Configuration Page . . . 41
3.4 The JDL Editor Page . . . 42
3.5 The Job Library Editor Page . . . 42
3.6 The Job Policies Page . . . 43
3.7 The Job Constraints/Facts Page . . . 44

4 The Resource Object 45

4.1 Resource Groups . . . 45
4.2 Resource Info/Groups Page . . . 45
4.2.1 Info Panel . . . 46
4.2.2 Groups Panel . . . 71
4.3 Provision Info Page . . . 71
4.4 Resource Log Page . . . 72
4.5 Resource Policies Page . . . 72
4.6 Resource Health Debugger Page . . . 72
4.7 Resource Constraints/Facts Page . . . 72
4.8 Resource Object Naming and Renaming . . . 73

5 The VM Host Object 75

5.1 Info Page . . . 75
5.1.1 Show Inherited Fact Values Check Box . . . 76
5.1.2 VM Host Information Panel . . . 76
5.1.3 Provisioning Adapter Config Panel . . . 79
5.1.4 Guest VM Monitor Information Panel . . . 79
5.2 Policies Page . . . 80
5.3 Health Debugger Page . . . 80

5.4 Constraints/Facts Page . . . 81
5.5 Action History Page . . . 81
5.6 VM Host Object Naming and Renaming . . . 81
5.7 Unique VM Host Cluster Facts . . . 82
5.7.1 Orchestration Server Facts in the VM Host Cluster Object . . . 82
5.7.2 Orchestration Server Facts in a VM Host Residing in a Cluster . . . 83
5.7.3 Orchestration Server Facts in VMs Hosted in Clusters . . . 84
5.8 vCPU Slots for VM Hosts . . . 85
5.8.1 Configuring vCPUs on VM Hosts . . . 85
5.8.2 Configuring vCPUs on VM Host Clusters . . . 86
5.8.3 Configuring vCPUs on VMs . . . 86

6 The Virtual Disk Object 89

6.1 Understanding the Virtual Disk Object . . . 89
6.1.1 Creating or Deleting a vDisk in the Orchestration Console . . . 89
6.1.2 Sharing Virtual Disks Among VM Hosts . . . 92
6.1.3 Moving Virtual Disks . . . 93
6.2 Viewing Virtual Disk Configuration in the Orchestration Console . . . 94
6.2.1 Virtual Disk Information Panel . . . 95
6.2.2 Virtual Disk Policies Tab . . . 97
6.2.3 Virtual Disk Health Debugger Tab . . . 97
6.2.4 Virtual Disk Constraints/Facts Tab . . . 97
6.2.5 Virtual Disk Object Naming and Renaming . . . 98
6.3 Managing Block Devices as VM Virtual Disks . . . 99
6.3.1 How Block Device Support Works . . . 99
6.3.2 Viewing the Physical Disk Configuration in the Orchestration Console . . . 101

7 The Virtual NIC Object 107

7.1 Understanding the Virtual NIC Object . . . 107
7.1.1 The Purpose of the Virtual NIC . . . 107
7.1.2 Creating or Deleting a vNIC in the Orchestration Console . . . 107
7.2 Viewing the Virtual NIC Configuration in the Orchestration Console . . . 109
7.2.1 Virtual NIC Info Panel . . . 109
7.2.2 Virtual NIC Policies Tab . . . 112
7.2.3 Virtual NIC Health Debugger Tab . . . 112
7.2.4 Virtual NIC Constraints/Facts Tab . . . 112
7.2.5 Virtual NIC Object Naming and Renaming . . . 113

8 The Network Group and its Virtual Bridge Objects 115

8.1 Understanding the Network Group and Virtual Bridge Objects . . . 115
8.1.1 Virtual Bridge Object . . . 115
8.1.2 The Purpose of the Virtual Bridge . . . 116
8.1.3 Creating or Deleting a vBridge in the Orchestration Console . . . 116
8.1.4 Virtual Bridge Object Naming and Renaming . . . 118
8.2 Viewing the Virtual Bridge Configuration in the Orchestration Console . . . 118
8.2.1 Virtual Bridge Info/Groups Tab . . . 119
8.2.2 Virtual Bridge Policies Tab . . . 120
8.2.3 Virtual Bridge Health Debugger Tab . . . 120
8.2.4 Virtual Bridge Constraints/Facts Tab . . . 120

9 The Repository Object 121

9.1 Right-Click Menu Actions on the Repository Object . . . 121
9.2 Repository Groups . . . 122

9.3 Repository Info/Groups Tab . . . 122
9.3.1 Info Panel . . . 122
9.3.2 Best Practices for Entering Repository File Paths . . . 127
9.3.3 Groups . . . 129
9.4 Repository Policies Tab . . . 129
9.5 Repository Health Debugger Tab . . . 129
9.6 Repository Constraints/Facts Tab . . . 129
9.7 The Repository Action History Tab . . . 129
9.8 Repository Object Naming and Renaming . . . 130
9.9 Shared Storage for Disk Images . . . 130
9.9.1 Setting Disk Discovery Facts . . . 131
9.9.2 Running the Discovery . . . 131
9.9.3 Sharing Disks Between VMs . . . 132
9.9.4 Attaching a Discovered Disk to a VM . . . 132
9.9.5 Using Attached Disks in the Guest OS . . . 133

10 The User Object 135

10.1 User Groups . . . 135
10.2 User Info/Groups Tab . . . 135
10.2.1 Info . . . 136
10.2.2 Groups . . . 141
10.3 User Policies Tab . . . 141
10.4 User Health Debugger Tab . . . 141
10.5 User Constraints/Facts Tab . . . 142
10.6 The User Action History Tab . . . 142

11 Miscellaneous Objects Displayed in the Explorer Tree 143

11.1 Policy Object . . . 143
11.1.1 Policy Constraints . . . 143
11.1.2 Policy Facts . . . 143
11.2 Computed Fact Objects . . . 143
11.3 Event Objects . . . 144
11.4 Metrics Objects . . . 144

12 The Orchestration Server Job Scheduler 145

12.1 Understanding the Job Scheduler View . . . 145
12.1.1 Navigating The Job Schedules Table . . . 146
12.2 Understanding the Job Schedule Editor . . . 148
12.2.1 Schedule Name . . . 149
12.2.2 Job . . . 149
12.2.3 User . . . 149
12.2.4 Priority . . . 150
12.2.5 Description . . . 150
12.2.6 Matching Resources . . . 150
12.2.7 Test Schedule Now . . . 150
12.2.8 Triggers . . . 150
12.2.9 Job Arguments . . . 155
12.2.10 User Environment . . . 156
12.2.11 Constraints . . . 156
12.3 Understanding Cron Syntax in the Job Scheduler . . . 157
12.3.1 Format . . . 157
12.3.2 Special Characters . . . 157
12.3.3 Examples of Cron Syntax . . . 159
12.3.4 Cron Scheduling Precautions . . . 160

13 The Policy Debugger **161**

13.1 Constraints Table View . . . 161
13.1.1 Match Context Area . . . 162
13.1.2 Constraint Type List . . . 163
13.1.3 Verbose Check Box . . . 163
13.1.4 Capable Resources Summary . . . 164
13.1.5 Constraints Column of the Constraints Table View . . . 164
13.1.6 Policy Column of the Constraints Table . . . 165
13.2 Facts Table View . . . 166
13.2.1 All Facts Check Box . . . 166
13.3 Policy Debugger Use Cases . . . 167

14 Grid Object Health Monitoring **169**

14.1 Health Facts . . . 169
14.1.1 Explicitly Set or Cleared by the Administrator . . . 169
14.1.2 Set by Using a Discovery Job . . . 170
14.1.3 Set by Using a Policy . . . 170
14.1.4 Set by Using a Computed Fact . . . 170
14.1.5 Set Automatically by Using a Health Constraint . . . 170
14.2 Health Events . . . 171
14.3 Health Debugger . . . 171
14.3.1 Constraints Table Panel . . . 172
14.3.2 Facts Table View . . . 175

15 Events **177**

15.1 Understanding the Orchestration Server Events System . . . 177
15.1.1 Event Notification . . . 177
15.1.2 Built-in Events . . . 177
15.1.3 Rule-Based Events . . . 178
15.2 Event Object Visualization and Management in the Orchestration Console . . . 180
15.2.1 Deploying a New Rule-Based Event . . . 180
15.2.2 Deploying a Pre-written Rule-Based Event . . . 181
15.2.3 Undeploying an Event . . . 181
15.2.4 Event Editor . . . 181
15.3 Event Debugger . . . 182
15.3.1 Constraints Table . . . 182
15.3.2 The Facts Table . . . 184

16 The Metrics Facility **185**

16.1 Metrics Facility Functionality . . . 185
16.2 Ganglia Metrics . . . 185
16.3 How Does the Metrics Facility Impact Orchestration Server Performance? . . . 186
16.3.1 I/O Contention . . . 187
16.3.2 Too Many Open Files . . . 187
16.4 RRD Definition Using Deployable .metric Files . . . 187
16.4.1 XML Format for Deployable .metric Definitions . . . 188
16.5 Query of Aggregated Metric Values . . . 189
16.5.1 Example of a JDL Query for Aggregated Metric Values . . . 189
16.5.2 Example of a Policy Constraint or Event Constraint Using Aggregated Metric Values . . . 189
16.5.3 Example of Using Non-aggregated (“Raw”) Historical Metric Values . . . 189
16.6 MetricsManager MBean API . . . 190
16.6.1 MBean Methods Exposed by the MetricsManager Facility . . . 190
16.6.2 The MetricsDeployer Facility . . . 191
16.7 Using the Metrics Facility in the Orchestration Console . . . 191

Part II Orchestration Server Tools **193**

17 Orchestration Server zos Command Line Tool **195**

17.1 List of zos Commands . . . 195
17.2 Getting Started with the zos Command . . . 196
17.2.1 Logging In to the Orchestration Server . . . 196
17.2.2 Running a Job . . . 197
17.2.3 Monitoring a Job . . . 197
17.3 Details, Usage, and Syntax Examples of zos Commands . . . 199
cancel . . . 200
cat . . . 202
copy . . . 203
delete . . . 205
dir . . . 206
event . . . 207
fail . . . 208
head . . . 209
info . . . 211
jobinfo . . . 213
joblist . . . 215
jobs . . . 216
listen . . . 218
log . . . 219
login . . . 221
logout . . . 223
mkdir . . . 224
move . . . 226
password . . . 227
pause . . . 228
priority . . . 230
resume . . . 232
run . . . 233
search . . . 235
status . . . 237
tail . . . 239

18 Orchestration Server zosadmin Command Line Tool **243**

18.1 List of zosadmin Commands . . . 243
18.2 Getting Started with the zosadmin Command . . . 244
18.2.1 Logging In . . . 244
18.2.2 Checking Login Status . . . 245
18.2.3 Logging Out . . . 245
18.3 Details, Usage, and Syntax Examples of zosadmin Commands . . . 246
auditclean . . . 247
auditcount . . . 248
auditreport . . . 249
cancelalljobs . . . 251
create . . . 252
deploy . . . 255
disconnect . . . 256
dump . . . 257
facts . . . 258
get . . . 259
init . . . 260
invoke . . . 261
list . . . 262
login . . . 263

logout . . . 265
nodes . . . 266
password . . . 267
redeploy . . . 268
rotatelogs . . . 269
sessions . . . 270
set . . . 272
start . . . 273
status . . . 275
stop . . . 276
undeploy . . . 277
upgrade . . . 278
users . . . 280
verify . . . 281

Part III Orchestration Provisioning Adapters 283

19 The VMware vSphere Provisioning Adapter 285

19.1 Provisioning Actions Supported by the vSphere Provisioning Adapter . . . 285
19.2 vSphere Server Changes Synchronized to Cloud Manager by vSphereUpdater . . . 286
19.2.1 VM Changes . . . 287
19.2.2 Repository Changes . . . 287
19.2.3 Network Changes . . . 287
19.2.4 Resource Pool Changes . . . 287
19.3 Apply Config and Save Config Actions on vSphere VMs . . . 288
19.3.1 Expected Results of the Apply Config Action on Windows Server VMs . . . 288
19.3.2 Expected Results of the Save Config Action on Windows Server VMs . . . 289
19.3.3 Expected Results of the Apply Config Action on Linux Server VMs . . . 290
19.3.4 Expected Results of the Save Config Action on Linux Server VMs . . . 291
19.4 Discovering or Dynamically Creating vSphere Port Groups . . . 291
19.4.1 How the Orchestration Server Models a vSphere Port Group and Exposes vSwitch Settings . . . 292
19.4.2 Dynamically Creating Port Groups . . . 292
19.5 Enabling VNC Access to vSphere 5 VM Guest Consoles . . . 293
19.5.1 Enabling VNC Access By Opening Multiple Firewall Ports . . . 294
19.5.2 Enabling VNC Access by Creating a Special Configuration File . . . 294
19.6 Known Configuration Limitations for vSphere VMs . . . 296

Part IV Orchestration Server and Provisioning Adapter Troubleshooting 297

20 Troubleshooting Cloud Manager Orchestration Issues 299

20.1 Troubleshooting Installation Issues . . . 299
20.2 Troubleshooting Orchestration Server Upgrade Issues . . . 300
20.3 Troubleshooting Orchestration Server Operation Issues . . . 300
20.4 Troubleshooting Orchestration Console Issues . . . 302
20.5 Troubleshooting Orchestration Agent Issues . . . 303
20.6 Troubleshooting General VM Management Issues . . . 304
20.7 Troubleshooting vSphere VM Provisioning Actions . . . 308
20.8 Multicast Troubleshooting . . . 312

21 Increasing the Kernel ARP Threshold Value on the Orchestration Server 313

21.1 Threshold Definitions . . . 313
21.2 Determining the Current Kernel Threshold Value . . . 313

21.3 Changing the Current Kernel Threshold Value . . . 314
21.3.1 Editing the /etc/sysctl.conf File . . . 314
21.3.2 Making Live Changes to the Threshold Values . . . 314

22 Orchestration Server Logs 315

Part V Application Server and Console Features and Functions 317

23 The File Upload Feature 319

23.1 Prerequisites to Using the Feature . . . 319
23.2 Feature Behaviors . . . 319
23.3 The “Preflight” Check . . . 320
23.4 Stages of the File Upload . . . 320
23.4.1 The File Uploads to the Application Server . . . 320
23.4.2 The File is Published from the Application Server to the Orchestration Grid Repository . . . 321
23.4.3 The CD-ROM Mounts on the VM . . . 321
23.5 Canceling a File Upload . . . 322
23.6 Configuring vSphere for HTTP Connections . . . 322

Part VI Rebranding the Cloud Manager Interface 325

24 The Cloud Manager Branding Infrastructure 327

24.1 The Theming Directory Structure . . . 327
24.2 Creating a New Theme . . . 327
24.3 The Theming Configuration File . . . 328
24.4 Theming Search Order . . . 329
24.4.1 More About the Locale Folder . . . 329

25 Rebranding the Application Server Console Interface 331

25.1 Additional Theming Infrastructure . . . 331
25.2 Theming Resources for the Cloud Manager Service Reseller . . . 331
25.3 Theming the Remote Console Logo . . . 333

26 Rebranding Reports 335

26.1 File Locations and Contents for Report Theming . . . 335

27 Rebranding Email Messages 337

27.1 File Locations and Contents for Mail Message Theming . . . 337
27.1.1 The main_skin.msg file . . . 337
27.1.2 The /images Folder . . . 337
27.1.3 The /styles Folder . . . 337
27.1.4 The /tasks Folder . . . 338
27.1.5 The .msg and .subject Files . . . 338
27.2 Tokens . . . 338
27.3 Email Testing Tools . . . 341

28 Rebranding the Cloud Marketplace Mobile App **343**

Part VII Application Server Troubleshooting **345**

29 Troubleshooting Cloud Manager Application Issues **347**

29.1 Business Services Issues 347

30 Cloud Manager Application Server Logs **351**

About This Book and the Library

The *Administrator Reference* provides authoritative, specialized information related to the components of the Cloud Manager product.

Intended Audience

This document is intended for IT staff, such as IT service providers and operators, as well as business and technical users, who use Cloud Manager to build, manage, and consume cloud computing resources. A knowledge of VMware hypervisor technology is assumed.

Information in the Library

The library for this product is available in HTML and PDF formats on the Cloud Manager Documentation (https://www.netiq.com/documentation/cloud-manager) website.

The Cloud Manager library provides the following information resources:

Release Notes

Provides information about new features and enhancements in the release, as well as any known issues.

Product Overview

Provides information about the features, functionality, and operational concepts of Cloud Manager.

Installation and Upgrade Guide

Provides detailed planning and installation information for the software, as well as upgrade information.

Procedures Guide

Provides conceptual information, an overview of the user interface, and step-by-step guidance for common tasks.

Administrator Reference

Provides detailed reference information about tools and interfaces used by this product.

Help

Provides context-sensitive information and step-by-step guidance for common tasks as you work in the interface.

Feedback

If you have suggestions for documentation improvements, click **comment on this topic** at the bottom of any page in the HTML version of the documentation. You can also email Documentation-Feedback@netiq.com. We value your input and look forward to hearing from you.

Additional Resources

We encourage you to use the following additional resources online:

- Cloud Manager Forum (https://forums.netiq.com/forumdisplay.php?13-Cloud-Manager): A web-based community of product users where you can discuss product functionality and advice with other product users.
- Cloud Manager Product page (https://www.netiq.com/products/cloud-manager/): A web-based product brochure that provides information about features, how to buy, technical specifications, frequently asked questions, and a variety of resources such as videos and white papers.
- NetIQ User Community (https://www.netiq.com/communities/): A web-based community with a variety of discussion topics.
- NetIQ Support Knowledgebase (https://www.netiq.com/support/kb/): A collection of in-depth technical articles.
- NetIQ Support Forums (https://forums.netiq.com/forum.php): A web location where product users can discuss NetIQ product functionality and advice with other product users.
- MyNetIQ (https://www.netiq.com/f/mynetiq/): A website offering product information and services, such as access to premium white papers, webcast registrations, and product trial downloads.

Technical Support

For specific product issues, please contact our Technical Support team at www.netiq.com/support.

Orchestration Server Components

This section provides information about the Cloud Manager Orchestration Server components.

- Chapter 1, “Interface Layout of the Orchestration Server Console,” on page 15
- Chapter 2, “The Orchestration Server and the Server Admin Objects,” on page 17
- Chapter 3, “The Job Object,” on page 31
- Chapter 4, “The Resource Object,” on page 45
- Chapter 5, “The VM Host Object,” on page 75
- Chapter 6, “The Virtual Disk Object,” on page 89
- Chapter 7, “The Virtual NIC Object,” on page 107
- Chapter 8, “The Network Group and its Virtual Bridge Objects,” on page 115
- Chapter 9, “The Repository Object,” on page 121
- Chapter 10, “The User Object,” on page 135
- Chapter 11, “Miscellaneous Objects Displayed in the Explorer Tree,” on page 143
- Chapter 12, “The Orchestration Server Job Scheduler,” on page 145
- Chapter 13, “The Policy Debugger,” on page 161
- Chapter 14, “Grid Object Health Monitoring,” on page 169
- Chapter 15, “Events,” on page 177
- Chapter 16, “The Metrics Facility,” on page 185

1 Interface Layout of the Orchestration Server Console

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

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

Figure 1-1 *The Cloud Manager Orchestration Console*

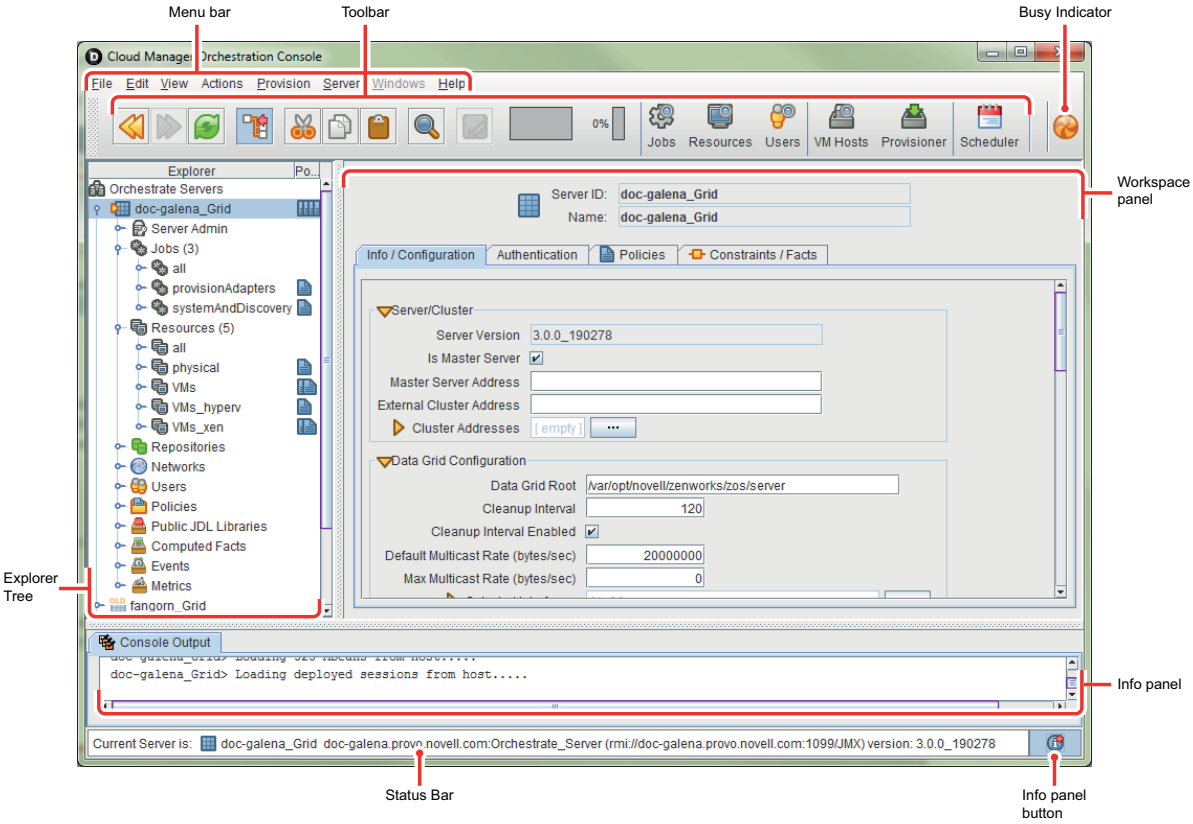

The following list describes the functional areas of the main Orchestration Console display.

- **Operations Menu bar:** The Operations Menu bar provides operations categorized under menus such as **File**, **Edit**, **View**, **Grid**, **Server**, **Windows**, and **Help**.
  - The **File** menu lets you save any changes you have made or exit the console.
  - The **Edit** menu lets you cut, copy, and paste items and choose general and server preferences for the console.
  - The **View** menu lets you manipulate the display of the different components of the console and refresh the Explorer and Workspace panels.

- The **Actions** menu lets you launch specific tools that create and delete users or user groups, computing resources, jobs, policies, and computed facts.
- The **Provision** menu lets you discover VM hosts and Repositories, VM Images, disks, and Repository capacities. you can also start and stop VMs and VM hosts, resync a VM's state or a VM hosts' state, reset the state of all VMs in the grid or cancel that resync, or can perform a rediscovery of all VMs in the grid.
- The **Server** menu lets you start a local server, log in to the server, create and display logs for logged-in servers, log out of the server, and shut down the server.
- The **Windows** menu lets you select console windows to display when you have more than one console window open. You can open the Explorer panel and the two tabs of the Info panel (**<Orchestrator> Log** and **Console Output**) in their own windows by right-clicking the tab and choosing **Open in window** in the pop-up menu.
- The **Help** menu provides access to the About box for the console. It also provides a link to Cloud Manager documentation on the web.

- **Console toolbar:** The Console toolbar has buttons for executing common tasks. The basic tasks are **Go Back**, **Go Forward**, **Refresh the view**, **Hide or Show the Explorer Panel**, **Cut**, **Copy**, **Paste**, **Save changes in workspace view**, and **Open the Find Dialog**.

  The toolbar also includes buttons that open monitoring views for **Jobs**, **Resources**, and **Users**.

  To the far right of the toolbar, a pinwheel icon indicates when the console is busy.
- **Explorer tree:** Sometimes referred to as the "Tree view," this panel displays a hierarchical tree that lets you navigate to different objects; you can click items in the tree to see their details. For example, you can display computing resources for a selected grid. When you click **Computing Resources** in the tree, its details appear in the Workspace panel with a list of active computing resources. You can edit the **Computing Resource** attributes in the workspace panel.
- **Workspace panel:** This panel displays a detailed view for an item you select in the Explorer tree. For example, if you select a computing resource under physical in the Explorer tree, the Workspace panel view changes to and "Admin view" to show the details for that resource. You can edit the properties of Grid object in the views displayed in the Workspace panel.

  The views in the panel change as you select different monitoring tools (**Jobs**, **Resources**, **Users**, **VM Hosts**, **Provisioner**) or the **Scheduler** tool in the Main toolbar.
- **Info panel:** The Info panel displays a variety of information, such as validation and error messages, log files, and query results. You can display or hide the Info panel by clicking the **Info panel** icon in the Status bar.
- **Status bar:** The status bar displays general identity information about the Orchestration Server where you are logged in.

For information about launching the console and using it for the first time, see "Launching the Orchestration Console and Logging in to the Orchestration Server"in the *Cloud Manager Installation and Upgrade Guide*.

2 The Orchestration Server and the Server Admin Objects

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

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

- Section 2.1, “Orchestration Server Object,” on page 17
- Section 2.2, “Server Admin Object,” on page 30

2.1 Orchestration Server Object

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

The Orchestration Console is version aware. When the console launches or when server discovery is manually run, the console recognizes both current Orchestration Server installations and old installations of discovered servers and displays their icons accordingly. This visual cue helps you to recognize when older Orchestration Servers need to be upgraded.

Figure 2-1 *Current and “Old” Server Objects*

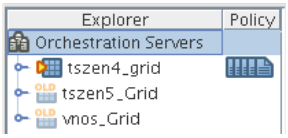

The tooltip for an Orchestration Server lists its RMI configuration, its IP address, the directory location where the server instance was installed, and its exact version number.

The icons to the right of a current Orchestration Server represent its policies, either those added by default upon server install and configuration, or those added later. A drop-down menu of all associated policies is opened when you right-click a policy icon. From there, you can select a policy to open in the Policy Editor. For more information about policies, see Section 11.1, “Policy Object,” on page 143.

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

- Section 2.1.1, “The Orchestration Server Info/Configuration Page,” on page 18
- Section 2.1.2, “Orchestration Server Authentication Page,” on page 26

- Section 2.1.3, “Orchestration Server Policies Page,” on page 30
- Section 2.1.4, “Orchestration Server Constraints/Facts Page,” on page 30

2.1.1 The Orchestration Server Info/Configuration Page

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

- “Server/Cluster Panel” on page 18
- “Data Grid Configuration Panel” on page 19
- “Security/TLS Configuration Panel” on page 20
- “Agent/User Session Configuration Panel” on page 22
- “Audit Database Configuration Panel” on page 22
- “Sentinel Server Configuration Panel” on page 23
- “Job Limits Panel” on page 25

Server/Cluster Panel

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

- Server Version
- Is Master Server
- Master Server Address
- External Cluster Address
- Cluster Addresses

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

Is Master Server: A check box that is selected if the server is the Master Server in a High Availability cluster configuration.

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

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

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

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

For more information about using Cloud Manager Orchestration components in a High Availability environment, see “Preparing the Cloud Manager Orchestration Server for SUSE High Availability Support” in the *Cloud Manager Installation and Upgrade Guide*.

Data Grid Configuration Panel

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

- Data Grid Root
- Cleanup Interval
- Cleanup Interval Enabled
- Default Multicast Rate
- Selected Interfaces
- Available Interfaces
- Total Packets Sent
- Total Packets Resent
- Total Resend Rate
- Current Packets Sent
- Current Packets Resent
- Current Resend Rate
- Current File Size
- Current Bytes Sent
- Current Percent Complete
- Skipped (Sparse) Bytes
- Current Receiver Count
- Current File Name

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

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

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

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

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

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

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

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

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

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

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

Security/TLS Configuration Panel

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

- TLS On Agent
- Forbid TLS for agents
- Allow TLS on the agents; default to falling back to unencrypted
- Allow TLS on the agents; default to TLS encrypted if not configured encrypted
- Make TLS mandatory on the agents
- TLS On Client
- Forbid TLS for clients
- Allow TLS on clients; default to falling back to unencrypted
- Allow TLS on agents; default to TLS encrypted if not configured encrypted
- Make TLS mandatory on the clients
- TLS Address

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

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

  This is the default setting for the Orchestration Server. More secure installations might require a setting to one of the higher levels below.
- **Allow TLS on the agents; default to TLS encrypted if not configured encrypted:** The server defaults to using encryption, but the agent can optionally disable encryption.
- **Make TLS mandatory on the agents:** The Orchestration Server rejects any connections that do not establish TLS encryption. This is the most secure encryption level because it ensures that all message communication between the node (that is, an agent) and the server is protected from tampering or interception.

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

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

  This is the default setting for the Orchestration Server. More secure installations might require a setting to one of the higher levels below.
- **Allow TLS on agents; default to TLS encrypted if not configured encrypted:** The server defaults to using encryption, but the user can optionally disable encryption.
- **Make TLS mandatory on the clients:** The Orchestration Server rejects any connections that do not establish TLS encryption. This is the most secure encryption level because it ensures that all message communication between the user's client programs and the server is protected from tampering or interception.

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

Agent/User Session Configuration Panel

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

- Agent Session Lifetime
- Agent Session Timeout
- Socket Keeps Agent Sessions Alive
- User Session Lifetime
- User Session Timeout
- Socket Keeps User Sessions Alive

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

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

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

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

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

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

Audit Database Configuration Panel

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

- JDBC Driver Name
- JDBC Library
- JDBC Connection URL
- Database Username
- Database Password
- Is Connected
- Connect

- Disconnect
- Clear Queue

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

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

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

Database Username: Specifies the user name for database authentication.

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

Is Connected: Indicates that the driver is successfully connected.

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

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

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

Sentinel Server Configuration Panel

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

For information about setting up a Sentinel Collector Server in your Orchestration Server environment, see "Integrating the Orchestration Server with a Sentinel Collector" in the *Cloud Manager Installation and Upgrade Guide*.

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

- Server Hostname
- Server Port Number
- Is Connected
- Log Channels
- Connect
- Disconnect
- Configure

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

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

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

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

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

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

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

For more information about Orchestration Server log levels, see "Orchestration Server Log Levels Mapped to Sentinel Log Levels" in the *Cloud Manager Procedures Guide*.

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

The following table shows the log channels you can choose and the Orchestration Server actions that prompt sending a log message through this channel.

Table 2-1 *Log Channels and the Occasions for Sending Messages Through Each*

Log Channel Name in the Orchestration Console (Sentinel Server Configuration Panel)	When Are Messages Sent to This Channel?
ActionStatusManager	◆ When the status of a Grid action is updated.
Audit	◆ When the Grid interacts with the audit database.
AuthLDAP **AuthZOS** **AuthenticationManager**	◆ On grid-wide authentication events.

Broker	◆ On job execution.
	  ◆ start
	  ◆ cancel

Event Manager **JobManager** **NodeManager** **UserManager** **repositoryManager** **vbridgeManager** **vdiskManager** **vnicManager**	◆ When a Grid object of the corresponding type is created, deleted, or its health changes to a bad state.
GroupManager	◆ When a member is added/removed in a Group.
JobScheduler	◆ When the job schedule or the job trigger deployment/undeployment.
MBeanServer	◆ When internal Grid Resources are updated.
PolicyManager	◆ On policy creation/deletion. ◆ On policy association/disassociation with any Grid object.
Sentinel	◆ When the Grid interacts with a Novell Sentinel server.
SessionManager	◆ On user or Resource login/logout.

Log Channel Name in the Orchestration Console (Sentinel Server Configuration Panel)	When Are Messages Sent to This Channel?
VmManager	◆ When actions are performed on VMs (provision, migrate, shutdown, clone, etc.). This could be initiated automatically or manually. ◆ When authorization fails during VM operation. ◆ When provisioning job fails.
computedFact	◆ When computed facts are created or updated or deleted.
deployer/computedFact **deployer/event** **deployer/facility** **deployer/jdlLibrary** **deployer/job** **deployer/library** **deployer/metric** **deployer/policy** **deployer/properties** **deployer/schedule** **deployer/service** **deployer/trigger** **deployer/xml**	◆ When a corresponding resource is deployed to or undeployed from the Grid.

Job Limits Panel

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

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

- max.active.jobs
- max.queued.jobs
- job.finishing.timeout

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

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

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

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

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

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

2.1.2 Orchestration Server Authentication Page

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

- "Resources Panel" on page 26
- "Users Panel" on page 26
- "Authentication Page" on page 28

Resources Panel

The resources in a Orchestration Server grid are actually Orchestration Agents that authenticate or "register" with the Orchestration Server.

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

Auto Upgrade Agents: This check box is already selected if you chose to enable the automatic upgrade of Orchestration Agents that communicate with your Orchestration Server. If you did not select this option during upgrade configuration, the check box is not selected.

If you select this check box, the associated Orchestration Agents are upgraded at intervals over a period of approximately five minutes. The upgrade happens without administrator approval.

If you deselect this check box at any time before or during the automatic upgrade, the upgrade process stops. Any agents that are not upgraded continue to be identified as "OLD" and are not usable with the newly upgraded server.

Users Panel

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

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

Enable LDAP Subpanel

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

- Enable LDAP
- Administrators
- Server Type
- Settings Subpanel
  - Directory Name
  - Servers

- Advanced
  - SSL
  - Start TLS
  - Query Account
  - Query Password
- Generic Settings Subpanel
  - Base Domain Name
  - User Attribute
  - User Filter
  - User Prefix
  - Group Attribute
  - Group Filter
  - Group Prefix
  - Group DNA Attribute
  - Nested DNA Attribute

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

Administrators: The Administrators list specifies the group names whose membership includes Orchestration Server administrators as returned by the specified authentication provider. You can add groups to this list by clicking the ... button to open an array editor dialog, which allows groups to be added, removed, and reordered. A group must be in the format `<provider>:<group|groupnocase>:<groupname>`, where the `<provider>` is either ZOS or LDAP. For example, adding `LDAP:groupnocase:XyZ` allows users reported by the LDAP server as members of a group xyz, or XYZ, xYz, etc. to authenticate as an administrator. To enforce to case-sensitive matching, use `LDAP:group:XyZ` instead. Non-case-sensitive matching is needed for Active Directory servers.

Server Type: This drop-down list lets you specify which authentication provider you want to use: **Active Directory Service** or **Generic LDAP Directory Service**.

- If you select **Active Directory Service**, specify the values in the **Settings** subpanel only.
- If you select **Generic LDAP Directory Service**, specify the values in the **Settings** subpanel (except **Advanced** settings) and the values in the **Generic Settings** subpanel.

Settings Subpanel: Set the values in this subpanel for the ADS authentication provider.

- **Directory Name:** The name of the Active Directory Service server.
- **Servers:** A list of strings containing `server:port` entries for a list of servers to be used.

  Each entry can be of one of three forms:

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

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

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

- **Advanced:** The settings in this subpanel are for more selective ADS authentication.
  - **SSL:** If the accompanying **Start TLS** check box is not selected and if the ADS server's SSL certificate has been installed on the Orchestration Server JVM, this option securely connects to the ADS server through SSL encryption.

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

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

- **Base Domain Name:** The Root DN of the LDAP server's directory tree. This must be obtained by the administrator, and is usually in the form of `dc=adsroot,dc=novell,dc=com`.
- **User Attribute:** The attribute on a user's entry that identifies his or her login account name. For ADS servers, this attribute is sAMAccountName.
- **User Filter:** The name of the filter to be used in the lookup for the user's LDAP distinguished name.

  For ADS, this prefix is `cn=Users`.
- **User Prefix:** The prefix used to define the LDAP subtree within the BaseDN tree that contains user accounts. If you leave this property blank, the Orchestration Server uses the BaseDN.
- **Group Attribute:** Specifies the attribute of a group entry describing the login name of that group.
- **Group Filter:** A filter to be used in the lookup for group memberships on some LDAP schemas. The filter can use either ${USER_NAME} or ${USER_DN} to substitute that value. For example: memberUid=${USER_NAME}.
- **Group Prefix:** The prefix used to define the LDAP subtree within the BaseDN tree that contains group accounts.

  This field is not used for Active Directory authentication.
- **Group DNA Attribute:** The directory root where all queries for a user's group memberships (stored as a list of "member of" attributes on the user's entry on an ADS server) are to occur.
- **Nested DNA Attribute:** The attribute of a group entry where subgroups can be queried.

Authentication Page

As a data center administrator, you often have to provide credentials and certificates as you interact with the different hypervisor technologies, such as the vSphere technologies. The Orchestration Server lets you store this data in a centralized, secure location in its Credential Manager. No clear text passwords are accessible.

NOTE: The Cloud Manager Orchestration Server uses Triple DES password-based encryption in its Credential Manager to encrypt stored credentials and certificates.

The Credential Manager requires information from the Authentication page of the Orchestration Server object in the Orchestration Console, The page includes the following sections:

- "Stored Credentials Panel" on page 29
- "Stored Certificates Panel" on page 29

Stored Credentials Panel

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

- **Name:** (Required) The name that you want to use to refer to this credential set.
- **User:** (Required) The user name with rights to administer objects in this grid.
- **Secret/Password:** (Required) The password that authenticates the user.
- **Type:** (Optional) A user-defined string that lets similar credentials be put into a category or group. For example, you might have a "type" of credential for the amazon-ec2 provisioning adapter and another type for the vsphere provisioning adapter.

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

This password is used to encrypt the stored passwords. By default the password is `CHANGE_THIS_PASSPHRASE`. We recommend that you select a new password to use for encrypting stored passwords.

Stored Certificates Panel

In order to trust certificates not signed by well-known certificate authorities, the Orchestration Server lets you store certificates that are trusted by Java.

NOTE: Public/Private key pairs can be stored as certificates. This is useful if you need to manage Amazon EC2 key pairs.

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

Currently, this functionality is used only by the Orchestration vsphere provisioning adapter.

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

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

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

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

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

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

or as

```
your.server.name
```

or as

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

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

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

2.1.3 Orchestration Server Policies Page

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

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

2.1.4 Orchestration Server Constraints/Facts Page

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

2.2 Server Admin Object

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

3 The Job Object

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

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

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

A Job object represents an individual job in the Explorer tree of the Orchestration Console. This object contains facts with attributes that are used for job and joblet control. Policies associated with the job also control the job. The Orchestration Console has an administrative (“admin”) view in the Explorer Panel that lets you edit these objects.

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

- Section 3.1, “Job Groups,” on page 31
- Section 3.2, “The Job Info/Groups Page,” on page 32
- Section 3.3, “The Job Configuration Page,” on page 41
- Section 3.4, “The JDL Editor Page,” on page 42
- Section 3.5, “The Job Library Editor Page,” on page 42
- Section 3.6, “The Job Policies Page,” on page 43
- Section 3.7, “The Job Constraints/Facts Page,” on page 44

3.1 Job Groups

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

3.2 The Job Info/Groups Page

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

- Section 3.2.1, “Info,” on page 32
- Section 3.2.2, “Groups,” on page 41

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

3.2.1 Info

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

- “Show Inherited Fact Values Check Box” on page 32
- “Job Control Settings” on page 32
- “Joblet Control Settings” on page 36
- “Automatic Resource Provisioning Settings” on page 37
- “Resource Preemption Settings” on page 38
- “Job Counts” on page 39
- “Job History” on page 39

Show Inherited Fact Values Check Box

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

Job Control Settings

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

The job type options (completion time algorithms) include:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

The queue type options (start time algorithms) include:

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

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

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

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

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

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

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

In the Fact Editor, this fact is listed as `jopb.cacheresourcematches.ttl`:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Joblet Control Settings

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Automatic Resource Provisioning Settings

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Resource Preemption Settings

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

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

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

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

Job Counts

Total Instances: Displays the total number of job instances of this type that exist in the Cloud Manager Orchestration Server system.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Job History

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

3.2.2 Groups

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

3.3 The Job Configuration Page

The **Job Configuration** tab of the Job admin view opens the Job Configuration page, which includes the LVM Discovery Settings panel. On this panel, you can configure settings that control the discovery of logical volume managers. The **Volume Group Patterns** fact (`job.volume_group_patterns`) provides a place for you to define regular expression patterns for Volume Group names. The Orchestration Server uses these patterns to create Repository objects.

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

```
<fact name="job.volume_group_patterns" description="Regular expression patterns
for Volume Group names to create a Repository Object for">
  <array>
    <string>NCM-*</string>
  </array>
</fact>
```

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

The admin view also includes a link to open and edit the configuration definition file.

3.4 The JDL Editor Page

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

Figure 3-1 *The JDL Editor*

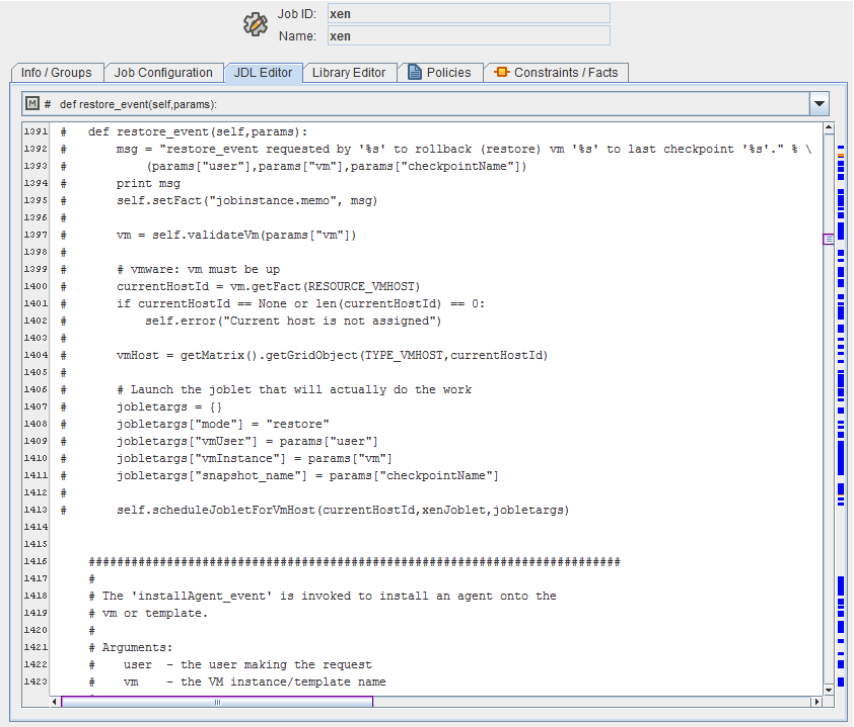

A drop-down list at the top of the editor includes the Java classes and their methods that are bookmarked in the code. Select any of these to go to the location in the code where they are invoked. Clickable colored blocks on the editor scroll bar perform a similar bookmarking function.

3.5 The Job Library Editor Page

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

Figure 3-2 *The Job Library Editor*

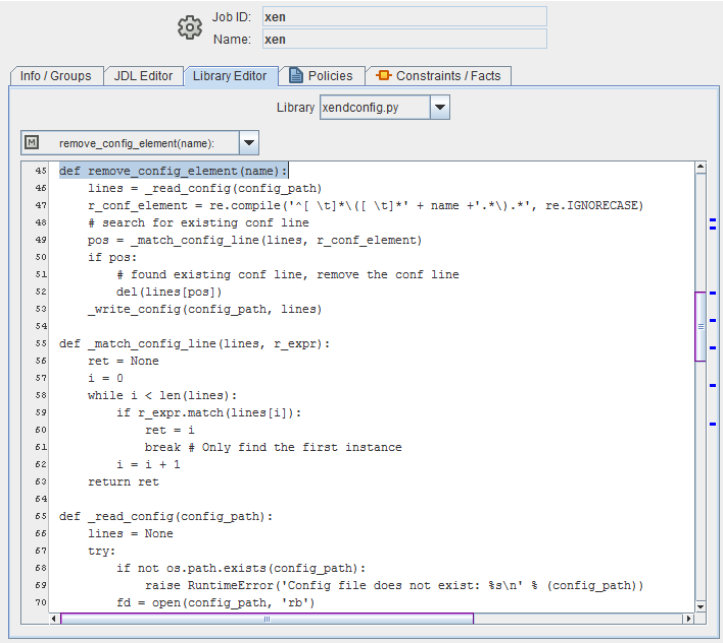

There are two drop-down lists located at the top of the Library Editor view. The first list displays the different libraries for the job, and the second list displays the methods that are bookmarked in the code. Select a method in the second drop-down list to go to the location in the library code where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

3.6 The Job Policies Page

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

You can modify a policy by using the Policy Grid object. For more information see Section 11.1, “Policy Object,” on page 143.

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

Figure 3-3 *The Policy Selection dialog*

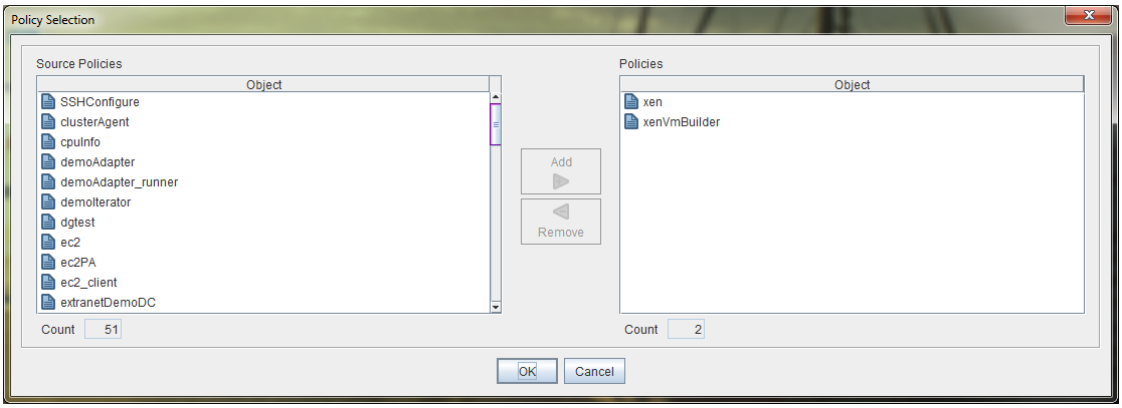

3.7 The Job Constraints/Facts Page

The **Constraints/Facts** tab opens a page that shows all of the effective constraints and facts for a Grid object. Each Grid object has an associated set of facts and constraints that define its properties. By changing the policy constraints and fact values for a job, you can change the behavior of the job and how the Orchestration Server allocates available system resources to it. The Orchestration Server assigns default values to each of the component facts. Facts with no mode shown can be changed at any time by the administrator. Facts with mode `r/o` have read-only values, which can be viewed by using the pencil icon, but changes cannot be made.

4 The Resource Object

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

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

- Section 4.1, “Resource Groups,” on page 45
- Section 4.2, “Resource Info/Groups Page,” on page 45
- Section 4.3, “Provision Info Page,” on page 71
- Section 4.4, “Resource Log Page,” on page 72
- Section 4.5, “Resource Policies Page,” on page 72
- Section 4.6, “Resource Health Debugger Page,” on page 72
- Section 4.7, “Resource Constraints/Facts Page,” on page 72
- Section 4.8, “Resource Object Naming and Renaming,” on page 73

4.1 Resource Groups

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

4.2 Resource Info/Groups Page

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

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

- Section 4.2.1, “Info Panel,” on page 46
- Section 4.2.2, “Groups Panel,” on page 71

4.2.1 Info Panel

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

- “Show Inherited Fact Values Check Box” on page 46
- “Resource Information” on page 46
- “VM Host Info” on page 50
- “Virtual Machine Configuration” on page 51
- “Provisioning Information” on page 55
- “OS Information” on page 64
- “CPU Information” on page 66
- “Memory Information” on page 67
- “Disk/Network Information” on page 68
- “Agent Information” on page 68
- “Agent Configuration” on page 69
- “Installed Components” on page 71

Show Inherited Fact Values Check Box

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

Resource Information

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

- Resource Type
- Resource Enabled
- Healthy
- Shutting Down
- Host Name
- Host Fully Qualified Name
- Password
- Host IP Address
- VNC IP Address
- VNC Port
- Billing Rate
- Bill For
- Power Factor
- Load Average
- CPU Load
- Joblet Slots
- Extra System Joblet Slots
- Joblets Active

- Became Idle On
- Total Joblets Started
- Total Completed Joblets
- Total Cancelled Joblets
- Total Failed Joblets
- Total Charge
- Total Wall Time
- Total Grid Time
- Sessions
- Provisionable Resource

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

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

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

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

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

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

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

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

Healthy: When this check box is selected (it has a value of true), the resource is considered to be in good health. You can set the health of the object by selecting or deselecting the health check box. Changing the value in this way has an immediate effect unless the value is overriden by an attached policy. For more information, see Chapter 14, “Grid Object Health Monitoring,” on page 169.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

```
<fact name="resource.becameidle" value="7/23/11 5:02 PM" type="Date" />
```

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

VM Host Info

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

- VM Host Containers
- VM Host Repositories

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

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

```
<fact name="resource.vmhosts">
  <array>
    <string>host1slesx_vsphere</string>
  </array>
</fact>
```

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

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

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

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

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

Virtual Machine Configuration

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

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

For information about VM provisioning based on options exposed in the Orchestration Console, see "Using the Right-Click Menu for Provisioning Actions" in *Cloud Manager Procedures Guide*.

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

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

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

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

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

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

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

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

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

VM Version: This fact is no longer used.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

```
<fact name="resource.vm.files">
  <dictionary>
    <dictelement key="config">
      <string>/vmfiles/vms/vpshere-vm123</string>
    </dictelement>
  </dictionary>
</fact>
```

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

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

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

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

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

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

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

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

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

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

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

Possible types include:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Construction Specification: The **VM Builder Specifications** field in this section displays a list of specifications that were used to build this VM. These specifications are interpreted by the provisioning adapter. .

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

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

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

Provisioning Information

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

- Provisioning Job
- Provisioned Instances
- Cloned Instances
- Instances
- Max Provisioned Instances
- Agent Shutdown Timeout
- Default Agent Idle Timeout
- Host Wait Timeout
- Preferred Host Wait
- Recommended Host
- Debug Provision Log
- Parent Template
- Current State
- Current Host
- Current Status

- Current Action
- Request Time
- Start Time
- Shutdown Time
- Host Wait Time
- Managing Job ID
- Automatic Provision
- Needs Resync

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

- Linux Autoprep Config
- Windows Sysprep Config

Settings

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

```
<fact name="resource.provision.vmhost" value="vmh6sles_vsphere" type="String" />
```

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

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

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

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

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

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

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

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

```
<fact name="resource.provision.time.request" value="8/24/11 4:36 PM" type="Date" /
>
```

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Linux Autoprep Config

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

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

This section includes the following settings:

- Linux Computer Name
- Linux Domain
- Linux Root Password

The section also includes a subpanel where you can set Network Autoprep Configuration information.

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

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

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

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

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

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

Linux Root Password: This value specifies the root password for the VM administrator.

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

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

Network Autoprep Config

This section includes the following settings:

- DNS Server IP Addresses
- DNS Suffixes
- Gateway IP Addresses

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

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

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

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

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

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

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

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

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

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

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

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

Windows Sysprep Config

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

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

The section includes the following settings/facts:

- Change SID
- Delete Accounts
- Admin Password
- Admin Password Plaintext
- Timezone
- Autologon
- Autologon Count
- Fullname
- Org Name

- Computer Name
- Product ID
- Run Once Command
- Workgroup
- Domain
- Domain Admin
- Domain Admin Password
- Domain Admin Password Plaintext
- Machine Object OU
- Machine Password
- Machine Password Plaintext
- License File Automode
- License File Autousers

The section also includes a subpanel where you can set Network Sysprep Configuration information.

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

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

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

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

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

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

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

In the Fact Editor, this fact is listed as `resource.provisioner.autoprep.sysprep.GuiUnattended.AdminPassword.value`:

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

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

In the Fact Editor, this fact is listed as `resource.provisioner.autoprep.sysprep.GuiUnattended.AdminPassword.plainText`:

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

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

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

In the Fact Editor, this fact is listed as `resource.provisioner.autoprep.sysprep.GuiUnattended.TimeZone`:

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

Autologon: If the value is true, the VM automatically logs into the Administrator account using Admin Password. If it is false, logon is prompted.

In the Fact Editor, this fact is listed as `resource.provisioner.autoprep.sysprep.GuiUnattended.AutoLogon`:

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

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

In the Fact Editor, this fact is listed as `resource.provisioner.autoprep.sysprep.GuiUnattended.AutoLogonCount`:

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

Fullname: The user's full name.

In the Fact Editor, this fact is listed as `resource.provisioner.autoprep.sysprep.UserData.FullName`:

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

Org Name: The organization name.

In the Fact Editor, this fact is listed as `resource.provisioner.autoprep.sysprep.UserData.OrgName`:

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

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

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

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

Product ID: The Windows product key.

In the Fact Editor, this fact is listed as `resource.provisioner.autoprep.sysprep.UserData.ProductID`:

```
<fact name="resource.provisioner.autoprep.sysprep.UserData.ProductID"
value="jklaieuqa4354" type="String" />
```

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

In the Fact Editor, this fact is listed as `resource.provisioner.autoprep.sysprep.GuiRunOnce.Command`:

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

Workgroup: Windows workgroup name. If the VM is joining a domain, use `JoinDomain`.

In the Fact Editor, this fact is listed as `resource.provisioner.autoprep.sysprep.Identification.JoinWorkgroup`:

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

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

In the Fact Editor, this fact is listed as `resource.provisioner.autoprep.sysprep.Identification.JoinDomain`:

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

Domain Admin: The Windows domain administrator name.

In the Fact Editor, this fact is listed as `resource.provisioner.autoprep.sysprep.Identification.DomainAdmin`:

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

Domain Admin Password: The Windows domain administrator account password.

In the Fact Editor, this fact is listed as `resource.provisioner.autoprep.sysprep.Identification.DomainAdminPassword.value`:

```
<fact
name="resource.provisioner.autoprep.sysprep.Identification.DomainAdminPassword.val
ue" value="cleanwindow" type="String" />
```

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

In the Fact Editor, this fact is listed as `resource.provisioner.autoprep.sysprep.Identification.DomainAdminPassword.plainText`:

```
<fact
name="resource.provisioner.autoprep.sysprep.Identification.DomainAdminPassword.pla
inText" value="false" type="Boolean" />
```

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

In the Fact Editor, this fact is listed as `resource.provisioner.autoprep.sysprep.Identification.MachineObjectOU`

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

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

In the Fact Editor, this fact is listed as `resource.provisioner.autoprep.sysprep.Identification.MachinePassword.value`

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

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

In the Fact Editor, this fact is listed as `resource.provisioner.autoprep.sysprep.Identification.MachinePassword.plainText`

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

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

In the Fact Editor, this fact is listed as `resource.provisioner.autoprep.sysprep.LicenseFilePrintData.AutoMode`:

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

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

In the Fact Editor, this fact is listed as `resource.provisioner.autoprep.sysprep.LicenseFilePrintData.AutoUsers`:

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

Network Sysprep Config

This section includes the following settings:

- DNS Suffixes

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

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

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

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

OS Information

This section includes the following settings:

- OS Name
- OS Version
- OS Version String
- OS Architecture
- OS Family
- OS Type
- OS Vendor

- OS Vendor Version
- OS Vendor String
- OS File Path Separator

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

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

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

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

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

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

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

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

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

OS Architecture: The operating system architecture (for example, **x86**, **amd64**, **i386**, or **sparc**).

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

```
<fact name="resource.os.arch" value="i386" type="String" />
```

OS Family: The operating system family name (for example, **windows**, **linux**, **solaris**, **unix**, **aix**, **mac**) of the resource, if known.

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

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

OS Type: This drop-down list lets you select the unique string identifier for each OS release, for example, **sles11**.

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

```
<fact name="resource.os.type" value="sles10" type="String" />
```

OS Vendor: (Read Only) The operating system vendor. SuSE for SUSE Linux Enterprise Server or SUSE Linux Enterprise Desktop.

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

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

OS Vendor Version: (Read Only) This field displays the vendor-defined version for the operating system. Ffor example, **11** for SUSE Linux Enterprise Server 11.

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

```
<fact name="resource.os.vendor.version" value="10" type="String" />
```

OS Vendor String: (Read Only) This field displays the full identification for the operating system that is supplied by the vendor. The osinfo system job must run for this value to be displayed.

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

```
<fact name="resource.os.vendor.string" value="Welcome to SUSE Linux Enterprise
Server 11 (i586) - Kernel   (\l)." type="String" />
```

OS File Path Separator: (Read Only) The resource operating system file separator.

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

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

CPU Information

This section includes the following settings:

- Number of CPUs
- CPU Speed (Mhz)
- CPU Vendor
- CPU Model
- CPU Architecture
- CPU HVM Support

Number of CPUs: (Read only) The number of CPUs available for this resource to use. For a VM resource, this fact remains undefined.

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

```
<fact name="resource.cpu.number" value="2" type="Integer" />
```

CPU Speed (Mhz): (Read only) The processor speed measured in Mhz. The cpuinfo job must run for this value to be displayed. For a VM resource, this fact remains undefined.

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

```
<fact name="resource.cpu.mhz" value="2594" type="Integer" />
```

CPU Vendor: (Read only) The name of the CPU vendor. The cpuinfo system job must run for this value to be displayed. For a VM resource, this fact remains undefined.

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

```
<fact name="resource.cpu.vendor" value="GenuineIntel" type="String" />
```

CPU Model: (Read only) The full vendor model number of the CPU. The cpuinfo system job must run for this value to be displayed. For a VM resource, this fact remains undefined.

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

```
<fact name="resource.cpu.model" value="Intel(R) Pentium(R) 4 CPU 2.60GHz"
type="String" />
```

CPU Architecture: The CPU architecture (for example, **x86**, **x86_64**, **sparc**) of this resource. For a VM resource, this fact remains undefined.

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

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

CPU HVM Support: This field is marked true if the CPU has hardware virtualization support.

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

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

Memory Information

This section includes the following settings:

- Virtual Memory (Mb)
- Virtual Available
- Physical Memory (Mb)
- Physical Available
- Swap Memory (Mb)
- Swap Available

Virtual Memory (Mb): (Read only) The total amount of virtual memory, measured in MB, on the resource. The memInfo system job must run for this value to be displayed

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

```
<fact name="resource.memory.virtual.total" value="4060" type="Integer" />
```

Virtual Available: (Read only) The amount of available virtual memory, measured in MB, on the resource. The memInfo system job must run for this value to be displayed.

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

```
<fact name="resource.memory.virtual.available" value="19951" type="Integer" />
```

Physical Memory (Mb): (Read only) The total amount of physical memory, measured in MB, on the resource. The memInfo system job must run for this value to be displayed

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

```
<fact name="resource.memory.physical.total" value="3889" type="Integer" />
```

Physical Available: (Read only) The amount of available physical memory, measured in MB, on the resource. The memInfo system job must run for this value to be displayed.

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

```
<fact name="resource.memory.physical.available" value="3565" type="Integer" />
```

Swap Memory (Mb): (Read only) The total amount of configured swap space, measured in MB, on the resource. The memInfo system job must run for this value to be displayed.

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

```
<fact name="resource.memory.swap.total" value="16386" type="Integer" />
```

Swap Available: (Read only) The total amount of free swap space, measured in MB, on the resource. The memInfo system job must run for this value to be displayed.

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

```
<fact name="resource.memory.swap.available" value="16386" type="Integer" />
```

Disk/Network Information

The facts in the **Disk/Network Information** section of the Info/Groups page are not currently functional and are not supported.

Agent Information

This section includes the following settings:

- Agent Version
- Agent Install Dir
- Agent Java Version
- Agent Java Runtime
- Agent Java Vendor
- Agent Java Home Dir
- Available Agent Memory
- Enhanced Exec Available
- Clustered Agent

Agent Version: (Read only) The Orchestration Agent version and build number that is installed on this resource. The string uses the following syntax:

major.minor.point_build

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

```
<fact name="resource.agent.version" value="2.0.2_70917" type="String" />
```

Agent Install Dir: (Read only) The name of the home directory of the Orchestration Agent installation files.

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

```
<fact name="resource.agent.home" value="/opt/novell/zenworks/zos/agent"
type="String" />
```

Agent Java Version: (Read only) The version of the Java JVM currently in use by the Orchestration Agent installed on this resource.

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

```
<fact name="resource.agent.jvm.version" value="1.5.0_17" type="String" />
```

Agent Java Runtime: (Read only) The version of the Java JVM runtime currently in use by the Orchestration Agent installed on this resource.

In the Fact Editor, this fact is listed as `resource.agent.jvm.runtime`:

```
<fact name="resource.agent.jvm.runtime" value="1.5.0_17-b04" type="String" />
```

Agent Java Vendor: (Read only) The name of the vendor of the Java JVM currently in use by the Orchestration Agent installed on this resource.

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

```
<fact name="resource.agent.jvm.vendor" value="Sun Microsystems Inc." type="String"
/>
```

Agent Java Home Dir: (Read only) The path to the home directory of the Java JVM currently in use by the Orchestration Agent installed on this resource.

In the Fact Editor, this fact is listed as `resource.agent.jvm.home`:

```
<fact name="resource.agent.jvm.home" value="/opt/novell/zenworks/zos/agent/jre"
type="String" />
```

Available Agent Memory: (Read only) The amount of memory, measured in MB, available to the Orchestration Agent installed on this resource.

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

```
<fact name="resource.agent.jvm.memory" value="127" type="Integer" />
```

Enhanced Exec Available: This check box is selected by default. When it is selected (it has a value of true), it indicates that the Orchestration Agent installed on this resource is able to use enhanced exec features, as opposed to unsupported agent installs, such as AIX.

In the Fact Editor, this fact is listed as `resource.agent.exec.installed`:

```
<fact name="resource.agent.exec.installed" value="true" type="Boolean" />
```

Clustered Agent: This check box is not selected by default. When you select it (it has a value of true), you specify that the agent is “clustered” on this VM resource. This means that it converts duplicate logins to failover logins.

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

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

Agent Configuration

This section includes the following settings:

- Gmond Port
- Datagrid Cache TTL
- Datagrid Cleanup Interval
- Exec Daemon Timeout
- Exec As Agent User Only
- Cleanup After Joblets
- Use Enhanced Exec
- Log Level
- Debug Logging

Gmond Port: The port that the agent uses for gmond. Port 8649 is the default port. A setting of zero (0) or less means that the value is not read.

In the Fact Editor, this fact is listed as `esource.agent.config.gmond.port`:

```
<fact name="resource.agent.config.gmond.port" value="8649" type="Integer" />
```

Datagrid Cache TTL: The amount of time (measured in minutes) that inactive files should remain in the agent’s datagrid cache. A setting of zero (0) turns off the cache.

In the Fact Editor, this fact is listed as `jresource.agent.config.datagrid.cache.lifetime`:

```
<fact name="resource.agent.config.datagrid.cache.lifetime" value="1440"
type="Integer" />
```

Datagrid Cleanup Interval: The amount of time (measured in minutes) that the Orchestration Server should wait between cleanup sweeps of the agent's datagrid cache.

In the Fact Editor, this fact is listed as `resource.agent.config.datagrid.cache.cleanupinterval`:

```
<fact name="resource.agent.config.datagrid.cache.cleanupinterval" value="60"
type="Integer" />
```

Exec Daemon Timeout: The amount of time (measured in seconds) that the enhanced exec daemon is to remain running. A setting of zero (0) specifies that the daemon is to remain running. The exec daemon is the non-Java component of the agent that is responsible for executing commands remotely.

In the Fact Editor, this fact is listed as `resource.agent.config.exec.daemon.timeout`:

```
<fact name="resource.agent.config.exec.daemon.timeout" value="300" type="Integer"
/>
```

Exec As Agent User Only: This check box is selected by default. When you select it (it has a value of true), you specify that the agent is to always run executables as the Agent User only. Selecting this check box overrides any job fact settings for the `job.joblet.runtype` fact.

In the Fact Editor, this fact is listed as `resource.agent.config.exec.asagentuseronly`:

```
<fact name="resource.agent.config.exec.asagentuseronly" value="true"
type="Boolean" />
```

Cleanup After Joblets: This check box is not selected by default. When you select it (it has a value of true), you specify that the agent on this resource is to clean up temporary directories created for each joblet. You can deselect this check box for debugging purposes; when you select it again, the cleanup process starts again, deleting temporary directories that were created while the setting was deactivated.

In the Fact Editor, this fact is listed as `resource.agent.config.joblet.cleanup`:

```
<fact name="resource.agent.config.joblet.cleanup" value="true" type="Boolean" />
```

Use Enhanced Exec: This check box is not selected by default. When you select it (it has a value of true), you specify that the agent on this resource is to use the enhanced exec feature of the agent, which is available for supported agent installations. Marking this fact as false causes the enhanced exec feature to not be used.

In the Fact Editor, this fact is listed as `resource.agent.config.exec.enhancedused`:

```
<fact name="resource.agent.config.exec.enhancedused" value="true" type="Boolean" /
>
```

Log Level: Lets you choose the level of agent logging in terms of the amount of detail (that is, the verbosity) you want to include in the agent log. The choices include:

- quiet
- normal
- verbose

In the Fact Editor, this fact is listed as `resource.agent.config.loglevel`:

```
<fact name="resource.agent.config.loglevel" value="normal" type="String" />
```

Debug Logging: This check box is not selected by default. When you select it (it has a value of true), you activate the debug function in the agent log, which is additive to the log level.

In the Fact Editor, this fact is listed as `resource.agent.config.logdebug`:

```
<fact name="resource.agent.config.logdebug" value="false" type="Boolean" />
```

Installed Components

Applications: A list of the names of applications (including the full version name) that are installed on this resource. This is useful for constraining joblets to run only on a resource with a particular application installed.

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

```
<fact name="resource.installed.apps">
  <array>
    <string>man-pages-2.39-0.9</string>
    <string>xorg-x11-fonts-scalable-6.9.0-50.45</string>
    <string>cifs-mount-3.0.24-2.23</string>
    <string>gdbm-1.8.3-243.2</string>
    <string>libaio-0.3.104-14.2</string>
    <string>libnl-1.0-18.4</string>
...
  </array>
</fact>
```

You can edit this array by clicking the ··· button to open an array editor, where you can add or remove the application name or change its order in the array of element choices.

4.2.2 Groups Panel

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

4.3 Provision Info Page

The **Provision Info** tab is displayed only for VM resource objects selected in the Explorer tree. The read-only fields displayed at the top of the Provision Info page summarize information related to the provisioning of the resource, whether that resource is a VM or VM template.

The page has several subtabs that open other pages that display further information about the VM:

- **Show Log:** A log that includes historical details about the history of the provisioning of the VM.
- **Host Assignment Log:** A log of VM host assignment constraint errors of the last migration or provision action.
- **Autoprep Data:** A list of autoprep/sysprep fact overrides used to prepare this VM instance. For more information about autoprep, see “Configuring Sysprep or Autoprep” in the *Cloud Manager Installation and Upgrade Guide*.
- **Policy Debugger:** Information to help you debug the policy and its resulting constraints and facts that are associated with this VM. For more information, see Chapter 13, “The Policy Debugger,” on page 161.

- **Action History:** A log that includes historical details about the history of the provisioning of the VM. For more information, see “Provisioning Actions and History” in the *Cloud Manager Procedures Guide*.

4.4 Resource Log Page

Open the **Resource Log** tab to view the contents of the log file for this resource. You can click **Refresh** to update the content. You can also select **Debug Logging** to activate the debug feature as part of the logging or change the **Log Level** to the level of detail that you want.

4.5 Resource Policies Page

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

You can modify a policy by using the Policy Grid object. For more information, see Section 11.1, “Policy Object,” on page 143.

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

Figure 4-1 *The Policy Selection dialog*

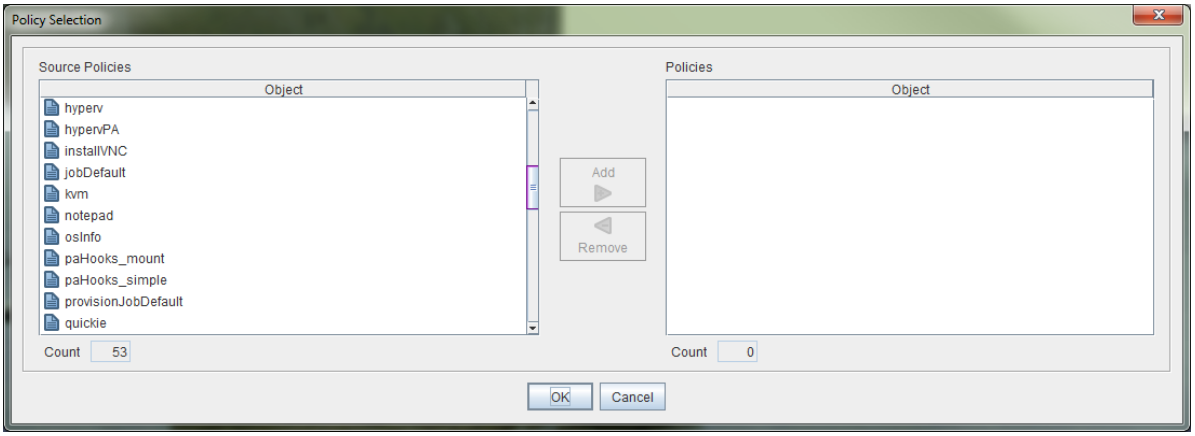

4.6 Resource Health Debugger Page

The Health Debugger is a common Admin view in the Orchestration Console for most Grid objects. For information about this tool, see Health Debugger in Chapter 14, “Grid Object Health Monitoring,” on page 169.

4.7 Resource Constraints/Facts Page

The **Constraints/Facts** tab opens a page that shows all of the effective constraints and facts for a Grid object. Each Grid object has an associated set of facts and constraints that define its properties. By changing the policy constraints and fact values for a job, you can change the behavior of the job and how the Orchestration Server allocates available system resources to it. The Orchestration Server

assigns default values to each of the component facts. Facts with no mode specified can be changed at any time by the administrator.Facts with mode `r/o` have read-only values, which can be viewed by using the edit pencil icon, but changes cannot be made.

For information about renaming Resource objects using the Fact Editor on this page, see Section 4.8, “Resource Object Naming and Renaming,” on page 73.

4.8 Resource Object Naming and Renaming

The Orchestration Server Resource Grid object type can include resources of various types, including physical machines, virtual machines (VMs), and VM templates, all of which are modeled differently in the Orchestration Console because of their varying roles in the Orchestration Server system. Some resource names are generated by the Orchestration Server system and can therefore receive generic, arbitrary names such as `mysql-1`, `mysql-2`, and so on. Resources you name at installation time or creation time might also change in their purpose or facilities.

As the quantity of these Resource objects grows in your grid, you might find it helpful or necessary to rename them, assigning more meaningful, intuitive names to suit the purpose of the object. The object’s “display name” is visible in the Orchestration Console interface, interfaces, and in optional zos and zosadmin commands.

NOTE: Resource object groups (that is, the folders that contain these Resource objects) can also be renamed. Objects such as jobs, events, and users cannot be renamed.

A Resource object’s name is stored in the `${objectType}.displayname` fact, which exists on every Grid object type, even those objects that cannot be renamed.

You can rename a Resource object in the Orchestration Console by using one of three methods:

- Right-click the Resource object in the Explorer tree, then select **Rename** to allow editing of the display name.
- Triple-click the Resource object in the Explorer tree to allow editing of the display name.
- In the Constraints/Facts page, select the resource object `.displayname` fact and then open the Fact Editor to enter a new value for that fact.

As you use one of these methods, you will notice that the fact value is pre-populated with the `${objectType}.id` fact. This functions as the name value for the object name until you decide to change it.

NOTE: Even after being renamed, the Resource object retains its associated resource ID in the `.id` fact. This is not editable.

For more information about making the Resource object display names visible from the zos or zosadmin command line, see Chapter 17, “Orchestration Server zos Command Line Tool,” on page 195 and Chapter 18, “Orchestration Server zosadmin Command Line Tool,” on page 243.

5 The VM Host Object

A VM host represents a VM host technology or hypervisor either installed on a physical resource or accessed by it (in the case of VMware). VM host objects can be used when making provisioning decisions for a resource.

The Cloud Manager Orchestration Server also supports the discovery of VMware vSphere clusters used for high availability in a VMware environment or managed by the VMware Distributed Resource Scheduler (DRS) after an Orchestration Agent has been deployed into such an environment. In this scenario, the Orchestration Server also allows you to determine when actions have taken place outside of the Orchestration environment, such as when the DRS moves a VM to an alternate host in the cluster or when an administrator moves a VM into a different resource pool (see "Setting Up Orchestration to Accommodate VMware DRS Clustering and Updates" in the *Cloud Manager Installation and Upgrade Guide*).

Although the VM host and the VM host Cluster are regarded as two different types of VM host object, and have differing icons, the discovered clusters are represented in the Explorer tree of the Orchestration Console as VM host objects.

NOTE: The Orchestration Console interface (that is, the fields in the admin view) for a VM host and a VM host Cluster are nearly identical. Facts unique to the VM host Cluster are listed in Section 5.7, "Unique VM Host Cluster Facts," on page 82.

This section includes the following information:

- Section 5.1, "Info Page," on page 75
- Section 5.2, "Policies Page," on page 80
- Section 5.3, "Health Debugger Page," on page 80
- Section 5.4, "Constraints/Facts Page," on page 81
- Section 5.5, "Action History Page," on page 81
- Section 5.6, "VM Host Object Naming and Renaming," on page 81
- Section 5.7, "Unique VM Host Cluster Facts," on page 82
- Section 5.8, "vCPU Slots for VM Hosts," on page 85

5.1 Info Page

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

- Section 5.1.1, "Show Inherited Fact Values Check Box," on page 76
- Section 5.1.2, "VM Host Information Panel," on page 76
- Section 5.1.3, "Provisioning Adapter Config Panel," on page 79
- Section 5.1.4, "Guest VM Monitor Information Panel," on page 79

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

5.1.1 Show Inherited Fact Values Check Box

Select this check box to show facts with overridden values supplied through attached or inherited policies. These fact values are read only (non-editable), although you can use the Policy Editor to modify the policy values themselves if you want to.

5.1.2 VM Host Information Panel

The **VM Host Information** panel on the Info page includes the following fields:

- Physical Resource
- VM Host Type
- VmHost Cluster
- Enabled
- Online
- Healthy
- Shutting Down
- Location
- Supports VM Migration
- Supports H/W HVM
- Accounting Group
- Max Hosted VMs
- Max Hosted vCPUs
- Max Virtual Memory
- Repositories
- Available VM Resource Groups
- Managing Job
- Needs Resync

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

Physical Resource: (Read Only) The name of the resource that houses this VM host container.

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

```
<fact name="vmhost.resource" value="vmh7sles" type="String" />
```

VM Host Type: This field displays a read-only fact, as discovered by the Orchestration Server. It identifies the VM host as a regular VM host (**vmhost**) or as a VMware cluster (**vmhostcluster**).

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

```
<fact name="vmhost.type" value="vmhostcluster" type="String" />
```

VmHost Cluster: (Conditional) This field displays a read-only fact, as discovered by the Orchestration Server. It identifies the VM host cluster to which this VM host belongs. The field is displayed on the information page only when the VM host is a member of a cluster.

In the fact editor, this fact is listed as `vmhost.cluster`:

```
<fact name="vmhost.cluster" value="esx35_cluster_vsphere" type="String" />
```

Enabled: This check box is selected by default. When it is selected (it has a value of true), the VM host is enabled, which means that VM instances can be provisioned on it.

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

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

Online: When this check box is selected (it has a value of true), the agent on the physical resource is online.

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

```
<fact name="vmhost.online" value="true" type="Boolean" />
```

Healthy: This check box is selected by default. When it is selected (it has a value of true), the VM host is designated as being in good health. You can set the health of the object by selecting or deselecting the health check box. Changing the value in this way has an immediate effect unless the value is overriden by an attached policy. For more information, see Chapter 14, “Grid Object Health Monitoring,” on page 169

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

```
<fact name="vmhost.health" value="false" type="Boolean" />
```

Shutting Down: When this check box is selected (it has a value of true), the VM host is attempting to shut down and does not accept provisioning requests.

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

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

Location: For the vsphere provisioning adapter, this is the `ManagedObjectReference` path for this VM host. For other provisioning adapters, this is an optional description of the physical location of the VM host.

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

```
<fact name="vmhost.location" value="" type="String" />
```

Supports VM Migration: When this check box is selected (it has a value of true), the VM host can support VM migration. The state of this fact can also depend on the migration capabilities of the provisioning adapter used to provision the VM.

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

```
<fact name="vmhost.migration" value="true" type="Boolean" />
```

Supports H/W HVM: When this check box is selected (it has a value of true), the hypervisor on the VM host can support hardware virtualization.

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

```
<fact name="vmhost.hvm" value="false" type="Boolean" />
```

Accounting Group: The default VM host group that you want to be adjusted for VM tracking statistics.

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

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

Max Hosted VMs: The maximum number of VM instances allowed on this VM host.

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

```
<fact name="vmhost.maxvmslots" value="8" type="Integer" />
```

Max Hosted vCPUs: The maximum number of virtual CPUs that this VM host can support.

In the Fact Editor, this fact is listed as `vmhost.vcpu.max`:

```
<fact name="vmhost.vcpu.max" value="8" type="Integer" />
```

For more information about vCPU slots, see Section 5.8, “vCPU Slots for VM Hosts,” on page 85.

Max Virtual Memory: The amount of memory (measured in MB) available to hosted VMs.

In the Fact Editor, this fact is listed as `vmhost.memory.max`:

```
<fact name="vmhost.memory.max" value="1000" type="Integer" />
```

Repositories: The list of repositories (VM disk stores) that are visible to this VM host.

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

```
<fact name="vmhost.repositories">
  <array type="String">
  </array>
</fact>
```

You can edit this array by clicking the ··· button to open the Choose Grid Objects dialog, where you can add, remove, or edit repositories in an array of repository choices.

Available VM Resource Groups: This field displays a list of resource groups containing VMs that are allowed to run on this VM host.

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

```
<fact name="vmhost.vm.available.groups">
  <array type="String">
  </array>
</fact>
```

You can edit this array by clicking the ··· button to open the Choose Grid Objects dialog, where you can add, remove, or edit the resource groups (element values) in an array of choices.

Managing Job: The ID of a running job that manages VM operations on this VM host. When this field is completed, the VM Manager prevents other jobs from initiating provisioning actions. The fact is cleared when the managing job ends.

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

```
<fact name="vmhost.controllingjob" value="" type="String" />
```

Needs Resync: When this check box is selected (it has a value of true), you specify that, at the next opportunity, this VM host is to be probed to resynchronize all the VMs that are managed here.

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

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

5.1.3 Provisioning Adapter Config Panel

- Adapter Job Name
- Username
- Password

Adapter Job Name: The name of the provisioning adapter job that manages VM discovery on this host. Do not change this value unless you have implemented your own discovery job.

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

```
<fact name="vmhost.provisioner.job" value="vsphere" type="String" />
```

Username: (Optional) The user name required for provisioning on the VM host.

In the Fact Editor, this fact is listed as `vmhost.provisioner.username`:

```
<fact name="vmhost.provisioner.username" value="" type="String" />
```

Password: (Optional) The password required for provisioning on the VM host.

In the Fact Editor, this fact is listed as `vmhost.provisioner.password`:

```
<fact name="vmhost.provisioner.password" value="" type="String" />
```

5.1.4 Guest VM Monitor Information Panel

- Current VM Count
- Available vCPUs
- Available Virtual Memory
- VM Image Counts
- Running VM Instances
- Load Index (Slots)
- Load Index (Memory)

Current VM Count: (Read Only) The current number of active VM instances.

In the Fact Editor, this fact is listed as `vmhost.vm.count`:

```
<fact name="vmhost.vm.count" value="0" type="Integer" />
```

Available vCPUs: The number of vCPUs available on this VM host.

In the Fact Editor, this fact is listed as `vmhost.vcpu.available`:

```
<fact name="vmhost.vcpu.available" value="8" type="Integer" />
```

For more information, see Section 5.8, “vCPU Slots for VM Hosts,” on page 85.

Available Virtual Memory: The amount of memory (measured in MB) available to new VMs.

In the Fact Editor, this fact is listed as `vmhost.memory.available`:

```
<fact name="vmhost.memory.available" value="1000" type="Integer" />
```

VM Image Counts: The dictionary of running instance counts for each running VM template.

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

```
<fact name="vmhost.vm.templatecounts">
  <dictionary>
    <dictelement key="ads">
      <time>12:00 AM</time>
    </dictelement>
  </dictionary>
</fact>
```

You can edit the dictionary elements by clicking the ··· button to open the VM Image Counts dialog, then adding or removing the names in the dictionary.

Running VM Instances: (Read Only) A list of active VM instances.

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

```
<fact name="vmhost.vm.instanceids">
  <array type="String">
  </array>
</fact>
```

Load Index (Slots): (Read Only) The current loading index of resource slots, which is a ratio of the active hosted VMs to the specified maximum number of VMs allowed on this host. Each provision VM takes up one slot. For more information, see *Max Hosted VMs*.

In the Fact Editor, this fact is listed as `vmhost.loadindex.slots`:

```
<fact name="vmhost.loadindex.slots" value="0.1250" type="Real" />
```

Load Index (Memory): (Read Only) The current loading index for memory, which is a ratio of the virtual memory consumed on this VM host to the specified maximum amount of memory allocated to this host.

In the Fact Editor, this fact is listed as `vmhost.loadindex.virtualmemory`:

```
<fact name="vmhost.loadindex.virtualmemory" value="0.0000" type="Real" />
```

5.2 Policies Page

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

NOTE: You can edit a policy by right-clicking a policy icon and selecting **Edit Policy**. Remember to click the Save button when your changes are complete.

5.3 Health Debugger Page

The Health Debugger is a common Admin view in the Orchestration Console for most Grid objects. For information about this tool, see Health Debugger in Chapter 14, “Grid Object Health Monitoring,” on page 169.

5.4 Constraints/Facts Page

The **Constraints/Facts** tab opens a page that shows all of the effective constraints and facts for a Grid object. Each Grid object has an associated set of facts and constraints that define its properties. By building, deploying, and running jobs on the Orchestration Server, you can individually change the functionality of any system resources by managing an object's facts and constraints. The Orchestration Server assigns default values to each of the component facts. Facts with no mode listed can be changed at any time by the administrator. Facts with mode `r/o` have read-only values, which can be viewed by using the pencil icon, but changes cannot be made.

For information about using the Fact Editor on this page to rename VM host objects, see Section 5.6, "VM Host Object Naming and Renaming," on page 81

5.5 Action History Page

The **Action History** tab is displayed in the administrative view of the Repository object. When you select the **Action History** tab, a table displays on the Action History page with a list of the history for all VM provisioning actions performed on this Grid object.

The Orchestration Server must be connected to an audit database for the **Include Audit Database** check box to be available. If the **Include Audit Database** check box is selected in this view, the action status is not polled. Click the Refresh button in the toolbar to retrieve and display fresh data.

For more details about the information listed on the Action History page, see "Action History in Admin Views of the Orchestration Console" in the *Cloud Manager Procedures Guide*.

5.6 VM Host Object Naming and Renaming

Some VM host names (or VM host Cluster names) are generated by the Orchestration Server and can therefore receive generic, arbitrary names such as `host2_demoAdapter`, `host3_demoAdapter`, and so on. VM hosts you name at creation time might also change later in their purpose or facilities.

As the quantity of these VM host objects grows in your grid, you might find it helpful or necessary to rename them, assigning more meaningful, intuitive names to suit the purpose of the object. The object's "display name" is visible in the Orchestration Console interface, the , and in optional zos and zosadmin commands.

NOTE: Resource object groups (that is, the folders that contain these VM host objects) can also be renamed. Objects such as jobs, events, and users cannot be renamed.

A VM host object's name is stored in the `${objectType}.displayname` fact, which exists on every Grid object type, even those objects that cannot be renamed.

You can rename a VM host object (or a VM host Cluster object) in the Orchestration Console by using one of three methods:

- Right-click the VM host object in the Explorer tree, then select **Rename** to allow editing of the display name.
- Triple-click the VM host object in the Explorer tree to allow editing of the display name.
- In the Constraints/Facts page, select the VM host object `.displayname` fact and then open the Fact Editor to enter a new value for that fact.

As you use one of these methods, you will notice that the fact value is pre-populated with the `${objectType}.id` fact. This functions as the name value for the object name until you decide to change it.

NOTE: Even after being renamed, the VM host object retains its associated resource ID in the `.id` fact. This is not editable.

For more information about making the Resource object display names visible from the zos or zosadmin command line, see Chapter 17, “Orchestration Server zos Command Line Tool,” on page 195 and Chapter 18, “Orchestration Server zosadmin Command Line Tool,” on page 243.

5.7 Unique VM Host Cluster Facts

There are several VM host cluster-related facts that are not found in a regular VM host object. This section contains detail about those facts.

- Section 5.7.1, “Orchestration Server Facts in the VM Host Cluster Object,” on page 82
- Section 5.7.2, “Orchestration Server Facts in a VM Host Residing in a Cluster,” on page 83
- Section 5.7.3, “Orchestration Server Facts in VMs Hosted in Clusters,” on page 84

5.7.1 Orchestration Server Facts in the VM Host Cluster Object

Any vSphere clusters discovered by the Orchestration Server are listed in the Orchestration Console as members of a convenience group (for example, a group named `clusters_vsphere`). The following table lists the read-only, cluster-related facts in a VM host Cluster object.

Table 5-1 *Cluster-Related Facts in a Cluster Object*

Fact Name	Type	Description
`vmhost.cluster.vmhosts`	String[]	This string array lists all of the VM hosts that are members of this cluster.
`vmhost.location`	String	The data center Managed Object Reference (MOR) path to the cluster.
`vmhost.vphere.cluster.das.admission_control_enabled`	Boolean	If this fact value is true, the VM host cluster has high availability configuration to use admission control. **NOTE:** High availability was originally named “Dynamic Availability Service” (DAS) in VMware. This is the origin of the `.das` prefix.
`vmhost.vphere.cluster.das.enabled`	Boolean	If this fact value is true, high availability is turned on in this cluster.

Fact Name	Type	Description
vmhost.vphere.cluster.drs.allow_behavior_override	Boolean	Whether VMs can specify their own placement behavior.
vmhost.vphere.cluster.drs.default_behavior	String	Specifies manual, partially automated, or fully automated VM placement.
vmhost.vphere.cluster.drs.enabled	Boolean	If this fact value is true, DRS is turned on in this cluster.
vmhost.vphere.pools	Dictionary	The mapping of the pool MOR path to dictionaries of pool configuration values.
vmhost.type	String	The “type” differentiation for clusters vs. VM hosts. The value is either vmhost or vmhostcluster.

5.7.2 Orchestration Server Facts in a VM Host Residing in a Cluster

The following table lists the read-only, cluster-related facts in a virtual or physical machine with an installed Orchestration Agent and residing in a cluster.

Table 5-2 *Cluster-Related Facts in a VM Host Residing in a vSphere Cluster*

Fact Name	Type	Description
vmhost.cluster	String[]	The ID of the cluster that contains this host.
vmhost.type	String	The “type” differentiation for clusters vs. VM hosts. The value is either vmhost or vmhostcluster.

5.7.3 Orchestration Server Facts in VMs Hosted in Clusters

The following table lists the read-only, cluster-related facts in a VM hosted in vSphere clusters.

Table 5-3 *Cluster-Related Facts in VMs Hosted in Clusters*

Fact Name	Type	Description
`resource.provision.vmhost`	String	The ID of the VM host cluster or VM host that contains this VM. When provisioning or migrating to a cluster, this value is initially the cluster ID. On subsequent resync or discovery or when using the vSphereUpdate daemon, this is set to the value of the selected VM host within the cluster.
`resource.vm.pool`	String	The MOR path of the resource pool that contains this VM. This value is also a key in the `vmhost.vsphere.pools` dictionary.
`resource.vm.vmhost.location`	String	The MOR path of the VM host cluster that contains this VM. This is always the cluster if the VM host is in a cluster; otherwise, it is the MOR path of the VM host itself.
`resource.vm.vsphere.cpu.limit`	Integer	The maximum amount (in MHz) of CPU resources to be used by this VM.
`resource.vm.vsphere.cpu.reservation`	Integer	The minimum amount (in MHz) of CPU resources guaranteed to this VM.
`resource.vm.vsphere.cpu.shares.level`	String	The relative amount (assigned a value of low, normal, high, or custom) of CPU allocated for this VM.
`resource.vm.vsphere.cpu.shares.custom`	Integer	The custom amount of relative CPU for this VM. This fact is valid only when the value for the `resource.vm.vsphere.cpu.shares.level` fact is “custom.”
`resource.vm.vsphere.memory.limit`	Integer	The maximum amount (measured in MB) of memory resources to be used by this VM.
`resource.vm.vsphere.memory.reservation`	Integer	The minimum amount (measured in MB) of memory resources guaranteed to this VM.
`resource.vm.vsphere.memory.shares.level`	String	The relative amount (assigned a value of low, normal, high, or custom) of memory specified for this VM.

Fact Name	Type	Description
`resource.vm.vsphere.memory.shares.custom`	Integer	The custom amount of relative memory for this VM. This fact is valid only when the value for the `resource.vm.vsphere.memory.shares.level` fact is “custom.”

5.8 vCPU Slots for VM Hosts

In the Orchestration Server, a vCPU represents a logical CPU. It provides a way to set up limits for allocating CPUs on VM hosts. These limits let you specify how many vCPUs should be hosted by each VM host so that you can control how much CPU processing power is available. If all vCPUs are in use, a subsequent provision can be denied or made to wait. This lets you ensure the quality of service you want to maintain in the data center.

When a VM is provisioned, the Orchestration Server runs a constraint check on every suitable VM host to determine if the number of available vCPUs on the VM host is sufficient for a VM. If a VM host with sufficient available vCPUs is not available, the provisioning request waits until one becomes available or (depending on VM facts) the request is denied.

Using vCPU facts differs from using the existing slot fact (`vmhost.maxvmslots`). The `maxvmslots` fact provides basic control of the number of VMs allocated to a VM host, which is useful for limiting VMs because of license restrictions or for generally limiting the VMs being managed. The vCPU facts are similar to the memory limit facts, giving you more control to avoid overloading a VM host and letting you ensure quality of service.

This section includes information about how the vCPU facts are used in the Orchestration Server.

- Section 5.8.1, “Configuring vCPUs on VM Hosts,” on page 85
- Section 5.8.2, “Configuring vCPUs on VM Host Clusters,” on page 86
- Section 5.8.3, “Configuring vCPUs on VMs,” on page 86

5.8.1 Configuring vCPUs on VM Hosts

There are two vCPU facts displayed on the VM host Info page in the Orchestration Console:

- **Max Hosted vCPUs:** This value (Integer) represents the maximum number of vCPUs that the VM host can support. The fact name is `vmhost.vcpu.max`.
- **Available vCPUs:** This value (Integer) represents the number of virtual CPUs available on this host. The fact name is `vmhost.vcpu.available`.

The `vmhost.vcpu.available` value changes when a VM is provisioned or shut down on that VM host. If the `vmhost.vcpu.max` fact is set to -1 (unlimited), the `vmhost.vcpu.available` value changes to -1 (unlimited). When it is set to unlimited, no counting occurs, so the Orchestration Server does not check vCPU limits.

When a VM host object is created (during discovery), the `vmhost.vcpu.max` value is set to the number of physical cores multiplied by a factor of 4. For example, on a VMware VM host that has eight physical cores, the Orchestration VM host discovery for vSphere creates a VM host object with a maximum vCPU of 32. The factor value of 4 represents partitioning a physical core to four vCPUs, which represents 25% capacity.

You can change this default value by creating a policy that sets the `vmhost.vcpu.max` fact value and associates the policy either to the VM host or to a Resource Group of VM hosts. In the preceding vSphere VM host example, if you wanted to partition the eight physical cores on a VM host to 50% capacity, you would set the maximum vCPUs to 16 by creating the following policy and then associating it to the VM host or to a Resource Group of VM hosts:

```
<policy>
   <vmhost>
       <fact name="vcpu.max" type="Integer" value="16" />
  </vmhost>
</policy>
```

If you want the Orchestration Server to allocate vCPUs without checking limits, you can set the maximum vCPUs to -1, which indicates an unlimited number. You would create the following policy to make that configuration setting:

```
<policy>
   <vmhost>
       <fact name="vcpu.max" type="Integer" value="-1 />
  </vmhost>
</policy>
```

5.8.2 Configuring vCPUs on VM Host Clusters

Because the VM host Cluster object represents a set of VM hosts, both `vmhost.vcpu.available` and `vmhost.vcpu.max` facts are sums of the underlying VM host objects. If all of the underlying VM hosts have their `vmhost.vcpu.max` values set to -1 (unlimited), then the corresponding `vmhost.vcpu.max` fact in the VM host Cluster is -1 (unlimited). This also means that no counting or checking occur for vCPU limits. This is used in scenarios where you rely on the underlying hypervisor to account for vCPUs rather than the Orchestration Server. For example, if you have set up vSphere with DRS and clustering, you do not need the server to do any checking.

NOTE: The Orchestration Server vsphere provisioning adapter sets the `vmhost.vcpu.max` fact value to -1 (unlimited) for DRS-enabled VM host clusters.

5.8.3 Configuring vCPUs on VMs

An existing vCPU fact (`resource.vm.vcpu`) specifies the number of vCPUs for a VM. This fact is set on VM image discovery. You can view it in the Orchestration Console on the VM object Admin view on the Info/Groups page. The Virtual Machine Configuration pane on this page has a **Host CPU Number** field where this fact is set.

In the VM Hosts Monitor view, the tooltip for a VM host displays the following:

`..Total vCPUs: xx, Available vCPUs: xx`

Figure 5-1 *Tool Tip Text for VM Host vCPU Slots*

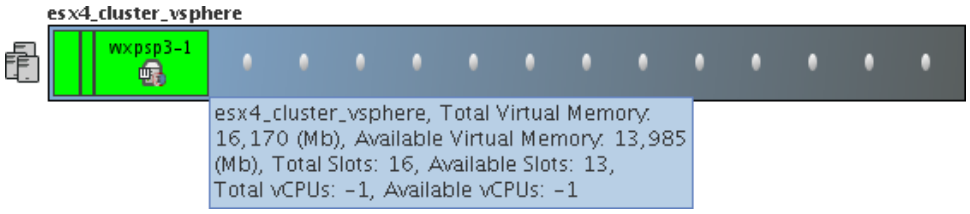

These values change when a VM is provisioned or shut down. For example, if the `vmhost.vcpu.max` fact is set to -1 (unlimited), then the tooltip shows “-1” for both total and available.

6 The Virtual Disk Object

This section includes the following information:

- Section 6.1, “Understanding the Virtual Disk Object,” on page 89
- Section 6.2, “Viewing Virtual Disk Configuration in the Orchestration Console,” on page 94
- Section 6.3, “Managing Block Devices as VM Virtual Disks,” on page 99

6.1 Understanding the Virtual Disk Object

A virtual disk (vDisk) represents a VM’s view of its storage devices, which could include any type of physical disk (such as a file-backed disk image, an ISO image file, a physical hard drive, a CD/DVD device, or a block device) associated to a VM. The vDisk objects are discovered, along with their associated VM, when a Discover VM Images job is run on a repository.

The vDisk is modeled as a Grid object, located as a subordinate to the VM Grid object in the Explorer tree of the Orchestration Console. In the Explorer Tree, a vDisk is given the form *vmname*`_vdisk<`*n*`>` where <*n*> represents the numerical order in which this vDisk was discovered, with 1 appended to the name of the first vDisk discovered or created. For example, `suse11_vdisk1` would be the name of the first disk discovered for a VM with the Grid ID `suse11`. Each additional vDisk is incremented by one, so the second vDisk in this example would be named `suse11_vdisk2`.

This section includes the following information:

- Section 6.1.1, “Creating or Deleting a vDisk in the Orchestration Console,” on page 89
- Section 6.1.2, “Sharing Virtual Disks Among VM Hosts,” on page 92
- Section 6.1.3, “Moving Virtual Disks,” on page 93

6.1.1 Creating or Deleting a vDisk in the Orchestration Console

This section includes the following information:

- “Creating and Configuring a Virtual Disk” on page 89
- “Creating a Sparse Virtual Disk” on page 91
- “Deleting a Virtual Disk” on page 91

Creating and Configuring a Virtual Disk

You might want to manually create a vDisk in the following scenarios:

- When you want to create a “blank” disk image file for the VM. In this scenario, the disk image does not actually reside on the local file system, but a disk image of the specified size (measured in MB) should be created at the location specified for use by the VM. This is essentially a blank file, until it is used by the VM.
- When the Orchestration Server might not have discovered the vDisk objects correctly, such as omitting a disk that should exist. You need to manually correct the incorrect discovery.

- A VM that already exists needs to have patches applied to it. The patches are delivered through an ISO file, which was not configured to be attached to the VM. This configuration lets the administrator configure the VM with access to the ISO disk image, then apply the patches, and then later delete the vDisk object, returning the VM to its original configuration.

  You need to manually add the vDisk, select the **Save Config** action in the Orchestration Console, then apply the patches to the running VM. Later, you shut down the VM, delete the vDisk object from the Orchestration Server, then select the **Save Config** action again.

  The scenario includes configuring the VM to use the existing ISO file (that is, creating the vDisk object, then selecting **Save Config**), and then deconfiguring the VM to no longer use the ISO file (that is, deleting the vDisk object, then selecting **Save Config**).

  In this scenario, only the vDisk object from the Orchestration Server is deleted, not the ISO file.

To create a virtual disk in the Orchestration Console, you can either right-click the VM where you want to create the vDisk, then select **Create Virtual Disk** (if you do this, you can skip to Step 4 below) or you can use the following procedure from the Orchestration Console menu:

1. In the Orchestration Console main menu, select **Actions** > **Create Virtual Disk** to display the Create a New Virtual Disk dialog.

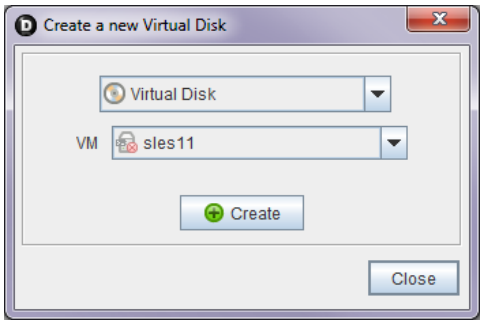

2. In the **VM** drop-down list, select the name of the VM where you want to add a vDisk, then click **Create**.
3. When you have created all of the vDisks you need, click **Close**.
4. Select a newly created vDisk object in the Explorer tree to view the Info/Groups page of the admin view.
5. On the Info/Groups page, configure the following settings:

   Type: Specify the vDisk type as the VM host sees it.

   Description: Describe the vDisk with any text that you choose

   Healthy: Designates the health state of the vDisk. Do not configure.

   Moveable: Specifies whether the disk image can be copied (relocated) with the VM when the VM is moved (relocated) to another repository. For more information, see “Moveable” later in this section.

   Mode: Specifies the mode of the vDisk as made available and supported by the provisioning adapter:

   - r = read only
   - w = read/write

   VM: Specifies the name of the VM that uses this vDisk.

Repository: The repository where this disk location path resides.This setting is important because the Orchestration Server uses it to find a suitable VM host for provisioning, building, or migration actions. The value of this setting can be **none**, which informs the server to ignore this vDisk when locating a suitable VM host.

Physical Disk: The name of the pDisk that this vDisk is associated with.

Location: The path (location) to the disk image.

- If you specify a location to a disk that already exists, the existing disk file is used and the VM configuration is modified (according to the value in `vdisk.location` fact) to use this existing disk.
- If you specify a path to a disk that does not exist (that is, if the value in `vdisk.location` is invalid), the action fails and an empty disk image file of the specified size is created. An error in the action status or job log is created.

Size: The size (measured in MB) of the disk image. Do not configure.

Sparse Disk: Designates whether the vDisk file is a sparse file. Do not configure.

Actual Size: The actual sparse size (measured in MB) of the vDisk file. Do not configure.

6 Click the Save button  in the toolbar to save the fact changes you made.

7 In the Explorer tree, right-click the VM object where you added the vDisk, then select **Save Config** to apply the changes to the VM's configuration.

Creating a Sparse Virtual Disk

Sparse disk creation is supported by the vSphere hypervisor. If you want to create your vDisk as a sparse file, you can use the procedure for "Creating and Configuring a Virtual Disk" on page 89 (see Step 5 on page 90 in particular).

You need to set the Sparse Disk fact to true, specify the amount of space (in MB) for the disk in the Size fact, and also specify the path to the repository where the sparse vDisk resides. Ensure that you perform the **Save Config** action to apply the vDisk changes to the VM's configuration.

Deleting a Virtual Disk

You might want to manually delete a vDisk in at least two scenarios:

- When the Orchestration Server might not have discovered the vDisk objects correctly, such as adding a disk that should not exist. The administrator needs to manually correct the incorrect discovery.
- A VM that already exists needs to have patches applied to it. The patches are delivered through an ISO file, which was not configured to be attached to the VM. This configuration lets the administrator configure the VM with access to the ISO disk image, then apply the patches, then later delete the vDisk object, returning the VM to its original configuration.

  The administrator needs to manually add the vDisk, run the Save Config command from the Orchestration Console, then apply the patches to the running VM. Later, the administrator shuts down the VM, deletes the vDisk object from the server, then performs the **Save Config** action again.

  The scenario includes configuring the VM to use the existing ISO file (that is, creating the vDisk object and selecting the **Save Config** action), then deconfiguring the VM to no longer use the ISO file (that is, deleting the vDisk object, then selecting the **Save Config** action).

  In this scenario, only the vDisk object from the server is deleted, not the ISO file.

To delete a virtual disk, you can either right-click the vDisk object in the Explorer, then select **Delete** (if you do this, you can skip to Step 4, below), or you can use the following procedure from the Orchestration Console:

1. In the Orchestration Console, select **Actions** > **Delete** > **Delete Virtual Disk** to display the Delete a Virtual Disk dialog.

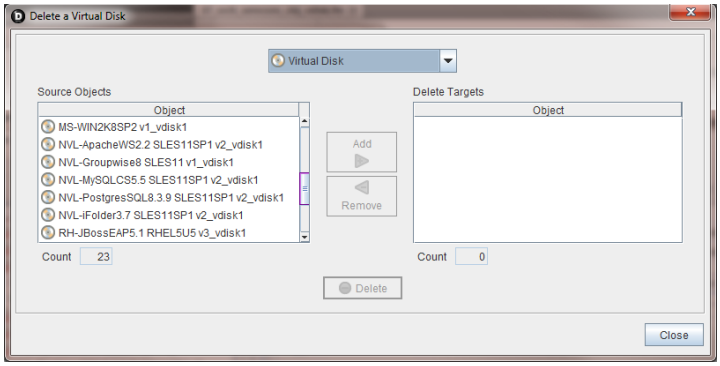

2. In the **Source Objects** list, select the name of the vDisk (hold down the Ctrl key to select multiple objects), then click **Add** to move these objects to the **Delete Targets** list.
3. When you have selected all of the vDisks you want to delete, click **Delete** to display the Delete dialog.

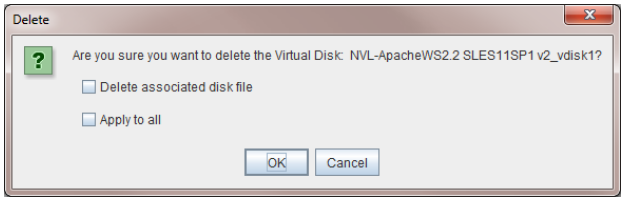

4. In the dialog, select **Apply to all** to delete all of the vDisk objects in the **Delete Targets** list, select **Delete associated disk file** if you want to delete the files associated with the vDisks you have selected, click **OK**, then click **Close**.
5. In the Explorer tree, right-click the VM object where you deleted the vDisk, then select **Save Config** to apply the changes to the VM's configuration.

NOTE: The **Save Config** action rewrites the configuration file for the VM (for example, `config.vsphere` for the vsphere provisioning adapter), but it does not delete any vDisk files on the file system. In this case, manual deletion of the vDisk file is required.

To delete a VM and its backing files, use the **Delete/Destroy Resource** action.

6.1.2 Sharing Virtual Disks Among VM Hosts

For a VM to be provisionable by other VM hosts, all of a VM's vDisks must be visible in the same way that the VM's default repository (`resource.vm.repository`) is visible to VM hosts. If a VM has multiple vDisks and each vDisk has a different associated repository, these repositories must also be visible from a potential VM host.

6.1.3 Moving Virtual Disks

When you move a VM to a new repository (see the *Move Disk Image* action for each hypervisor at Chapter 19, “The VMware vSphere Provisioning Adapter,” on page 285), all of that VM’s moveable vDisk images (see “Moveable:” on page 95) are moved with it to be co-located in the same repository. The Orchestration Server uses the aggregated size of each moveable vDisk to determine if the designated repository has enough space for all of the disk images. vDisks that are marked as not moveable stay in place and are not used in the calculation for the VM disk size.

The following illustration further explains this concept:

Figure 6-1 *Example of Moving Virtual Disks with the VM*

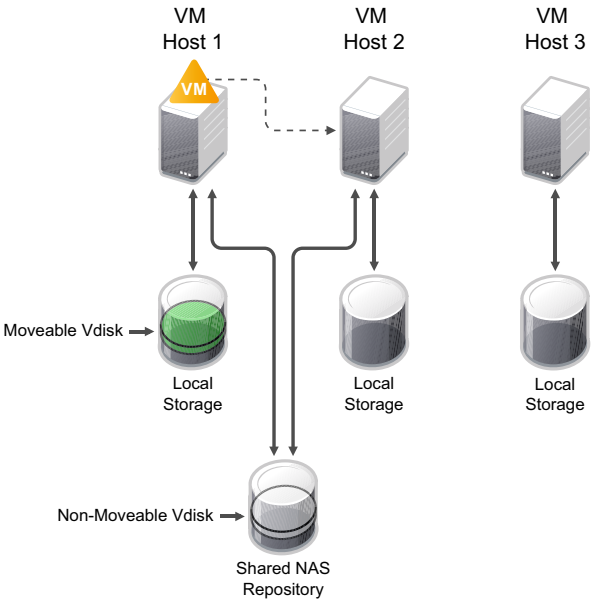

- VM host 1, VM host 2, and VM host 3 all have their own local storage repositories.
- VM host 1 has a vDisk located on it. It is designated as a moveable vDisk.
- VM host 1 and VM host 2 are also connected to a shared NAS storage repository.
- The local repository connected to VM host 1 has a vDisk located on it. It is designated as a moveable vDisk.
- The shared NAS repository has a vDisk located on it. It is designated as a non-moveable vDisk.

NOTE: Shared repositories are not created on discovery. They must be manually created and the sharing (visibility) configured.

- VM host 1 has a VM located on it.
- VM host 3 cannot communicate with the NAS repository; its `vmhost.repositories` fact does not include the NAS repository in the array, so that repository is not visible to VM host 3.

If you want to move the VM from VM host 1 to another VM host, the server manifests the following behavior:

- The vDisk sizes used by the VM (on local storage and shared storage) are aggregated and compared to free space available on the repositories.

- The only vDisk that is allowed to move is the moveable disk. This disk would be copied to either the shared NAS repository or the local storage on VM host 2.
- VM host 3 is not considered because it does not have access to the non-moveable disk on the NAS repository.

6.2 Viewing Virtual Disk Configuration in the Orchestration Console

You can visually expose a vDisk Grid object in the Orchestration Console in two ways:

- In the Explorer Tree, select a VM Resource object, select the **Info/Groups** tab in the admin view to open the VM Info Groups page, then scroll to the **Virtual Machine Configuration** panel on that page. Right-click the vDisk icon in that panel to display the four tabs in the Virtual Disk admin view.
- In the Explorer Tree, click the expand/collapse icon of a VM Grid object, identify the vDisk icon, then select the icon to display the four tabs in the Virtual Disk admin view.

Figure 6-2 *The Virtual Disk Info/Groups Page*

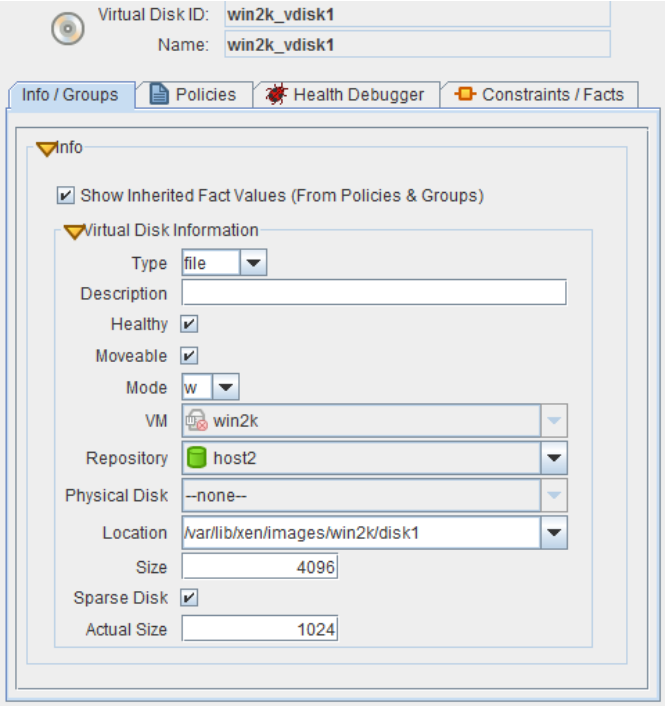

The page that opens under the **Info** tab includes fields where you can configure the general information and attributes (facts) of the vDisk.

NOTE: Whenever you make changes to vDisk object facts, the write icon is superimposed on the object's icon, signifying that the object has been changed. If you want to save the changes you have made, click the Save button on the Orchestration Console toolbar.

This section includes the following additional information:

- Section 6.2.1, “Virtual Disk Information Panel,” on page 95
- Section 6.2.2, “Virtual Disk Policies Tab,” on page 97
- Section 6.2.3, “Virtual Disk Health Debugger Tab,” on page 97
- Section 6.2.4, “Virtual Disk Constraints/Facts Tab,” on page 97
- Section 6.2.5, “Virtual Disk Object Naming and Renaming,” on page 98

6.2.1 Virtual Disk Information Panel

The Virtual Disk Information panel on the Info page includes the following fields:

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

Type: This drop-down list lets you select one of the vDisk types as the VM host sees it:

- **file:** Specifies this vDisk as a file-backed disk.
- **block:** Specifies this vDisk as a block device.

In the Fact Editor, this fact is listed as `vdisk.type`:

```
<fact name="vdisk.type" value="file" type="String" />
```

Description: Describes the vDisk with any text that you choose.

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

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

For a vDisk discovered and managed by the vsphere provisioning adapter, this field is populated with a display name label obtained by vSphere and mapped to this vDisk by the Orchestration Server.

Healthy: For a vDisk managed by any supported hypervisor, this check box is selected by default, which designates the vDisk as being in good health.

NOTE: We recommend that you do not change the **Healthy** value from its default.

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

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

Moveable: When this check box is selected (its value is true), the vDisk is moveable, which means that the disk image can be copied to a different repository when the VM moves.

In the Fact Editor, this fact is listed as `vdisk.moveable`:

```
<fact name="vdisk.moveable" value="true" type="Boolean" />
```

If **Moveable** is not selected, the disk image must stay at its current location because it cannot be copied or moved. By default upon discovery, if the server sees that this vDisk is an ISO image, the fact is set to false because it is assumed that the administrator doesn’t want to copy ISO images from one location to another.

Whenever you want to prevent a vDisk from being moved, you can deselect this check box.

The vDisk is not deleted during a VM **Delete/Destroy** action if this check box is deselected.

NOTE: In Cloud Manager 2.0, block type disks (that is, pDisks) are not moveable, even if you change this setting.

Mode: Specifies the mode of the vDisk as made available and supported by the provisioning adapter:

- r = read only
- w = read/write

In the Fact Editor, this fact is listed as `vdisk.mode`:

```
<fact name="vdisk.mode" value="w" type="String" />
```

VM: (Read Only) Specifies the name of the VM that uses this vDisk.

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

```
<fact name="vdisk.vm" value="mysql" type="String" />
```

This is a fact junction referencing the associated VM. Conversely, the `resource.vm.vdisks` fact visible from the VM Grid object is a fact junction showing the associated vDisks associated with the VM.

Repository: The storage location containing the vDisk image on the VM host.

Changing this fact after discovery only corrects a possible incorrectly discovered fact. Changing the storage location does not move the vDisk.

In general terms, a **block** type Repository represents a container for physical devices you can use for VM disks. If you select this type, the physical disks (Logical Volumes) on that VM host are listed.

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

```
<fact name="vdisk.repository" value="zos" type="String" />
```

Physical Disk: Indicates the name of the pDisk to which this vDisk is associated.

In the Fact Editor, this fact is listed as `vdisk.pdisk`:

You can edit this fact only if a **block** type Repository is selected for the `vdisk.repository` fact. When you make that selection, a drop-down list of Physical Disks becomes available for this field.

```
<fact name="vdisk.pdisk" value="" type="String" />
```

Although there might be several pDisks available in the block Repository, you can select only one.

Location: For file-backed disks, this fact represents the file system path to the vDisk image in the specified repository.

For example, a vDisk located on an NFS repository datastore would show the URI to the NFS share with the path to the disk appended to it.

For **block** type disks, this fact contains the URI to the block device, for example `/dev/hdc`, which could represent a CD/DVD tray on a VM host.

In the Fact Editor, this fact is listed as `vdisk.location`:

```
<fact name="vdisk.location" value="/vmfiles/vms/vsphere-vm123/disk1" type="String"
/>
```

Size: The size (measured in MB) of this vDisk image.

In the Fact Editor, this fact is listed as `vdisk.size`:

```
<fact name="vdisk.size" value="2048" type="Integer" />
```

The disk size value for each moveable vDisk on a VM is aggregated by the Orchestration Server into the `resource.vm.vdisksize` fact, which is used to determine if the VM can relocate from one repository to another, given that the new repository has enough free space to store the VM.

Sparse Disk: When this check box is selected (its value is true), this is a sparse (thin) backed vDisk.

In the Fact Editor, this fact is listed as `vdisk.sparse`:

```
<fact name="vdisk.sparse" value="true" type="Boolean" />
```

Actual Size: The actual (sparse) size (measured in MB) of this virtual disk.

In the Fact Editor, this fact is listed as `vdisk.size.actual`:

```
<fact name="vdisk.size.actual" value="1024" type="Integer" />
```

6.2.2 Virtual Disk Policies Tab

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

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

6.2.3 Virtual Disk Health Debugger Tab

The Health Debugger is a common admin view in the Orchestration Console for most Grid objects. For information about this tool, see “Health Debugger” on page 171.

6.2.4 Virtual Disk Constraints/Facts Tab

To support constraining a VM’s provisioning actions based on more than one disk’s repository (that is, more than just `resource.vm.repository`), the vDisk can be referenced in constraints. The vDisk constraints are used to assign VM hosts during actions such as provisioning, building, or migrating. You can write constraints against attributes of disks (such as the repository where the vDisk resides) and against the available VM host repositories.

The **Constraints/Facts** tab opens a page that shows all of the effective constraints and facts for a Grid object. Each Grid object has an associated set of facts and constraints that define its properties. In essence, by building, deploying, and running jobs on the Orchestration Server, you can individually change the functionality of any and all system resources by managing an object’s facts and constraints.

The Orchestration Server assigns default values to each of the component facts. The following table lists the possible fact modes and their function:

Table 6-1 *vDisk Fact Modes Displayed in the Constraints/Facts Tab*

vDisk Mode Type	Mode Function
blank (no mode displayed)	read/write, not deleteable
`del`	read/write, deleteable

vDisk Mode Type	Mode Function
`dynamic`	read/write, not deleteable
`dynamic, r/o`	read only, not writable
`r/o`	read only

IMPORTANT: Several custom facts for the vDisk object can be added at discovery time, and vary according to the VM technology that manages the respective vDisk.

For more information about using the Fact Editor on this page to rename the Virtual Disk object, see Section 6.2.5, “Virtual Disk Object Naming and Renaming,” on page 98.

6.2.5 Virtual Disk Object Naming and Renaming

Some resource names are generated by the Orchestration system and can therefore have generic, arbitrary names such as `mysql-vdisk1`, `mysql-vdisk2`, and so on. A Virtual Disk (vDisk) you name at creation time might also change later in its purpose or facilities.

The object’s display name is visible in the Orchestration Console interface, the Cloud Manager Web Console and Mobile Clients interfaces, and in optional zos and zosadmin commands. As the number of these vDisk objects grows in your grid, you might find it helpful or necessary to rename them, assigning more meaningful, intuitive names to suit the purpose of the object.

NOTE: Resource object groups (that is, the folders that contain these vDisk objects) can also be renamed. Objects such as jobs, events, and users cannot be renamed.

A vDisk object’s name is stored in the `${objectType}.displayname` fact, which exists on every Grid object type, even those objects that cannot be renamed.

You can rename a vDisk object in the Orchestration Console by using one of three methods:

- Right-click the vDisk object in the Explorer tree, then select **Rename** to allow editing of the display name.
- Triple-click the vDisk object in the Explorer tree to allow editing of the display name.
- In the Constraints/Facts page, select the vDisk object `.displayname` fact and then open the Fact Editor to specify a new value for the fact.

As you use one of these methods, notice that the fact value is prepopulated with the `${objectType}.id` fact. This functions as the name value for the object name until you decide to change it.

NOTE: Even after being renamed, the vDisk object retains its associated resource ID in the `.id` fact. This is not editable.

For more information about making the Resource object display names visible from the zos or zosadmin command line, see Chapter 17, “Orchestration Server zos Command Line Tool,” on page 195 and Chapter 18, “Orchestration Server zosadmin Command Line Tool,” on page 243.

6.3 Managing Block Devices as VM Virtual Disks

Cloud Manager Orchestration Server can use a block device (an addressable, physical location) for storage on VMware vSphere managed hosts, but it manages such devices as physical disks (pDisks) attached to VM virtual disks (vDisks) to provide better performance compared to a file-backed virtual disk.

A vDisk represents a VM's view of its storage devices, but a pDisk represents a VM Host's view of its physical storage devices allocated for VM usage.

Managing block devices as pDisks also saves the additional setup and management that normal block devices require.

IMPORTANT: Block devices defined in the Amazon EC2 environment are not discovered or supported by the Orchestration Server in Cloud Manager.

- Section 6.3.1, "How Block Device Support Works," on page 99
- Section 6.3.2, "Viewing the Physical Disk Configuration in the Orchestration Console," on page 101

6.3.1 How Block Device Support Works

The following information is included in this section:

- "Creating New Block Type Repositories" on page 99
- "Discovering pDisks" on page 99
- "Attaching an Available pDisk to a VM" on page 100
- "Limitations of Block Device Support in This Release" on page 100

Creating New Block Type Repositories

When you run the **Discover Hosts and Repositories** action by using the appropriate the Orchestration Server provisioning adapter job, the Orchestration Server discovers the block device on the VM host and creates a "block" type Repository Grid object for it in the Explorer tree . The repository is to be used only as a container for physical disks.

For more information about repository Grid objects, see Chapter 9, "The Repository Object," on page 121.

Discovering pDisks

After the initial VM host and Repository discovery action, some block devices (considered to be a "physical disk" or "pDisk" in the the Orchestration Server model) might be discovered already attached to a VM, while some are not yet discovered and added under the block repository container. Further discovery requires a **Discover Disks** action on the new block Repository Grid object. The action launches a provisioning adapter job that discovers the physical disks available for the repository. When this job finishes, new pDisk objects  are created in the Explorer tree under the block Repository object.

You can also create a pDisk in a block Repository by right-clicking the block Repository object and selecting the **Create Physical Disk** action. If you create a pDisk without using the discovery, you need to provide information (that is, facts) about it for the Orchestration Server. For more information, see

Section 6.3.2, “Viewing the Physical Disk Configuration in the Orchestration Console,” on page 101. If you create this object without having a corresponding Logical Volume on the host system, you must create it there so that block device support for this pDisk works in the Orchestration Server.

Attaching an Available pDisk to a VM

After you discover a pDisk (that is, a Logical Volume) using **Discover Disks** on the block type Repository, that pDisk is not associated with any VM (workload) that Cloud Manager can access. You need to create this association by attaching the pDisk to a vDisk on an available VM.

1. In the Explorer tree of the Cloud Manager Orchestration Console, select a VM to which you want to attach a pDisk.
2. Right-click the VM object and select **Create Virtual Disk** to create a new vDisk object to be associated with the VM.
3. Attach the pDisk to the vDisk:
   - **3a** In the admin view of the vDisk object, select the **Info/Groups** tab to open the **Info** page for the object.
   - **3b** Select the **Repository** drop-down menu to display the list of available repositories, then select the appropriate block type repository from the list.
   - **3c** Select the **Physical Disk** drop-down menu to display the list of available pDisks, then select the appropriate pDisk form the list. The default value, **none**, does not select any physical disk.
   - **3d** Click the **Save** icon  on the toolbar to save the fact changes you have made.

     When you perform this action, the vDisk facts are automatically populated with the corresponding values from the pDisk object. These values become read-only, because they are retrieved from the pDisk.
   - **3e** In the Explorer tree, right-click the VM object associated with the new vDisk you created, then select the **Save Config** action to reconfigure this VM.

     NOTE: Remember that when you save the configuration, you only update the object model in the Orchestration Server; you have not affected the VM in the hypervisor. To do so, you must perform the **Save Config** action.

Limitations of Block Device Support in This Release

Block device support for this release has the following limitations:

- Orchestration Services does not allow the creation or deletion of pDisks. In other words, no LUN or volume management is possible from the Orchestration Console.
- Although a VM move is supported (block device disks are marked as unmoveable), cloning a VM that includes the pDisk is not supported if the `pdisk.vdisks.max fact` is set to 1. Cloning works if the fact is set to a value greater than 1.
- Not all variations of possible block-based data disks can be discovered on a VM host.

6.3.2 Viewing the Physical Disk Configuration in the Orchestration Console

You can visually expose a pDisk Grid object in the Explorer Tree of the Orchestration Console by selecting the expand/collapse icon of a block Repository object, identifying the pDisk icon, then selecting the icon to display the four tabs in the Virtual Disk admin view.

Figure 6-3 *The Physical Disk Info/Groups Page*

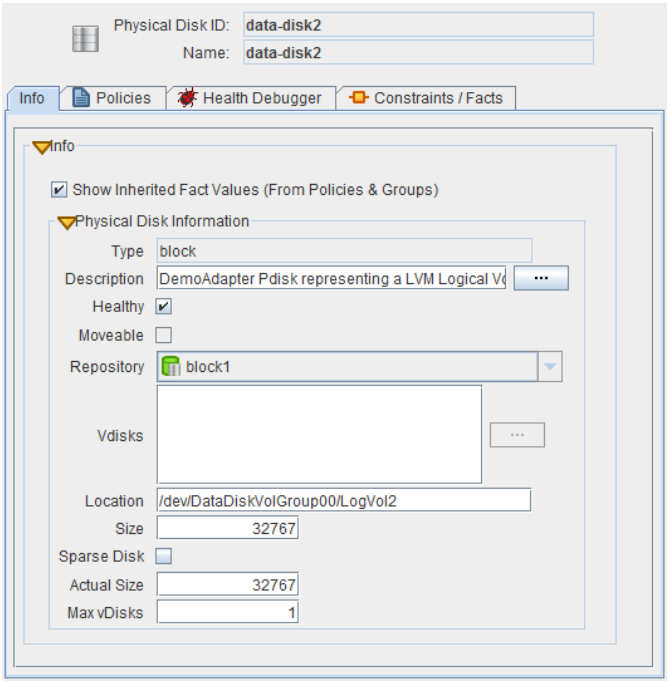

The page that opens under the **Info** tab includes fields where you can configure the general information and attributes (facts) of the pDisk.

NOTE: Whenever you make changes to pDisk object facts, the write icon is superimposed on the object's icon, signifying that the object has been changed. If you want to save the changes you have made, click the **Save** icon on the Orchestration Console toolbar.

This section includes the following additional information:

- "The Physical Disk Info Panel" on page 102
- "The Physical Disk Policies Tab" on page 103
- "The Physical Disk Health Debugger Tab" on page 103
- "The Physical Disk Constraints/Facts Tab" on page 104
- "Physical Disk Object Naming and Renaming" on page 104

The Physical Disk Info Panel

The Physical Disk Info panel on the Info page includes the following fields:

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

Show Inherited Fact Values: Select this check box to show facts with overridden values supplied through attached and/or inherited policies. Such fact values are read only (non-editable).

Type: This field is read-only. It specifies this pDisk as a block device, which is the pDisk type as the VM host sees it:

In the Fact Editor, this fact is listed as `pDisk.type`:

```
<fact name="pDisk.type" value="block" type="String" />
```

Description: Describes the pDisk with any text that you choose.

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

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

Because this is a free form field, you can enter any text you want here. For a pDisk discovered and managed by the vSphere provisioning adapter, this field is populated with a display name label obtained by vSphere and mapped to this pDisk by the Orchestration Server..

Healthy: For a pDisk managed by vsphere, this check box is selected by default, which designates the vDisk as being in good health.

NOTE: We recommend that you do not change the **Healthy** value from its default.

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

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

Moveable: This fact is marked false by default. Do not change this fact.

In the Fact Editor, this fact is listed as `pDisk.moveable`:

```
<fact name="pdisk.moveable" value="false" type="Boolean" />
```

Repository: The storage location (a Volume Group) containing the pDisk image on the VM host. The field is pre-populated with a block type repository name. It cannot be edited.

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

```
<fact name="pdisk.repository" value="vsphere-repo" type="String" />
```

Vdisks: The list of virtual disks that are associated with this physical disk.

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

```
<fact name="pdisk.vdisks">
  <array type="String">
  </array>
</fact>
```

This list box is populated when you attach the pDisk to a VM.

Location: For block type disks, this fact contains the URI to the block device, for example `/mnt/shared/iscsi/`, which could represent an iSCSI share on a VM host.

In the Fact Editor, this fact is listed as `pdisk.location`:

```
<fact name="pdisk.location" value="/dev/vspherevg/vspherelv2" type="String" />
```

Size: The size (measured in MB) of this pDisk image.

In the Fact Editor, this fact is listed as `pdisk.size`:

```
<fact name="pdisk.size" value="32767" type="Integer" />
```

Sparse Disk: This check box is marked true if the image is a sparse (thin) backed vDisk.

In the Fact Editor, this fact is listed as `pdisk.sparse`:

```
<fact name="pdisk.sparse" value="false" type="Boolean" />
```

Actual Size: The actual (sparse) size (measured in MB) of this virtual disk.

In the Fact Editor, this fact is listed as `pdisk.size.actual`:

```
<fact name="pdisk.size.actual" value="8192" type="Integer" />
```

Max vDisks: This fact represents the maximum number of vDisks that are allowed concurrent access to this pDisk. Its value is set to 1 by default. To allow multiple instances of the pDisk, you need to set this value accordingly. A value of -1 indicates that an unlimited number of vDisks have access.

One practical application of this setting is for specifying the pDisks in an OCFS2 environment. If you are using OCFS2, change this value of this fact to match the number of hosts, with each host representing an OCFS2 node.

In the Fact Editor, this fact is listed as `pdisk.vdisks.max`:

```
<fact name="pdisk.vdisks.max" value="1" type="Integer" />
```

In order to allow multiple instances of the pDisk, this value has to be set to some positive number other than 1.

The Physical Disk Policies Tab

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

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

The Physical Disk Health Debugger Tab

The Health Debugger is a common Admin view in the Orchestration Console for most Grid objects. For information about this tool, see Section 14.3, “Health Debugger,” on page 171.

The Physical Disk Constraints/Facts Tab

To support constraining a VM's provision actions based on more than one disk's repository (that is, more than just `resource.vm.repository`), the pDisk can be referenced in constraints. The pDisk constraints are used to assign VM hosts during actions such as provisioning, building, or migrating. You can write constraints against attributes of disks (such as the repository where the pDisk resides) and against the available VM host repositories.

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

IMPORTANT: Several custom facts for the pDisk object can be added at discovery time, and they vary according to the VM technology that manages the respective pDisk.

For more information about using the Fact Editor on this page to rename the Physical Disk object, see Section 6.2.5, "Virtual Disk Object Naming and Renaming," on page 98.

Physical Disk Object Naming and Renaming

Some resource names are generated by the Orchestration Server and can therefore have generic, arbitrary names such as `mysql-pDisk1`, `mysql-pDisk2`, and so on. A Physical Disk (pDisk) you name at creation time might also change later in its purpose or facilities.

The object's display name is visible in the Orchestration Console interface and in optional zos and zosadmin commands. As the number of these pDisk objects grows in your grid, you might find it helpful or necessary to rename them, assigning more meaningful, intuitive names to suit the purpose of the object.

NOTE: Resource object groups (that is, the folders that contain these pDisk objects) can also be renamed. Objects such as jobs, events, and users cannot be renamed.

A pDisk object's name is stored in the `${objectType}.displayname` fact, which exists on every Grid object type, even those objects that cannot be renamed.

You can rename a pDisk object in the Orchestration Console by using one of three methods:

- Right-click the pDisk object in the Explorer tree, then select **Rename** to allow editing of the display name.
- Triple-click the pDisk object in the Explorer tree to allow editing of the display name.
- In the Constraints/Facts page, select the pDisk object `.displayname` fact and then open the Fact Editor to specify a new value for the fact.

As you use one of these methods, notice that the fact value is pre-populated with the `${objectType}.id` fact. This functions as the name value for the object name until you decide to change it.

NOTE: Even after being renamed, the pDisk object retains its associated resource ID in the `.id` fact. This is not editable.

For more information about making the Resource object display names visible from the zos or zosadmin command line, see Chapter 17, “Orchestration Server zos Command Line Tool,” on page 195 and Chapter 18, “Orchestration Server zosadmin Command Line Tool,” on page 243.

7 The Virtual NIC Object

This section includes the following information:

- Section 7.1, “Understanding the Virtual NIC Object,” on page 107
- Section 7.2, “Viewing the Virtual NIC Configuration in the Orchestration Console,” on page 109

7.1 Understanding the Virtual NIC Object

A virtual network interface card (vNIC) represents the configuration of a VM connected to a network. A VM can be configured to have multiple vNICs. When a VM is provisioned, each of its associated vNICs can be attached to a virtual network bridge in order to gain connectivity to a specified network. The vNIC objects are discovered, along with their associated VM, when a Discover VM Images job has been run on a repository.

The vNIC is modeled as a Grid object, located as a subordinate to the VM Grid object in the Explorer tree of the Orchestration Console. A vNIC is given the form of *vmname*_vnic<*n*> where <*n*> is appended to indicate the order of discovery or creation of the vNIC. For example, `redhat_vnic1` would be the name of the first NIC discovered for a VM with the Grid ID `redhat`. Each additional vNIC is incremented by one, so the second vNIC in this example would be named `redhat_vnic2`.

- Section 7.1.1, “The Purpose of the Virtual NIC,” on page 107
- Section 7.1.2, “Creating or Deleting a vNIC in the Orchestration Console,” on page 107

7.1.1 The Purpose of the Virtual NIC

A vNIC represents the network interface configuration for a virtual machine. A vNIC is linked to a network by connecting to a virtual network bridge (vBridge). A group of vBridge objects is represented as a Network group in the Explorer Tree. By convention, during VM host discovery, any vBridges that are configured with the same name are assumed to be part of the same network.

For more information, see Chapter 8, “The Network Group and its Virtual Bridge Objects,” on page 115.

7.1.2 Creating or Deleting a vNIC in the Orchestration Console

Although a vNIC is generally discovered on a VM, you can also manually create or delete a vNIC. This section includes the following information:

- “Creating a Virtual NIC” on page 108
- “Deleting a Virtual NIC” on page 108

Creating a Virtual NIC

You can manually create a vNIC anytime you want to give a VM access to a network configured on the VM host.

To create a vNIC, you can select the VM Grid object in the Explorer tree, then right-click and select **Create Virtual NIC**. You can also use the following alternate method:

1. In the Orchestration Console main menu, select **Actions** > **Create** > **Create Virtual NIC** to display the Create a New Virtual NIC dialog.
2. In the **VM** drop-down list, select the name of the VM to which you want to add a vNIC, then click **Create**.
3. When you have created all of the vNICs you need, click **Close**.
4. Select the newly created vNIC object in the Explorer Tree to view the Info/Groups page of the admin view.
5. On the Info/Groups page, configure the following settings:

   MAC Address: The MAC Address assigned to this vNIC. If this field is left empty, or if it contains an asterisk (*), a MAC address is autogenerated for this vNIC.

   Network: The network (vBridge group) that should be used when provisioning the VM. When a VM is provisioned, any vBridge contained within the specified network group can be used for attaching the VNIC to the network.

   Although the vNIC can be formally created at this point, you can also configure Autoprep or Sysprep facts used to prepare a personalized version of the VM that can be provisioned later.
6. Click the **Save** button to save the fact changes you have made.
7. In the Explorer tree, right-click the VM object to which you added the vNIC, then select **Save Config** to apply the changes to the VM's configuration.

IMPORTANT: If you do not run the **Save Config** action after configuring the vNIC, any vNICs that were added to the grid or vNIC settings that were modified could be lost.

Deleting a Virtual NIC

Although it is uncommon, you might want to manually delete a vNIC when you no longer want the VM to have access to a specified network on the VM host. For example, if the initial need for connecting the VM to a network no longer exists or if the network is going to become private and the VM should not have access, you would delete the virtual NIC that allows connectivity.

To delete a virtual NIC, you can select the vNIC object in the Explorer tree, then right-click and select **Delete**, or you can use the following procedure:

1. In the Orchestration Console main menu, select **Actions** > **Delete > Delete Virtual NIC** to display the Create a New Virtual NIC dialog.
2. In the **Source Objects** list, select the name of the vNIC (hold down the Ctrl key to select multiple vNICs), then click **Add** to move these objects to the **Delete Targets** list.
3. When you have selected all of the vNICs you want to delete, click **Delete** to display the Delete dialog.
4. In the dialog, select **Apply to all** to delete all of the vNIC objects in the **Delete Targets** list, click **OK**, then click **Close**.
5. In the Explorer tree, right-click the VM object where you deleted the vNIC, then select **Save Config** to apply the changes to the VM's configuration.

NOTE: You must run the **Save Config** action to confirm the deletion of the vNIC. If you do not run this action, the vNIC is not deleted from the VM's configuration.

7.2 Viewing the Virtual NIC Configuration in the Orchestration Console

You can visually expose a vNIC Grid object  in the Orchestration Console in two ways:

- In the Explorer Tree, select a VM Resource object, then select the **Info/Groups** tab in the Admin View to open the VM Info Groups page, then scroll to the **Virtual Machine Configuration** panel on that page. You can right-click the vNIC icon in that panel to display the four tabs in the Virtual Disk Admin view.
- In the Explorer Tree, click the expand/collapse icon of a VM Grid object, identify the vNIC icon, then select the icon to display the four tabs in the Virtual NIC Admin view.

The page that opens under the **Info/Groups** tab includes fields where you can configure the general information and attributes (facts) of the vNIC.

NOTE: Whenever you make changes to vNIC object facts, the write icon is superimposed on the object's icon , signifying that the object has been altered. If you want to save the changes you have made, you need to click the **Save** button  on the Orchestration Console toolbar.

This section includes the following additional information:

- Section 7.2.1, "Virtual NIC Info Panel," on page 109
- Section 7.2.2, "Virtual NIC Policies Tab," on page 112
- Section 7.2.3, "Virtual NIC Health Debugger Tab," on page 112
- Section 7.2.4, "Virtual NIC Constraints/Facts Tab," on page 112
- Section 7.2.5, "Virtual NIC Object Naming and Renaming," on page 113

7.2.1 Virtual NIC Info Panel

The Virtual NIC Info panel on the Info/Groups page includes the following sections:

- "Show Inherited Fact Values Check Box" on page 109
- "Virtual NIC Information" on page 109
- "Autoprep/Sysprep Network Adapter" on page 111

Show Inherited Fact Values Check Box

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

Virtual NIC Information

The Virtual NIC Information panel on the Info page includes the following fields:

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

Description: A free-form field you can use to add any description about this vNIC.

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

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

Healthy: In most cases, this check box is selected by default, which designates the vNIC as being in good health.

We recommend that you do not change the **Healthy** value from its default.

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

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

MAC Address: Specifies the MAC address assigned to this vNIC. An empty string implies an auto-generated MAC address, as does an asterisk (*).

When the VM appears on the network, this will be its MAC identifier. A MAC address must be unique on the network to avoid routing conflicts.

There are some situations when you might want to define a static MAC address. For example, if a VM uses DHCP, you might want the DHCP service on the network to give the VM a static address. When the VM boots up and attempts to get an IP address, it contacts the DHCP server, which sees its statically-defined MAC address and then provides the same IP address (not a new one) each time the VM boots up. In this way, the VM is consistently configured with the same IP address. You might also want to define a static MAC address for audit trails or other security reasons.

In the Fact Editor, this fact is listed as `vnic.mac`:

```
<fact name="vnic.mac" value="" type="String" />
```

vBridge: Specifies the name of the actual host bridge used by this vNIC. When the VM is not running, this field is blank. When a VM is provisioned, a vBridge is chosen for this vNIC based on the available VM hosts and the specified network group. When the VM is running, the associated vBridge is identified in this field. For more information, see Chapter 8, “The Network Group and its Virtual Bridge Objects,” on page 115.

In the Fact Editor, this read-only fact is listed as `vnic.vbridge`:

```
<fact name="vnic.vbridge" value="" type="String" />
```

VM: Specifies the name of the VM resource that uses this vNIC.

In the Fact Editor, this read-only fact is listed as `vnic.vm`:

```
<fact name="vnic.vm" value="" type="String" />
```

This is a fact junction referencing the associated VM. Conversely, the `resource.vm.vnics` fact visible from the VM Grid object is a fact junction showing the associated vNICs associated with the VM. A vNIC cannot be shared between two VMs. Each VM has its own vNIC objects.

Network: The network (vBridge group) that should be used by this vNIC when the VM is provisioned. This fact is used in combination with the VM host placement constraints to choose a suitable vBridge at provision time.

In the Fact Editor, this fact is listed as `vnic.network`:

```
<fact name="vnic.network" value="eth1" type="String" />
```

Autoprep/Sysprep Network Adapter

VMs can be prepared for provisioning by configuring the facts in this panel. Click **Define** for each field if the value has not been previously configured.

NOTE: When you change any of the settings in this panel, you need to right-click the VM and select **Personalize** for the changes to take effect.

- **MAC Address:** The MAC address of the interface. Specify an asterisk (*) to generate a new MAC address. If the value is not set, the existing `vnic.mac` is used.

  IMPORTANT: An unset **MAC Address** fact generates a new MAC address. This is contrary to the current tooltip text.

  In the Fact Editor, this fact is listed as `vnic.provisioner.autoprep.MACAddress`:

  ```
  <fact name="vnic.provisioner.autoprep.MACAddress" value="" type="String" />
  ```

- **Use DHCP:** When this check box is selected (it has a value of true), the VM is configured to retrieve its network settings from a DHCP server. If the check box is not selected (it has value of false), you should ensure that the IP address, subnet mask, and gateway address facts are defined.

  In the Fact Editor, this fact is listed as `vnic.provisioner.autoprep.UseDHCP`:

  ```
  <fact name="vnic.provisioner.autoprep.UseDHCP" value="false" type="Boolean" />
  ```

- **IP Address:** The IP address for the adapter.

  In the Fact Editor, this fact is listed as `vnic.provisioner.autoprep.IPAddress`:

  ```
  <fact name="vnic.provisioner.autoprep.IPAddress" value="" type="String" />
  ```

- **Subnet Mask:** The subnet mask for this adapter.

  In the Fact Editor, this fact is listed as `vnic.provisioner.autoprep.subnetMask`:

  ```
  <fact name="vnic.provisioner.autoprep.subnetMask" value="" type="String" />
  ```

- **Gateway IP Addresses:** (Windows sysprep only) A list of the gateway IP addresses available to the interface.

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

  ```
  <fact name="vnic.provisioner.autoprep.Gateways">
    <array type="String">
    </array>
  </fact>
  ```

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

- **DNS from DHCP:** (Optional. SUSE VM only) When this check box is selected (it has a value of true), the SUSE VM is configured to retrieve its DNS server settings from DHCP.

  In the Fact Editor, this fact is listed as `vnic.provisioner.autoprep.DNSFromDHCP`:

  ```
  <fact name="vnic.provisioner.autoprep.DNSFromDHCP" value="false"
  type="Boolean" />
  ```

- **DNS Server IP Addresses:** (Windows VM only) The adapter's list of DNS servers used for name lookup.

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

```
<fact name="vnic.provisioner.autoprep.DNSServers">
  <array type="String">
  </array>
</fact>
```

- **DNS Domain:** (Windows VM only) The adapter's DNS domain name.

  In the Fact Editor, this fact is listed as `vnic.provisioner.autoprep.DNSDomain`:

  ```
  <fact name="vnic.provisioner.autoprep.DNSDomain" value="" type="String" />
  ```

- **Primary WINS Server:** (Windows VM only) The name of the adapter's primary WINS server.

  In the Fact Editor, this fact is listed as `vnic.provisioner.autoprep.primaryWINS`:

  ```
  <fact name="vnic.provisioner.autoprep.primaryWINS" value="" type="String" />
  ```

- **Secondary WINS Server:** (Windows VM only) The name of the adapter's secondary WINS server.

  In the Fact Editor, this fact is listed as `vnic.provisioner.autoprep.secondaryWINS`:

  ```
  <fact name="vnic.provisioner.autoprep.secondaryWINS" value="" type="String" />
  ```

- **NetBIOS:** (Windows VM only) The NetBIOS options for this VM. Options include:
  - **EnableNetBIOSviaDhcp**
  - **EnableNetBIOS**
  - **DisableNetBIOS**

  In the Fact Editor this fact is listed as `vnic.provisioner.autoprep.netBIOS`:

  ```
  <fact name="vnic.provisioner.autoprep.netBIOS" value="" type="String" />
  ```

7.2.2 Virtual NIC Policies Tab

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

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

7.2.3 Virtual NIC Health Debugger Tab

The Health Debugger is a common Admin view in the Orchestration Console for most Grid objects. For information about this tool, see "Health Debugger" on page 171.

7.2.4 Virtual NIC Constraints/Facts Tab

The **Constraints/Facts** tab opens a page that shows all of the effective constraints and facts for a Grid object. Each Grid object has an associated set of facts and constraints that define its properties. In essence, by building, deploying, and running jobs on the Orchestration Server, you can individually change the functionality of any system resources by managing an object's facts and constraints. The

Orchestration Server assigns default values to each of the component facts. Facts with no displayed mode can be changed at any time by the administrator. Facts with mode `r/o` have read-only values, which can be viewed by using the pencil button, but changes cannot be made.

For information about using the Fact Editor on this page to rename the Virtual NIC object, see Section 7.2.5, “Virtual NIC Object Naming and Renaming,” on page 113.

7.2.5 Virtual NIC Object Naming and Renaming

Some resource names are generated by the Orchestration Server and can therefore receive generic, arbitrary names such as `mysql-vnic1`, `mysql-vnic2`, and so on. A Virtual NIC (vNIC) you name at creation time might also change later in its purpose or facilities.

The object’s display name is visible in the Orchestration Console, the Cloud Manager Web Console and Mobile Clients interfaces, and in optional zos and zosadmin commands. As the number of these vNIC objects grows in your grid, you might find it helpful or necessary to rename them, assigning more meaningful, intuitive names to suit the purpose of the object.

NOTE: Resource object groups (that is, the folders that contain these vNIC objects) can also be renamed. Objects such as jobs, events, and users cannot be renamed.

A vNIC object’s name is stored in the `${objectType}.displayname` fact, which exists on every Grid object type, even those objects that cannot be renamed.

You can rename a vNIC object in the Orchestration Console by using one of three methods:

- Right-click the vNIC object in the Explorer tree, then select **Rename** to allow editing of the display name.
- Triple-click the vNIC object in the Explorer tree to allow editing of the display name.
- In the Constraints/Facts page, select the vNIC object `.displayname` fact, then open the Fact Editor to enter a new value for that fact.

As you begin to use one of these methods, you will notice that the fact value is prepopulated with the `${objectType}.id` fact. This functions as the name value for the object name until you decide to change it.

NOTE: Even after being renamed, the vNIC object retains its associated resource ID in the `.id` fact. This is not editable.

For more information about making the Resource object display names visible from the zos or zosadmin command line, see Chapter 17, “Orchestration Server zos Command Line Tool,” on page 195 and Chapter 18, “Orchestration Server zosadmin Command Line Tool,” on page 243.

8 The Network Group and its Virtual Bridge Objects

This section includes the following information:

- Section 8.1, “Understanding the Network Group and Virtual Bridge Objects,” on page 115
- Section 8.2, “Viewing the Virtual Bridge Configuration in the Orchestration Console,” on page 118

8.1 Understanding the Network Group and Virtual Bridge Objects

- Section 8.1.1, “Virtual Bridge Object,” on page 115
- Section 8.1.2, “The Purpose of the Virtual Bridge,” on page 116
- Section 8.1.3, “Creating or Deleting a vBridge in the Orchestration Console,” on page 116
- Section 8.1.4, “Virtual Bridge Object Naming and Renaming,” on page 118

8.1.1 Virtual Bridge Object

In the Cloud Manager Orchestration Server, a group of Virtual Bridge (vBridge) objects is called a “network.” It represents the networks (actually, LANs and VLANs) that are available to that VM host, and can be shared across multiple VM hosts. Network groups are created automatically during VM and VM host discovery as their virtual networking settings are determined. By default, the server automatically groups vBridges with the same name into a corresponding network with the same name, assuming that vBridges with similar names usually refer to the same network. After discovery, you can freely reassign the vBridges to other networks or multiple networks, depending on your organization’s physical network topology.

VLAN Fact

A Virtual LAN (VLAN) provides a “logical” LAN topology for a group of machines that does not need to depend on the switch hardware to which the machines (virtual or physical) are directly connected. Modern “smart” network switches can keep track of a VLAN ID on network traffic passing though them. Using this ID, the switches transparently route such traffic to the hosts configured to use the VLAN specified by the ID.

Using VLANs can reduce the costs of an organization’s physical network infrastructure. For example, if your organization requires multiple subnets or multi-homed hosts, you won’t need to buy new equipment for each new subnet or install multiple physical NICs on the physical machines; this can be done with VLANs using the same common physical network layer. VLANs allow greater flexibility in managing the network topology.

Where possible, the Orchestration Server discovers that a VLAN ID is in use on a network. The Virtual LAN Identifier fact is found on a Network object:

```
<fact name="group.vlanid" value="" type="String" />
```

The fact is populated with a positive integers value upon discovery of an already-existing VM host configuration. The server can track a VLAN ID for each Network object to allow correct management of individual VLANs as full-fledged networks.

The server applies a graphic overlay to the Network icon  to signify that the network was discovered as a VLAN. In cases where VLAN discovery is not accomplished through automation, you can also set the VLAN fact value manually.

8.1.2 The Purpose of the Virtual Bridge

The vBridge is discovered, along with its associated VM, when a discovery job runs on a VM host. A virtual bridge (vBridge) acts as a “virtual” Ethernet segment contained entirely within the software of a VM host. The virtual NIC (vNIC) devices on that host’s VMs can each be assigned to one of the VM host’s vBridges as if they were physically connected to a LAN.

Virtual bridges can also be associated with one or more physical NICs to combine the virtual LANs on these hosts into one common virtual LAN. This combined LAN is referred to as a “network” in the Orchestration Console. Association of virtual bridges is also frequently done to include a “virtual” LAN on a VM host in the organization’s overall physical LAN topology so the VMs can access other systems in the organization as if they were directly connected to the switches.

The following points might help you understand the relationship of these objects:

- Associating a vNIC to a vBridge is like plugging a physical NIC into an Ethernet switch.
- Associating a vBridge with a physical NIC is like connecting two physical switches together using their uplink ports.
- A vBridge can become part of a vLAN by associating it with a physical NIC device that itself is configured as a VLAN.
- A VM host can have multiple vBridges, each of which can be connected to separate physical networks (for example, through 802.1Q VLAN tagging).
- When a VM is provisioned, a virtual NIC must be connected to a virtual bridge in order for the virtual NIC to be usable.

In the Explorer Tree, a vBridge is given the form *vmhostname*_eth*n* where *n* represents the numerical order in which this vBridge was discovered on the VM host, with 0 being appended to the name of the first vBridge discovered or created. For example, `host1_eth0` might be the name of the first bridge. Each additional vBridge is incremented by one, so the second vBridge in this example would be named `host1_eth1`.

8.1.3 Creating or Deleting a vBridge in the Orchestration Console

This section includes the following information:

- “Creating a Virtual Bridge” on page 117
- “Deleting a Virtual Bridge” on page 117

Creating a Virtual Bridge

You might want to manually create a vBridge if, for some reason, the discovery process did not find one of the vBridges in a network. Creating a new vBridge in the Orchestration Console does not “add” a physical bridge to the network, it only helps to model a physical bridge that was not previously discovered.

To create a vBridge, you can select the Network object in the Explorer tree, then right-click and select **New Virtual Bridge**. You can also use the following alternate method:

1. From the Orchestration Console main menu, select **Actions** > **Create > Create Virtual Bridge** to display the Create a New Virtual Bridge dialog.
2. In the **Source Groups** list, select the Network object where you want to add a vBridge, then click **Add** to move it to the **Target Groups** list.
3. In the **New Virtual Bridge Name** field, specify the name you want to use to identify this vBridge. The data you enter here is completely free form. When the Orchestration Server discovers and names vBridges, it associates them by name as the default practice. This is done for convenience in locating the objects. A vBridge can be named anything, provided that it is a legal name for a vBridge on the VM host’s operating system.

   It is not necessary to add the `<vmhostname>_` prefix on the name. This prefix is automatically prepended when the new object is created. For example, if you select bridge name `sales` on VM host `vmh1`, the actual object is created as `vmh1_sales`.
4. Click **Create**, then click **Close**.
5. In the Explorer tree, expand the Network group where you created the new vBridge object.
6. Select the new vBridge object to open its Info/Groups admin view.
7. Configure the settings described in Section 8.2.1, “Virtual Bridge Info/Groups Tab,” on page 119.
8. Click the **Save** button to save the fact changes you have made.

Deleting a Virtual Bridge

Although it is uncommon, you might want to manually delete a vBridge when you no longer want the VM to have access to a specified network on the VM host. For example, if the initial need for connecting the VM to a network no longer exists or the network is going to become private and the VM should not have access you would delete the vBridge that allows connectivity.

To delete a virtual bridge, you can select the vBridge object in the Explorer tree, then right-click and select **Delete**, or you can use the following procedure:

1. In the Orchestration Console main menu, select **Actions** > **Delete > Delete Virtual Bridge** to display the Create a New Virtual Bridge dialog.
2. In the **Source Objects** list, select the name of the vBridge (hold down the Ctrl key to select multiple), then click **Add** to move these objects to the **Delete Targets** list.
3. When you have selected all of the vBridges you want to delete, click **Delete** to display the Delete dialog.
4. In the dialog, select **Apply to all** to delete all of the vBridge objects in the **Delete Targets** list, click **OK**, then click **Close**.

8.1.4 Virtual Bridge Object Naming and Renaming

Some resource names are generated by the Orchestration Server and can therefore receive generic, arbitrary names such as `host1-eth1`, `host2-eth1`, and so on. A Virtual Bridge (vBridge) you name at creation time might also change later in its purpose or facilities.

The object's display name is visible in the Orchestration Console, the Cloud Manager Web Console and Mobile Clients interfaces, and in optional zos and zosadmin commands. As the number of these vBridge objects grows in your grid, you might find it helpful or necessary to rename them, assigning more meaningful, intuitive names to suit the purpose of the object.

NOTE: A Network object (that is, the group that contains these vBridge objects) can also be renamed. Objects such as jobs, events, and users cannot be renamed.

A vBridge object's name is stored in the `${objectType}.displayname` fact, which exists on every Grid object type, even those objects that cannot be renamed.

You can rename a vBridge object in the Orchestration Console by using one of three methods:

- Right-click the vBridge object in the Explorer tree, then select **Rename** to allow editing of the display name.
- Triple-click the vBridge object in the Explorer tree to allow editing of the display name.
- In the Constraints/Facts page, select the vBridge object `.displayname` fact and then open the Fact Editor to enter a new value for that fact.

As you use one of these methods, you will notice that the fact value is prepopulated with the `${objectType}.id` fact. This functions as the name value for the object name until you decide to change it.

NOTE: Even after being renamed, the vBridge object retains its associated resource ID in the `.id` fact. This is not editable.

For more information about making the Resource object display names visible from the zos or zosadmin command line, see Chapter 17, "Orchestration Server zos Command Line Tool," on page 195 and Chapter 18, "Orchestration Server zosadmin Command Line Tool," on page 243.

8.2 Viewing the Virtual Bridge Configuration in the Orchestration Console

You can visually expose a vBridge Grid object  in the Explorer tree by selecting a Network object to expand it, then selecting the vBridge you want to view in the Info/Groups page.

The page that opens under the **Info/Groups** tab includes fields where you can configure the general information and attributes (facts) of the vBridge.

NOTE: Whenever you make changes to vBridge object facts, the write icon is superimposed on the object's icon, signifying that the object has been altered. If you want to save the changes you have made, you need to click the **Save** button  on the Orchestration Console toolbar.

8.2.1 Virtual Bridge Info/Groups Tab

The page that opens under the **Info/Groups** tab includes two collapsible sections where you can configure the general information and attributes of the vBridge.

- "Virtual Bridge Info Panel" on page 119
- "Virtual Bridge Groups Panel" on page 120

Virtual Bridge Info Panel

The **Info** panel on the **Info/Groups** page includes the following information:

- "Show Inherited Fact Values Check Box" on page 119
- "Network Information" on page 119

Show Inherited Fact Values Check Box

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

Network Information

The **Network Information** panel on the Info/Groups page for the vBridge Grid object includes the following fields:

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

Description: Enter a description of the vBridge Grid object.

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

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

Vbridge Enabled: This check box is selected by default. When it is selected (it has a value of true), the vBridge is enabled and Virtual NICs can be attached to it.

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

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

Healthy: This check box is selected by default. When it is selected (it has a value of true), the virtual bridge is designated as being in good health. You can set the health of the object by selecting or deselecting the health check box. Changing the value in this way has an immediate effect unless the value is overridden by an attached policy (this follows the normal rules of policy inheritance). For more information, see Chapter 14, "Grid Object Health Monitoring," on page 169.

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

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

Attached Virtual NICs: This list box lists the Virtual NICs that are attached to this vBridge.

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

```
<fact name="vbridge.vnics">
  <array type="String">
  </array>
</fact>
```

The list includes vNICs from currently running VMs only. When a VM is not running, the Orchestration Server does not consider its vNICS as “currently attached.” This is because vNICs are configured to associate with Network objects, and the attached vBridge might change, depending on which host the VM has been provisioned. The list can also change dynamically if the VM is migrated to another host (and that host’s vBridge) on the same network.

Virtual Bridge Groups Panel

This section of the **Info/Groups** page lists the groups of vBridge objects (called “Networks”) in the grid to which this vBridge is associated. Click **Choose** to open the Network Selection dialog. In this dialog, you can choose which networks to display in the Explorer tree by selecting a group and then clicking **Add** or **Remove** to move it to or from the **Source Networks** list.

8.2.2 Virtual Bridge Policies Tab

The **Policies** tab opens a page that contains a policy viewer for each of the policies associated with a Grid object. Click **Choose** to associate an existing policy with a vBridge.

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

8.2.3 Virtual Bridge Health Debugger Tab

The Health Debugger is a common Admin view in the Orchestration Console for most Grid objects. For information about this tool, see Section 14.3, “Health Debugger,” on page 171.

8.2.4 Virtual Bridge Constraints/Facts Tab

The **Constraints/Facts** tab opens a page that shows all of the effective constraints and facts for a Grid object. Each Grid object has an associated set of facts and constraints that define its properties. In essence, by building, deploying, and running jobs on the Orchestration Server, you can individually change the functionality of any system resources by managing an object’s facts and constraints. The Orchestration Server assigns default values to each of the component facts. Facts with no mode can be changed at any time by the administrator. Facts with mode `r/o` have read-only values, which can be viewed by using the edit pencil button, but changes cannot be made.

9 The Repository Object

Repositories are storage areas for VM image files and VM template files.

If a VM's files are stored on a particular host server, that VM must be run from that host server. If a VM's files are stored on a shared repository, that VM can be run on any host server that has access to the shared repository.

Host servers can have multiple repositories associated with them, and some repository types can be associated with multiple host servers as shared repositories. A host server can be associated with repositories stored locally on its server and with shared repositories stored on other machines.

The default size for all repositories is unlimited. To control disk space usage, you can change this default. For more information, see Capacity (MB): in "Repository Information Subpanel" on page 122.

The repository groups and their constituent repository objects are displayed in the Explorer panel and the accompanying Repository Admin View of the Cloud Manager Orchestration Console.

- Section 9.1, "Right-Click Menu Actions on the Repository Object," on page 121
- Section 9.2, "Repository Groups," on page 122
- Section 9.3, "Repository Info/Groups Tab," on page 122
- Section 9.4, "Repository Policies Tab," on page 129
- Section 9.5, "Repository Health Debugger Tab," on page 129
- Section 9.6, "Repository Constraints/Facts Tab," on page 129
- Section 9.7, "The Repository Action History Tab," on page 129
- Section 9.8, "Repository Object Naming and Renaming," on page 130
- Section 9.9, "Shared Storage for Disk Images," on page 130

9.1 Right-Click Menu Actions on the Repository Object

The Repository object displayed in the Explorer tree has three available actions in the right-click menu:

- **Discover VM Images:** Use this action on a newly discovered or newly created repository to populate it with the VMs residing in the mapped datastore location.
- **Discover Disks:** Use this action to discover ISO disks stored in this repository.
- **Delete:** Use this action to delete the Repository object from the Explorer tree in the Orchestration Console.

  Before you delete the Repository object, you need to delete all the VMs contained in that repository; otherwise, an error message is displayed.
- **Rename:** Use this action to rename the Repository object. For more information, see Section 9.8, "Repository Object Naming and Renaming," on page 130.

9.2 Repository Groups

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

9.3 Repository Info/Groups Tab

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

- Section 9.3.1, “Info Panel,” on page 122
- Section 9.3.2, “Best Practices for Entering Repository File Paths,” on page 127
- Section 9.3.3, “Groups,” on page 129

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

9.3.1 Info Panel

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

- ““Show Inherited Fact Values” Check Box” on page 122
- “Repository Information Subpanel” on page 122
- “SAN Adapter Configuration” on page 127

“Show Inherited Fact Values” Check Box

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

Repository Information Subpanel

The Repository Information panel on the Info/Groups page includes the following fields:

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

Description: The nature or purpose of this repository.

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

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

Repository Enabled: This check box is selected by default. When it is selected (it has a value of true), VMs can be moved to this repository or they can be provisioned from it.

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

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

Healthy: This check box is selected by default. When it is selected (it has a value of true), the repository is designated as being in good health. You can set the health of the object by selecting or deselecting the health check box. Changing the value in this way has an immediate effect unless the value is overridden by an attached policy. For more information, see Chapter 14, “Grid Object Health Monitoring,” on page 169

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

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

Type: Select the repository type for this Repository object by selecting an option from the drop-down list.

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

```
<fact name="repository.type" value="local" type="String" />
```

The following table includes information about the various repository types:

Table 9-1 *Repository Types and Descriptions*

Repository Type	Description	Orchestration Console Icon
`local`	The default repository on a host server. Each host server starts with its own local repository, which has the same name as the server’s Resource Grid object.	
`NAS` (Network Attached Storage)	Represents a NAS device connected to host servers (for example, NFS mount). This NAS device must be mounted and available on all host servers associated with this Repository Grid object.	
`SAN` (Storage Area Network)	Represents an iSCSI or Fibre Channel SAN. Currently supported only with the vsphere provisioning adapter.	
`datagrid`	The shared datagrid repository (named `zos`) is located on the Orchestration Server and is accessible to all host servers in the datagrid. By default, each host server has access to the `zos` datagrid repository.	
`virtual`	Represents an externally managed virtual disk (for example, Amazon EC2).	

IMPORTANT: If you have a vSphere environment with an iSCSI datastore based on an ESX 3.5 (or previous) host, the Orchestration Console incorrectly displays its type as **local** rather than **SAN**. This misrepresentation affects the accuracy of the VM host assignment (how the repositories are scored in the plan), and could possibly affect VM migration validation.

To work around this issue, set the type to **SAN**. You need to check that this setting is retained when another VM host or Repository discovery is executed.

SAN ID: (SAN repositories only) The SAN ID (the Virtual Fabric ID) for this repository.

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

```
<fact name="repository.id" value="test1" type="String" />
```

Root Location: The repository's logical root location. You can also think of this as the base location for all VM files and subdirectories contained within this repository.

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

```
<fact name="repository.location" value="/" type="String" />
```

The table below provides some examples you can consider as you enter a shared root path in this field. For more information, see "Best Practices for Entering Repository File Paths" on page 127.

Table 9-2 *Repository Types and Root Location Examples*

Repository Type	Root Location Examples
`local`	• `/` (`root`) • `c:/vm`
`NAS` (Network Attached Storage)	This is the mount point that is assumed to be the same on every host server with a connection to this NAS. • `/u` `/mnt/`*`myshareddisk`*
`SAN` (Storage Area Network)	Not required.
`datagrid`	`grid:///vms`
`virtual`	• `/` (`root`) • `c:/vm`

VM Config Search Path: The relative path (from `repository.location`) to be used during discover of VM configuration files. This fact also implicitly includes the `resource.preferredpath` fact.

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

```
<fact name="repository.searchpath">
  <array>
    <string>/vmfiles/vms</string>
  </array>
</fact>
```

The ··· button opens the Attribute element values dialog, where you can add, remove, or edit the path (element values) in an array of path choices.

The table below provides some examples you can consider as you enter a search path in this field.

Table 9-3 *Repository Types and VM Config Search Path Examples*

Repository Type	VM Config Search Path Examples
`local`	`/vmfiles/vms`
`NAS` (Network Attached Storage)	Each of these is the relative path from the location to search for VM configuration files. Null specifies to search the whole mount. ◆ *`my_vms`* ◆ *`saved_vms`* ◆ null (no path entry)
`SAN` (Storage Area Network)	Not required.
`datagrid`	`grid:///vms`
`virtual`	

Preferred Storage Path: The relative path (from `repository.location`) where you want the Orchestration Server to place the VM files after a move or a clone operation.

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

```
<fact name="repository.preferredpath" value="" type="String" />
```

IMPORTANT: If you use this field, do not include a leading forward slash (/) in the path. For more information, see “Best Practices for Entering Repository File Paths” on page 127.

Table 9-4 *Repository Types and Preferred Storage Path Examples*

Repository Type	Preferred Storage Path Examples
`local`	`/vmfiles/vms` for VMs
`NAS` (Network Attached Storage)	*`my_vms`*
`SAN` (Storage Area Network)	Not required.
`datagrid`	`grid:///vms`
`virtual`	

Disk Discovery Path: The directories to search for disk image files

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

```
<fact name="repository.disks.paths">
  <array type="String">
  </array>
</fact>
```

Disk Discovery Patterns: The patterns used to discover disk image files.

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

```
<fact name="repository.disks.patterns">
  <array>
    <string>*.iso</string>
  </array>
</fact>
```

Capacity (MB): The maximum amount (measured in MB) of storage space on the repository. The default (-1) designates an unlimited amount of space.

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

```
<fact name="repository.capacity" value="-1" type="Integer" />
```

Used Space (MB): The amount (measured in MB) of storage space used for VMs.

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

```
<fact name="repository.usedspace" value="0" type="Integer" />
```

Free Space (MB): The amount (measured in MB) of storage space available to new VMs. The value is always set to `-1`, which designates an unlimited amount of space.

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

```
<fact name="repository.freespace" value="-1" type="Integer" />
```

Actual Used Space (MB): The actual amount (measured in MB) of storage space used for VMs.

In the Fact Editor, this fact is listed as `repository.usedspace.actual`:

```
<fact name="repository.usedspace.actual" value="0" type="Integer" />
```

Actual Free Space (MB): The actual amount (measured in MB) of storage space available to new VMs. A value of -1 means unlimited.

Efficiency: Enter an efficiency coefficient that the Orchestration Server uses to calculate the cost of moving VM disk images to and from the repository. This value is multiplied by the disk image size (in MB) to determine an efficiency score. A score of zero (0) means no cost (very efficient).

NOTE: A fact not visible in the UI (except in the Fact Editor) is `repository.capacity.set`, a Boolean flag used by provisioning adapters to indicate that the repository capacity has been discovered.

Stored VMs: The VM images stored in this repository. The list is aggregated from individual VM facts.

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

```
<fact name="repository.vmimages">
  <array type="String">
  </array>
</fact>
```

This fact designated as readonly and is not editable.

Compatible VM Hosts: The VM hosts capable of using this repository. The list is aggregated from individual VM facts.

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

```
<fact name="repository.vmhosts">
  <array>
    <string>vmhost123_vshpere</string>
  </array>
</fact>
```

You can edit this array by clicking the ··· button to open the Attribute element values dialog, where you can add, remove, or edit the VM host IDs (element values) in an array of VM host ID choices.

Accessed By Provision Adapters: The provisioning adapter jobs that are allowed access to VMs on this repository.

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

```
<fact name="repository.provisioner.jobs">
  <array>
    <string>vsphere123</string>
  </array>
</fact>
```

You can edit this array by clicking the ··· button to open the Choose Grid Objects dialog, where you can add or remove provisioning adapters for the array of provisioning adapter choices.

NOTE: In the Fact Editor, you edit the provisioning adapter array by using the **Attribute Element Values** dialog.

SAN Adapter Configuration

SAN Adapter Vendor: (SAN repositories only) The name of the vendor of the SAN.This should be adapter specific, such as `iqn`, `npiv`, `emc`. An empty field (that is, no value in the string) indicates that bind/unbind is a no-op (no operation performed).

In the Fact Editor, this fact is listed as `repository.san.vendor`:

```
<fact name="repository.san.vendor" value="" type="String" />
```

SAN Transport: (SAN repositories only) From the drop-down list, select **iSCSI** or **Fibre Channel** to indicate the type of SAN transport this repository uses.

In the Fact Editor, this fact is listed as `repository.san.type`:

```
<fact name="repository.san.type" value="" type="String" />
```

9.3.2 Best Practices for Entering Repository File Paths

Use the following guidelines in scenarios where you need VM repositories.

- “Creating a Repository to Use with New VMs” on page 128
- “Creating a Repository to Use with Existing VMs” on page 128
- “Creating a Repository for Existing VMs with Shared Root Locations and Separate Configuration Directories” on page 128

Creating a Repository to Use with New VMs

If you are creating a repository for new VMs that you will eventually provision:

1. In the **Root Location** field, specify the location for the new repository.

   Example: `/vms_new`

2. In the **Preferred Storage Path** field, specify the path to your image file store (relative to the root location path). This becomes the path for VM configuration files and VM image files when you associate a VM with this repository.

   Example: `images` (no leading forward slash)

   Because the fields are concatenated, the provisioning adapter searches for the existing VM files in `/vms_new/images`.

Creating a Repository to Use with Existing VMs

Use this procedure when you already have VMs in your grid and a store for the VM configuration and disk image files already exists.

1. In the **Root Location** field, specify the shared location for this repository.

   Example: `/vms_new`

2. In the **VM Config Search Path** field, specify the search path to your existing configuration file store (relative to the root location path).

   Example: `old_config` (no leading forward slash)

   Because the fields are concatenated, the provisioning adapter searches for the existing VM configuration files in `/vms_new/old_config`.

3. In the **Preferred Storage Path** field, specify the path to your existing image file store (relative to the root location path). This also becomes the path for VM configuration files and VM image files when you associate a VM with this repository.

   Example: `all_images` (no leading forward slash)

   Because the fields are concatenated, the provisioning adapter searches for the existing VM files in `/vms_new/all_images`.

Creating a Repository for Existing VMs with Shared Root Locations and Separate Configuration Directories

Use this procedure when you want to create a repository for existing VMs that have a shared root path but separate configuration file directories such as `/vms_new/old_config1` and `/vms_new/old_config2`).

1. In the **Root Location** field, specify the shared location for this repository.

   Example: `/vms_new`

2. In the **VM Config Search Path** field, specify the search paths to your existing configuration file store (relative to the Root Location path).

   Example: Adjacent to the **VM Config Search Path** field, click , click **Add element**, enter `old_config1` (no leading forward slash), click **OK**, click **Add element** again, specify `old_config2` (no leading forward slash), then click **OK**.

   Because the fields are concatenated, the provisioning adapter searches for the existing VM configuration files in the array consisting of `/vms_new/old_config1` and `/vms_new/old_config2`.

3 In the **Preferred Storage Path** field, specify the path to your existing image file store (relative to the root location path).This path also becomes the path for VM configuration files and VM image files after a move or clone when a VM has been associated with this repository.

Example: `all_images` (no leading forward slash)

Because the fields are concatenated, the provisioning adapter searches for the existing VM files in `/vms_new/all_images`.

9.3.3 Groups

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

9.4 Repository Policies Tab

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

NOTE: You can edit a policy by right-clicking a policy icon, selecting **Edit Policy,** then clicking the **Save** button.

9.5 Repository Health Debugger Tab

The Health Debugger is a common Admin view in the Orchestration Console for most Grid objects. For information about this tool, see Section 14.3, “Health Debugger,” on page 171.

9.6 Repository Constraints/Facts Tab

The **Constraints/Facts** tab opens a page that shows all of the effective constraints and facts for a Grid object. Each Grid object has an associated set of facts and constraints that define its properties.

By building, deploying, and running jobs on the Orchestration Server, you can individually change the functionality of any system resources by managing an object’s facts and constraints. The Orchestration Server assigns default values to each of the component facts, although they can be changed at any time by the administrator, unless they are read-only. Facts with mode `r/o` have read-only values, which can be viewed by using the pencil icon, but changes cannot be made.

9.7 The Repository Action History Tab

The **Action History** tab is displayed in the administrative view of the Repository object. When you select the **Action History** tab, a table displays a list of the history for all actions performed on this Grid object.

The Orchestration Server must be connected to an audit database for the **Include Audit Database** check box to be available. If the **Include Audit Database** check box is selected in this view, the action status is not polled. Click the **Refresh** button in the toolbar to retrieve and display fresh data.

For more details about the information listed on the Action History page, see "Action History in Admin Views of the Orchestration Console" in the *Cloud Manager Procedures Guide*.

9.8 Repository Object Naming and Renaming

Some resource names are generated by the Orchestration Server and can therefore have generic, arbitrary names such as `host1`, `host2`, `host3`, and so on. Repositories you name at creation time might also change later in their purpose or facilities.

The object's display name is visible in the Orchestration Console, the Cloud Manager Web Console and Mobile Clients interfaces, and in optional zos and zosadmin commands. As the number of these Repository objects grows in your grid, you might find it helpful or necessary to rename them, assigning more meaningful, intuitive names to suit the purpose of the object.

NOTE: Repository object groups (that is, the folders that contain these Repository objects) can also be renamed. Objects such as jobs, events, and users cannot be renamed.

A Repository object's name is stored in the `${objectType}.displayname` fact, which exists on every Grid object type, even those objects that cannot be renamed.

You can rename a Repository object in the Orchestration Console using one of three methods:

- Right-click the Repository object in the Explorer tree, then select **Rename** to allow editing of the display name.
- Triple-click the Repository object in the Explorer tree to allow editing of the display name.
- In the Constraints/Facts page, select the Repository object `.displayname` fact and then open the Fact Editor to enter a new value for that fact.

As you use one of these methods, you will notice that the fact value is pre-populated with the `${objectType}.id` fact. This functions as the name value for the object name until you decide to change it.

NOTE: Even after being renamed, the Repository object retains its associated resource ID in the `.id` fact. This is not editable.

For more information about making the Resource object display names visible from the zos or zosadmin command line, see Chapter 17, "Orchestration Server zos Command Line Tool," on page 195 and Chapter 18, "Orchestration Server zosadmin Command Line Tool," on page 243.

9.9 Shared Storage for Disk Images

The Orchestration Server can discover disk files (such as ISOs) inside storage represented by a Repository Grid object, making it easier to specify a disk file that you want to attach to a VM.

This section includes the following information:

- Section 9.9.1, "Setting Disk Discovery Facts," on page 131
- Section 9.9.2, "Running the Discovery," on page 131
- Section 9.9.3, "Sharing Disks Between VMs," on page 132
- Section 9.9.4, "Attaching a Discovered Disk to a VM," on page 132
- Section 9.9.5, "Using Attached Disks in the Guest OS," on page 133

9.9.1 Setting Disk Discovery Facts

Before the Orchestration Server can discover disks, you need to set two facts in the Repository object. These facts are listed in the table below.

Table 9-5 *Required Facts for ISO Disk Discovery*

Fact Name	Purpose
`repository.disks.paths`	The directories to search for disk image files
`repository.disks.patterns`	The patterns used to discover disk image files (the default pattern is `*.iso`).

You can use the built-in array editor in the Orchestration Console to modify these facts. For example, to modify the `repository.disks.paths` fact,

1. In the Orchestration Console, select a Repository object that represents your shared storage.
2. Select the **Info/Groups** tab to open **Info** page in the admin view.
3. In the **Disk Discovery Path** field, click the array editor button ··· to open the **Attribute element values** dialog.

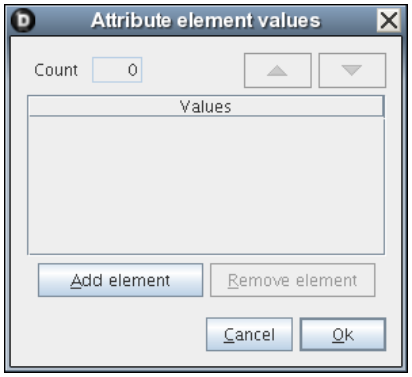

4. On the dialog, click **Add element** to open the editor where you can add a string value to the `repository.disks.paths` fact. For example, you could enter `/data/isos` as a possible search directory for disk image files.
5. Click the **Save** icon on the main toolbar to save your fact changes.

A similar array editor is available for changing the values in the repository.disks.patterns fact found in the **Disk Discovery Patterns** field of the Repository **Info** page of the Orchestration Console.

9.9.2 Running the Discovery

The **Discover Disks** action is included in the right-click menu of a Repository Grid object and in the **Provision** main menu of the Orchestration Console. The right-click action lets you discover disks on a single repository, while the action under the **Provision** menu lets you select multiple repositories for disk discovery.

When you select this action, a provisioning adapter function runs to search for and retrieve ISO disk filenames stored on the selected repository or repositories. If the selected provisioning adapter does not support the **Discover Disks** action, the option is dimmed in the menu.

NOTE: The Orchestration Server can discover any disk, but its default is to search for filenames with a `.iso` suffix.

After discovery, the disk filenames are stored in the `repository.disks` fact of the Repository object, located in the **Stored disks** field of the Repository's **Info/Groups** page in the admin view.

Figure 9-1 *The Disk Discovery Portion of the Info/Groups Page for the Repository Object*

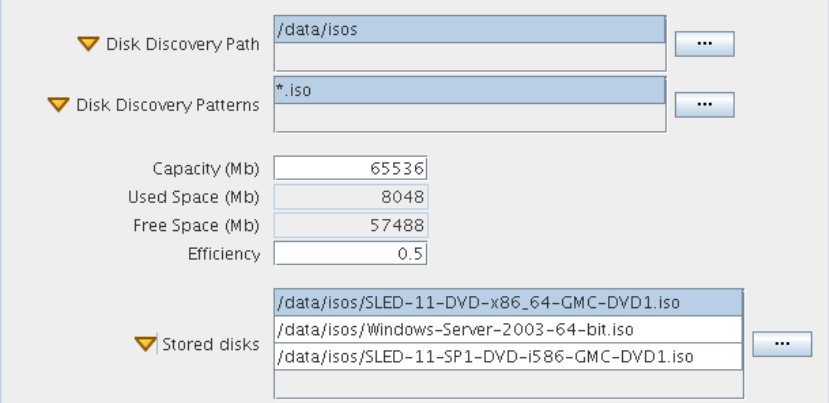

This is an array of disks, each of which is a valid value for the `vdisk.location` vDisk fact and stored in this repository. You can edit the filenames of these disks or add new filenames as values in the array.

9.9.3 Sharing Disks Between VMs

To share ISOs between VMs, you would need to store the ISOs in a shared location that is accessible on all VM hosts where those VMs are to run. In such a scenario, you would set up a directory on all VM hosts where `/data/isos/` is a mount point to some SAN or NFS storage of ISO files. For example, an ISO file might have this path: `/data/isos/SLES-11-SP1-DVD-x86_64-DVD1.iso`.

9.9.4 Attaching a Discovered Disk to a VM

You can attach a discovered disk to a VM.

1. In the Orchestration Console, select the VM to which you want to attach a disk, then right-click to display the available actions for the VM.
2. Select **Create Virtual Disk** to create a new disk. The Info/Groups page for the disk is displayed in the admin view.

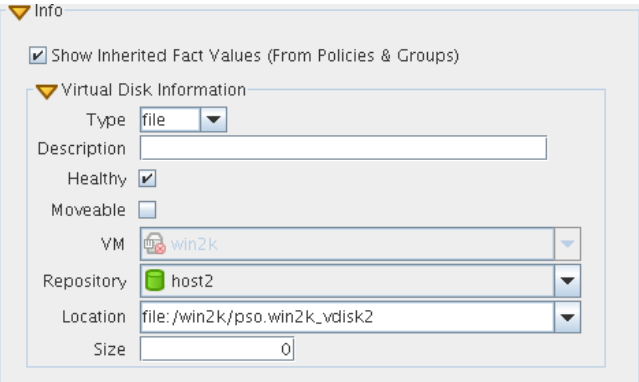

The **Location** drop-down menu lists the ISO file names that were previously discovered in the Repository associated to this vDisk.

3 In the **Location** field (that is, the `vdisk.location` fact), select the disk that you want to attach to the VM. You can also manually edit the filename of the disk you want to attach.

4 Click the **Save** icon on the main toolbar to save your changes.

5 Commit the changes to the VM configuration.

5a (If the VM is running) In the Explorer tree, right-click the VM object, select **Apply Config.**

5b (If the VM is not running) In the Explorer tree, right-click the VM object, select **Save Config**.

If the VM state or the provisioning adapter do not support either of these actions, the options are dimmed.

9.9.5 Using Attached Disks in the Guest OS

After the disk is attached to a VM, you might need to take additional steps for the disk to be usable.

Linux Guest OS: Use the `mount` command to make the disk visible to the guest OS. Use the `umount` command to unmount the disk prior to changing or deleting the underlying PSO vDisk object.

Windows Guest OS: If the disk is not immediately visible, you need to reboot your guest OS. In general, the you can assume that disk is visible if it is replacing an already-mounted disk that is visible in a CD-ROM device.

10 The User Object

A User object represents an individual account that is allowed to connect to the Cloud Manager Orchestration Server. Administrator users are also allowed to connect by using the `zosadmin` command line and the Orchestration Console user interfaces.

You can use the Orchestration Console user interface to manually create a User object. You can create objects automatically if authentication through LDAP or Active Directory is enabled, or optionally if auto-registration is configured.

The user object icon and the red square user object icon.

- Section 10.1, “User Groups,” on page 135
- Section 10.2, “User Info/Groups Tab,” on page 135
- Section 10.3, “User Policies Tab,” on page 141
- Section 10.4, “User Health Debugger Tab,” on page 141
- Section 10.5, “User Constraints/Facts Tab,” on page 142
- Section 10.6, “The User Action History Tab,” on page 142

10.1 User Groups

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

10.2 User Info/Groups Tab

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

- Section 10.2.1, “Info,” on page 136
- Section 10.2.2, “Groups,” on page 141

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

10.2.1 Info

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

- “Show Inherited Fact Values Check Box” on page 136
- “User Information” on page 136
- “Personal Information” on page 137
- “Job Information” on page 138
- “Accounting Information” on page 138
- “Job Control” on page 139
- “Quota Information” on page 140

Show Inherited Fact Values Check Box

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

User Information

The User Information panel on the Info/Groups page includes the following fields:

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

Account Enabled: This check box is selected by default.When the check box is selected, the user is allowed to log in and run jobs.

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

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

Online: When this check box is selected (it has a value of true), the user is currently logged in to the server.

In the Fact Editor, this fact is listed as `user.online`:

```
<fact name="user.online" value="false" type="Boolean" />
```

Healthy: This check box is selected by default. When it is selected (it has a value of true), the user is designated as being in good health. You can set the health of the object by selecting or deselecting the health check box. Changing the value in this way has an immediate effect unless the value is overridden by an attached policy. For more information, see Chapter 14, “Grid Object Health Monitoring,” on page 169

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

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

External Groups: A list of external groups (for example, LDAP) that this user belongs to.

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

```
<fact name="user.external.groups">
  <array type="String">
  </array>
</fact>
```

You can edit this array by clicking the ··· button to open the Attribute element values dialog, where you can add, remove, or edit the name and value for every user environment you want to use.

Personal Information

First Name: The user's first name.

In the Fact Editor, this fact is listed as `user.name.first`:

```
<fact name="user.name.first" value="" type="String" />
```

Last Name: The user's last name.

In the Fact Editor, this fact is listed as `user.name.last`:

```
<fact name="user.name.last" value="" type="String" />
```

Password: The user's hashed login password.

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

```
<fact name="user.password" value="kNLj1_Fc96C3ajVXcqQEGZBRrbivgxhhzK3TKLpP"
type="String" />
```

IMPORTANT: The System User password should not be changed. If you try to change this password by using the Orchestration Console, the password does not actually change and no warning is posted that the change was unsuccessful. There is no indication of authentication failure until you try to log in with the "new" password.

EMail: The user's email address.

In the Fact Editor, this fact is listed as `user.name.email`:

```
<fact name="user.name.email" value="" type="String" />
```

City: The name of the city where the user is located.

In the Fact Editor, this fact is listed as `user.location.city`:

```
<<fact name="user.location.city" value="" type="String" />
```

State: The name of the state or province where the user is located.

In the Fact Editor, this fact is listed as `user.location.state`:

```
<fact name="user.location.state" value="" type="String" />
```

Country: The name of the country where the user is located.

In the Fact Editor, this fact is listed as `user.location.country`:

```
<fact name="user.location.country" value="" type="String" />
```

Site: The name of the site (for example, a campus or building) where the user works.

In the Fact Editor, this fact is listed as `user.location.site`:

```
<fact name="user.location.site" value="" type="String" />
```

Environment: A list of default user environment variable names and values that the Orchestration Server sets when executing joblets remotely.

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

```
<fact name="user.env">
  <dictionary>
    <dictelement key="dfadsafd">
      <string>safdaf</string>
    </dictelement>
  </dictionary>
</fact>
```

Job Information

Total Job Count: The total number of jobs that this user has historically initiated on this Orchestration Server.

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

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

Active Jobs: The number of top-level jobs run with this user account that are in an active state.

In the Fact Editor, this fact is listed as `user.jobs.active`:

```
<fact name="user.jobs.active" value="0" type="Integer" />
```

Queued Jobs: The number of top-level jobs run with this user account that are currently in a queued state.

In the Fact Editor, this fact is listed as `user.jobs.queued`:

```
<fact name="user.jobs.queued" value="0" type="Integer" />
```

Total Jobs: The total number of top-level jobs run by this user account.

In the Fact Editor, this fact is listed as `user.jobs.total`:

```
<fact name="user.jobs.total" value="0" type="Integer" />
```

Active Sessions: The number of currently active sessions (that is, connections) that the user has established with the Orchestration Server.

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

```
<fact name="user.sessions" value="1" type="Integer" />
```

Accounting Information

Total Spending: The total cost of computing resources by this user.

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

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

Average Spending Rate: The computed moving average spending (in dollars per hour) over the last hour of activity for this user.

In the Fact Editor, this fact is listed as `user.account.spendrate`:

```
<fact name="user.account.spendrate" value="-0.0006" type="Real" />
```

Maximum Spending Rate: An amount (in dollars per hour) to be used by the Resource Scheduler to throttle the rate at which computing cycles are consumed by the user. A value of less than or equal to zero (<=0) turns the feature off.

In the Fact Editor, this fact is listed as `user.account.maxspendrate`:

```
<fact name="user.account.maxspendrate" value="0.0000" type="Real" />
```

Default Accounting Group: Lets you select the default User Group to be billed for work conducted by this user.

In the Fact Editor, this fact is listed as `user.accounting.group`:

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

Total Wall Time: The total amount of wall time (in seconds) consumed by this user.

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

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

Total Grid Time: The total amount of grid time (in gcycles, which is a normalized average of compute cycles) consumed by this user.

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

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

Job Control

Default Priority Value: A numerical representation of the default priority at which this user's job runs, with 1 being the lowest priority and 9 being the highest priority.

In the Fact Editor, this fact is listed as `user.priority.default`:

```
<fact name="user.priority.default" value="7" type="Integer" />
```

Default Priority: The string representation of the default priority at which this user can run a job. The value is matched to the integer value in `user.priority.default`.

In the Fact Editor, this fact is listed as `user.priority.default.string`:

```
<fact name="user.priority.default.string" value="high" type="String" />
```

Maximum Priority Value: A numerical representation of the maximum priority at which this user's job can run, with 1 being the lowest priority and 9 being the highest priority. Only the system user can run jobs at priority 10.

In the Fact Editor, this fact is listed as `user.priority.max`:

```
<fact name="user.priority.max" value="5" type="Integer" />
```

Datagrid Maximum History: The maximum number of job instance directories that should be kept in the datagrid for this user.

In the Fact Editor, this fact is listed as `user.datagrid.maxhistory`:

```
<fact name="user.datagrid.maxhistory" value="25" type="Integer" />
```

Job Preemption Enabled: Select this check box if you want to allow the user to preempt jobs that have a priority less than the priority of the running job instance.

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

```
<fact name="user.preemption.enabled" value="false" type="Boolean" />
```

Max Preemption Priority: The highest job priority band from which this user is allowed to preempt resources. The value acts as a delta from the current job instance priority. The maximum preemptible priority is always less than or equal to `user.priority.max`.

In the Fact Editor, this fact is listed as `user.preemption.priority.delta`:

```
<fact name="user.preemption.priority.delta" value="0" type="Integer" />
```

Resources Stealing Enabled: Select this check box to allow the user to steal resources that are running jobs that have a priority less than the priority of the running job instance.

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

```
<fact name="user.stealing.enabled" value="false" type="Boolean" />
```

Max Stealing Priority: The highest job priority band from which this user is allowed to steal resources. The value acts as a delta from the current job instance priority, and must be less than zero (<0).

In the Fact Editor, this fact is listed as `user.stealing.priority.delta`:

```
<fact name="user.stealing.priority.delta" value="-1" type="Integer" />
```

Privileged Job Groups: A list of Job Groups with jobs and joblets that this user is allowed to run on resources that have reached their slot maximum or that are provisioned resources that are reserved for another user or job.

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

```
<fact name="user.privilegedjobgroups">
  <array type="String">
  </array>
</fact>
```

You can edit this array by clicking the ··· button to open the Choose Grid Objects dialog, where you can add or remove Job Groups in an array of choices. The Job Groups can be added to or removed from a list of Source Grid Objects to a list of Target Grid Objects (or vice versa).

Quota Information

Account Balance Remaining: The balance (measured in dollars) that remains available for this user since the last reset. You can use this value to implement quotas on your server.

In the Fact Editor, this fact is listed as `user.account.balance`:

```
<fact name="user.account.balance" value="0.0000" type="Real" />
```

Job Counter: The number of jobs this user has initiated since the last reset. You can use this value to implement quotas on your server.

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

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

Time Remaining: The amount of wall time (measured in seconds) remaining for use by this user since last the reset. You can use this value to implement quotas on your server.

In the Fact Editor, this fact is listed as `user.account.time`:

```
<fact name="user.account.time" value="0" type="Integer" />
```

Grid Time Remaining: The amount of grid time (measured in gcycles) remaining for use by this user since last the reset. You can use this value to implement quotas in your grid.

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

```
<fact name="user.account.gcycles" value="0" type="Integer" />
```

10.2.2 Groups

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

10.3 User Policies Tab

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

You can modify policies by using the Policy Grid object. For more information, see Section 11.1, “Policy Object,” on page 143.

If you click **Choose** on the **Policy** tab, a Policy Selection dialog is launched, where you can add or remove individual policies to be applied to the selected User Grid object.

Figure 10-1 *The Policy Selection dialog*

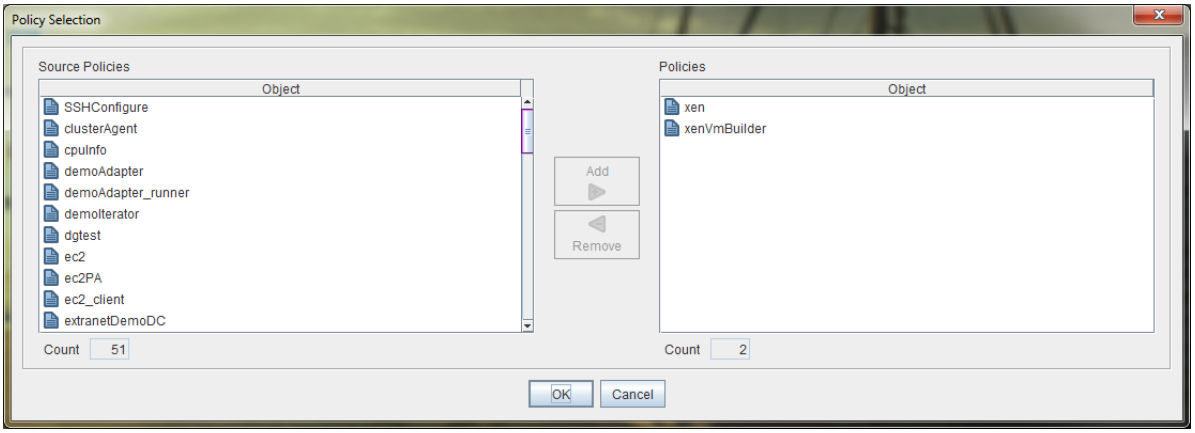

10.4 User Health Debugger Tab

The Health Debugger is a common Admin view in the Orchestration Console for most Grid objects. For information about this tool, see Section 14.3, “Health Debugger,” on page 171.

10.5 User Constraints/Facts Tab

The **Constraints/Facts** tab opens a page that shows all of the effective constraints and facts for a Grid object. Each Grid object has an associated set of facts and constraints that define its properties.

By building, deploying, and running jobs on the Orchestration Server, you can individually change the functionality of any system resources by managing an object's facts and constraints. The Orchestration Server assigns default values to each of the component facts. Facts with no displayed mode can be changed at any time by the administrator. Facts with mode `r/o` have read-only values, which can be viewed by using the pencil button, but changes cannot be made.

10.6 The User Action History Tab

The **Action History** tab is displayed in the administrative view of the User object. When you select the **Action History** tab, a table displays a list of the history for provisioning actions performed on this Grid object (assuming that it is provisionable, for example, a VM or VM template).

The Orchestration Server must be connected to an audit database for the **Include Audit Database** check box to be available. If the **Include Audit Database** check box is selected in this view, the action status is not polled. Click the **Refresh** button in the toolbar to retrieve and display fresh data.

For more details about the information listed on the Action History page, see "Action History in Admin Views of the Orchestration Console" in the *Cloud Manager Procedures Guide*.

11 Miscellaneous Objects Displayed in the Explorer Tree

The Explorer panel (tree) of the Orchestration Console displays some miscellaneous objects of importance:

- Section 11.1, “Policy Object,” on page 143
- Section 11.2, “Computed Fact Objects,” on page 143
- Section 11.3, “Event Objects,” on page 144
- Section 11.4, “Metrics Objects,” on page 144

11.1 Policy Object

XML is used to define Orchestration Server policies. A policy can be deployed to the server and associated with any grid object. The `policy` element is the root element for policies. Policies contain constraints and fact definitions for grid objects.

You can edit a policy by clicking its icon in the Explorer tree view, making your changes, then clicking the save icon. The **Where Used** tab of the Policy Editor lists the jobs where the selected policy is associated.

11.1.1 Policy Constraints

A policy can define a collection of constraints that are applied appropriately based on context. For example, a resource constraint can limit the selection of a resource to a subset based on resource group membership, or any number of other fact-based evaluations.

You can use `/opt/novell/zenworks/zos/server/examples/customNode.policy` on the Orchestration Server as an example policy file with constraints.

11.1.2 Policy Facts

The XML fact element defines a fact to be stored in the grid object’s fact namespace. The name, type and value of the fact are specified as attributes. For list or array fact types, the element tag defines list or array members. For dictionary fact types, the `dict` tag defines dictionary members.

You can see an example policy with an XML representation for all the fact types on the Orchestration Server at `/opt/novell/zenworks/zos/server/examples/allTypes.policy`.

11.2 Computed Fact Objects

Computed facts are used when you want to run JDL to generate the value for a fact. Although computed facts are not jobs, they use the same JDL syntax. You can see examples of computed facts on the Orchestration Server at `/opt/novell/zenworks/zos/server/examples/activejobs.cfact` and `/opt/novell/zenworks/zos/server/examples/repostiory.cfact`.

11.3 Event Objects

An Event object in the Explorer tree represents a user-described set of rules that can be associated with a schedule trigger or handled by long-running jobs written to respond to events.

For more information about using events, see “Event Triggers” on page 151 and Section 14.2, “Health Events,” on page 171.

11.4 Metrics Objects

The Orchestration Server Metrics Facility collects, aggregates, and allows simple fact-based retrieval of metric values by jobs and computed facts (via JDL), policy constraints, and Event triggers on a per-resource basis.

A Metrics object is deployed in the Explorer tree. Use the right-click menu to display the “deploy” and “undeploy” actions. Pre-defined `.metric` files are located in the `/opt/novell/zenworks/zos/server/components/metrics` folder, or you can create a new `.metric` file and paste in an XML file.

NOTE: You can also use the zosadmin deploy command to deploy a `.metric` file. For example:

```
zosadmin deploy load_one.metric
```

Metrics objects are listed by their deployment name, which may or may not be the same as the name of the actual metric. This potentially allows multiple, separately deployable, RRD definitions for a single “instantaneous” metric, with different aggregation periods defined.

12 The Orchestration Server Job Scheduler

You can use the Job Scheduler in the Orchestration Server to automatically start deployed jobs on your grid by using either time or event triggers.

You can think of the functionality provided by the time triggers as being similar to a distributed `cron` system (in fact, time triggers can be described in `cron` syntax). This triggering, coupled with the job control functions in the Orchestration Server, allows for the sophisticated automation of routine data center tasks.

For example, suppose you want to periodically harvest a large log file in a coordinated way from a farm of several hundred machines. First, you could create an Orchestration Server job that uses the datagrid for file movement. The job control options specify that the job should run on not more than three machines at once and sweep across the entire grid. You would then create a schedule to run this job at the desired interval.

As another example, you could use the Job Scheduler to trigger a discovery job every time a new resource is added to the grid. In this case, the job developer writes the discovery job to discover and set facts about the resource. Next, you would create a schedule to run this job on the `RESOURCE_ONLINE` built-in trigger. In fact, this type of triggered job is currently used in the standard set of deployed discovery jobs to detect specific resource CPU and OS information.

Yet another example would be to run a job on server startup that sends a notification email to an administrator.

- Section 12.1, “Understanding the Job Scheduler View,” on page 145
- Section 12.2, “Understanding the Job Schedule Editor,” on page 148
- Section 12.3, “Understanding Cron Syntax in the Job Scheduler,” on page 157

12.1 Understanding the Job Scheduler View

This section includes information to help you understand the functions of the Job Scheduler and how to use it to launch Orchestration Server jobs.

- Section 12.1.1, “Navigating The Job Schedules Table,” on page 146

Click **Scheduler** on the main toolbar of the Orchestration Console to open the Job Scheduler view.

Figure 12-1 *Job Scheduler View of the Orchestration Console*

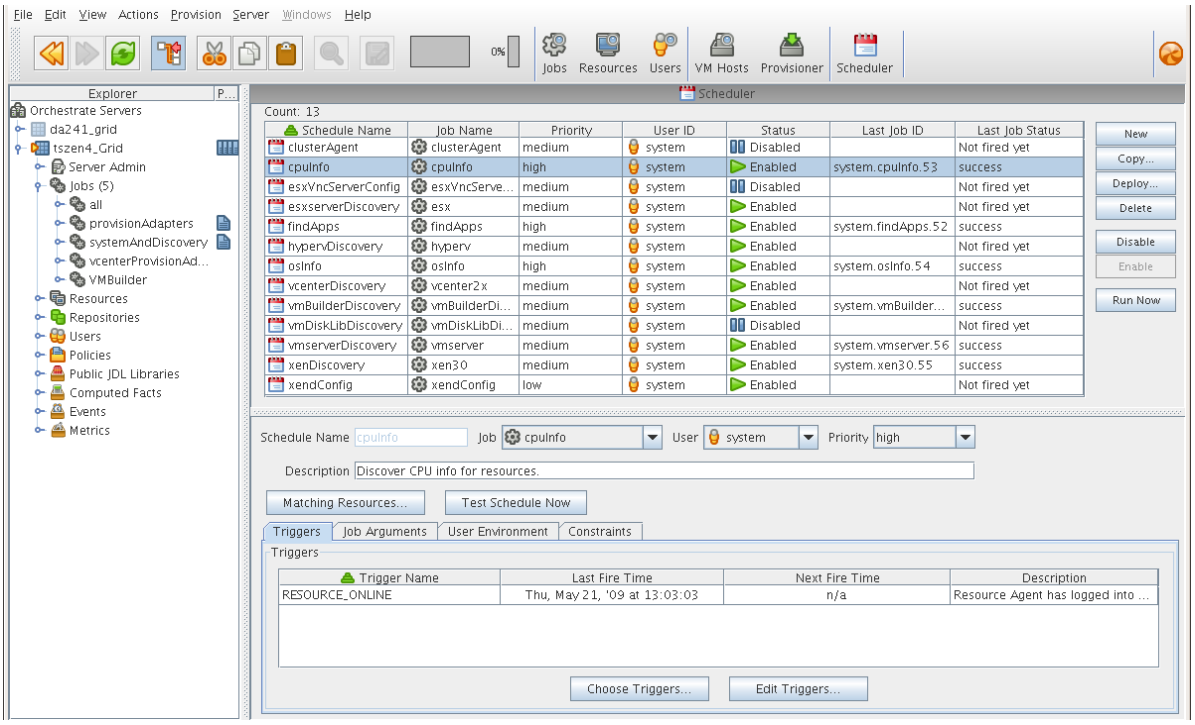

12.1.1 Navigating The Job Schedules Table

The Orchestration Server includes several predefined and predeployed discovery jobs that have predefined launch schedules. Among these jobs are the `cpuinfo`, `findapps`, `osinfo`, and other jobs, depending on the options (that is, the server profile) you chose and the configuration you used during the installation. After installation, these jobs are listed by name in a table in the Job Scheduler view.

Figure 12-2 *The Job Schedules Table in the Job Scheduler View*

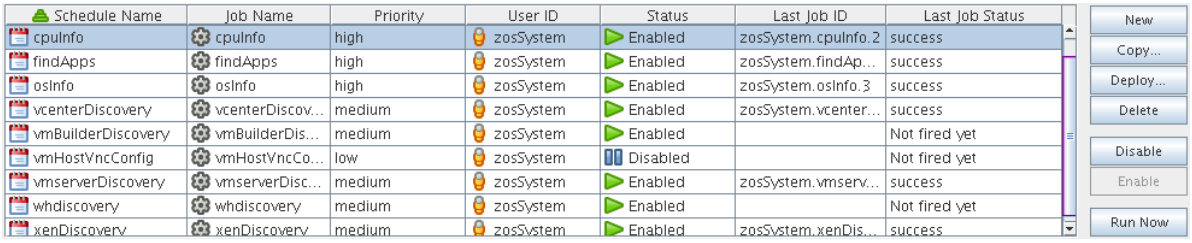

By default, the server uses schedule names that are similar to the job name so that schedules are easy to match (although this is not required). The schedules list shows all of the existing job schedules that accompany predefined jobs, along with the schedules that you create in the Job Scheduler.

NOTE: The Job Scheduler view is not a real-time monitor view of jobs, so if a job attribute (for example, Last Job Status or Last Fire Time) has changed, it might not be displayed until you click **Refresh**.

The Job Schedules Table has functionality that lets you decide how you want to display information about the job schedules:

- You can drag any column in the table to move it left or right in the table according to your preference.
- You can mouse over any column heading in the table to view tooltip text about the purpose of the data in that column.
- You can right-click any column heading in the table to open the Job Scheduler Column Editor dialog.

Figure 12-3 *Job Scheduler Column Editor dialog*

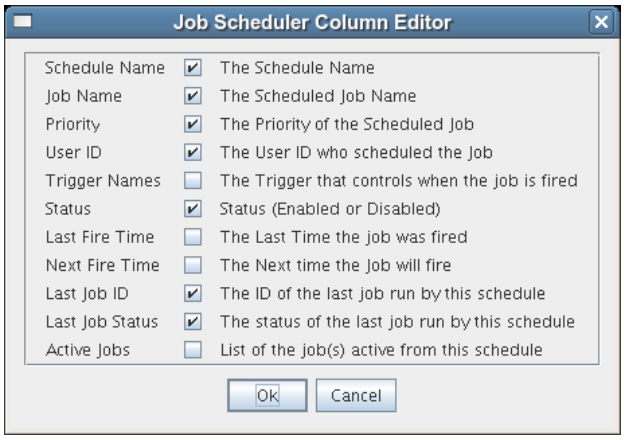

You can select any column heading in this dialog to display it in the Job Schedules Table. The columns display the attributes of a previously configured job schedule. As the figure shows, this dialog also includes text that clarifies the purpose of the data in each column.

In the Job Scheduler view, there are seven function buttons next to the Job Schedules Table (see Figure 12-2 on page 146) that let you take action on any schedule you select inside the table. (Only one schedule at a time can be selected.)

- **New:** Opens a dialog where you can create a new schedule. When you create a new schedule, the Job Scheduler adds a new line to the Job Schedules Table. When the new line is added, you can use the Job Schedule Editor to edit the attributes of the schedule. A new schedule must be given a unique schedule name.

  The Job Scheduler forces a new schedule to be created in the **Disabled** state to prevent it from running while it is being defined. You click **Enable** when a job is ready to be used.
- **Copy:** Copies a schedule you have selected in the Job Schedules Table. Clicking this button opens a dialog where you rename the copy. If you want to create a schedule that is similar but not identical to an existing schedule, use this button to save time in adding attributes to a job schedule configuration. A copy of a schedule must be given a unique schedule name.
- **Deploy:** Opens a dialog where you can select a schedule (that is, a deployable `.sched` file) to deploy.
- **Delete:** Deletes the selected schedule from the Job Schedules Table. You cannot recover a deleted job schedule.

NOTE: Deleting a schedule that was deployed as part of a `.job` or `.sar` displays a confirmation dialog. Deleting the schedule undeploys all contents of the `.job` or `.sar` that contains the schedule.

- **Disable:** Disables the selected schedule in the Job Schedules Table. The jobs associated with the schedule are not re-run, but any currently running instances of this job continue to run.
- **Enable:** Enables a disabled job schedule.
- **Run Now:** Forces the specified schedule to run immediately. This updates statistics such as **Last Fire Time**.

Removed Jobs or Users: Scheduler Behavior

If a job or a user is undeployed or removed from the Orchestration Server, the Job Schedules Table continues to list the schedule previously associated to that removed grid object, but the removed grid object no longer displays the icon that represents the object (job or user).

Figure 12-4 *User Object and Job Object Icons Are Not Displayed*

Schedule Name	Job Name	Priority	User ID	Status	Last Job ID	Last Job Status
CpuDiscovery	cpuInfo	high	zosSystem	Disabled		Not fired yet
clusterAgent	quickie	medium	zosSystem	Disabled		Not fired yet
findApps	findApps	high	zosSystem	Enabled	zosSystem.findAp...	success
osInfo	osInfo	medium	system	Enabled	zosSystem.osInfo.48	success
vcenterDiscovery	vcenterDiscov...	medium	zosSystem	Enabled	zosSystem.vcenter...	success
vmHostVncConfig	vmHostVncCo...	low	zosSystem	Disabled		Not fired yet
vmserverDiscovery	vmserverDisc...	medium	zosSystem	Enabled	zosSystem.vmserv...	success
xenDiscovery	xenDiscovery	medium	zosSystem	Enabled	zosSystem.xenDis...	success

In the preceding figure, the **CpuDiscovery** schedule displays no Job icon for the **cpuInfo** job in the schedules table. Even though the job has been undeployed, the schedule is still listed.

In the **osinfo** schedule, the **system** user has no User icon. That user has been removed from the server.

If you choose a new user or job to be associated with a schedule, a deleted or undeployed user or job is never displayed in the pop-up menu for that schedule again.

12.2 Understanding the Job Schedule Editor

The Job Schedule Editor is located immediately below the Job Schedules Table in the Job Scheduler view.

Figure 12-5 *The Job Schedule Editor in the Job Scheduler View*

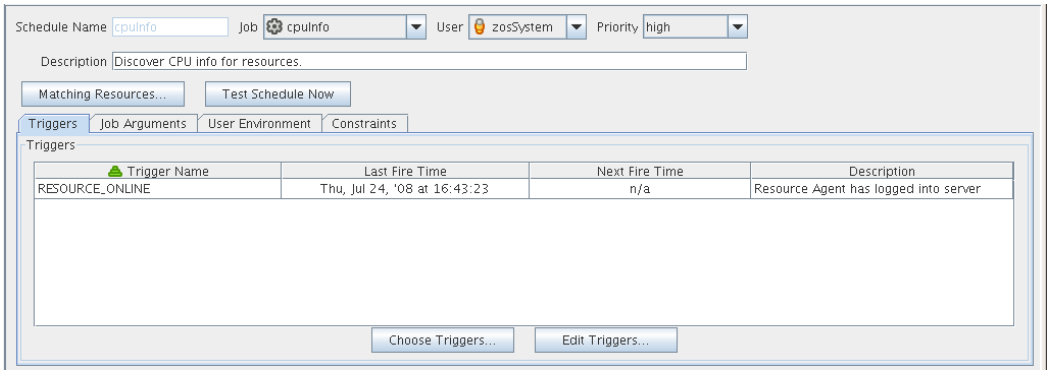

There are several times when you use this part of the Job Scheduler tool:

- When you create a new schedule by clicking **New**.

- When you modify the attributes of an existing schedule (available when you select a schedule in the table).
- When you create a copy of an existing schedule by clicking **Copy**.

The Job Schedule Editor lets you create or modify a job schedule by specifying its attributes.

You use the following controls and data when you create or modify a job schedule:

- Section 12.2.1, “Schedule Name,” on page 149
- Section 12.2.2, “Job,” on page 149
- Section 12.2.3, “User,” on page 149
- Section 12.2.4, “Priority,” on page 150
- Section 12.2.5, “Description,” on page 150
- Section 12.2.6, “Matching Resources,” on page 150
- Section 12.2.7, “Test Schedule Now,” on page 150
- Section 12.2.8, “Triggers,” on page 150
- Section 12.2.9, “Job Arguments,” on page 155
- Section 12.2.10, “User Environment,” on page 156
- Section 12.2.11, “Constraints,” on page 156

12.2.1 Schedule Name

When you create a new schedule, the unique name you specify is displayed in this field. If you select a schedule from the Job Schedules Table, the name of the schedule is displayed in this field. The field is not editable, because schedules cannot be renamed after they are created. (You can use a copy if this is required.)

12.2.2 Job

When you create a new schedule, you need to associate a deployed job with it. You can select the job you want to run from this drop-down list.

If you want to use a previously created schedule to run a different job, you can change the job here.

12.2.3 User

When you create a new schedule, you need to associate a user with it. The user represents the user for whom the job will run. The choice of user might affect the permissions, privileges and constraints of the job. You can select the user from this drop-down list.

If you want a different user to run a job on a previously created schedule, you can change the user here.

If you decide to change the user who runs the job, check the **Priority** field to ensure that the priority you want is selected.

12.2.4 Priority

When you create a new schedule and associate a job and a user with it, a list of possible run priorities becomes available in this drop-down list. The list of priorities varies, depending on the user that is specified in the previous field. In this field, you select the priority of the job that is to be run so that if other jobs are to start concurrently or are competing for resources, the Orchestration Server can determine which job takes priority.

12.2.5 Description

For predeployed jobs, this field contains a default description of what the job's schedule does. The field is editable, so you can enter a description of your own for job schedules that you create.

12.2.6 Matching Resources

This button displays a list of resources where the job runs now or where it could run. This list is useful for checking the context of constraints that might have been affected by a choice of user or by manually specifying additional constraints under the **Policy** tab. The list is also useful to verify that a discovery job (that is, one that is triggered by the **Run on Resource Start** option) runs on the preferred set of machines.

12.2.7 Test Schedule Now

Click this button to test the new or modified schedule you are working with. The test runs the new or modified schedule without permanently saving the current configuration of the schedule or recording statistics. This control differs from the **Run Now** control in the Job Schedules Table, which runs a saved (persisted) schedule, disregarding any unsaved modifications that have been made to it in the Job Schedule Editor.

12.2.8 Triggers

When you click the **Triggers** tab in the Job Scheduler view, the following page opens:

Figure 12-6 *The Schedule Triggers Page in the Job Scheduler*

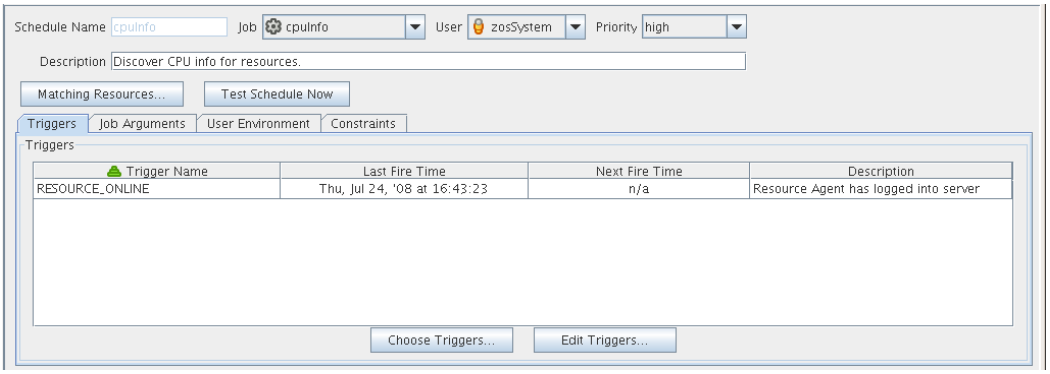

In this view, you can add or define the triggers you want to associate with job schedules. A trigger is the signal to the Job Scheduler to initiate, or "fire" a schedule at a given time or at the occurrence of a given event. Job Scheduler triggers can be classified with regard to two conditions: events and time.

The Triggers table on this page has functionality that lets you decide how you want to display information about the triggers:

- You can drag any column in the table to move it left or right in the table according to your preference.
- You can mouse over any column heading in the table to view tooltip text about the purpose of the data in that column.
- You can right-click any column heading in the table to open the Triggers table column editor dialog.

Figure 12-7 *Trigger Table Column Editor dialog*

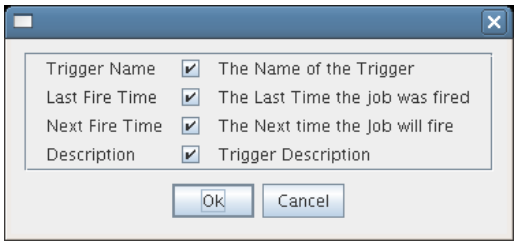

You can select any column heading in this dialog to display it in the Triggers table. The columns display the attributes of a previously configured Triggers table. As the figure shows, this dialog also includes text that clarifies the purpose of the data in each column.

You can create as many triggers as you want to meet any scheduling situation you might have. Multiple time triggers can be associated with a schedule and multiple schedules can use the same trigger. The triggers you create are retained by the Job Scheduler for you to choose from when you create a schedule for a job. The currently associated triggers are displayed in the list along with a description.

Choose Triggers

This button opens a dialog where you can choose both predefined and user defined time triggers to associate with this job schedule.

In this dialog, you can click **Add** to move a selected trigger to the active, scheduled triggers that are to be associated with this job schedule. You can also click **Remove** to unassociate a trigger.

When a trigger is moved to the scheduled list, it becomes associated to the job schedule and it is displayed in the Job Scheduler view.

Most example jobs in the Orchestration Server are associated with event triggers. The dialog can also list other job schedule triggers that are based on time.

Event Triggers

An Event trigger is the signal to the Job Scheduler to initiate, or “fire” a job when a given event occurs. An Event can be one of three types:

- **Event objects:** These objects are user defined events that are fired when an event rule is triggered. If an event object is deployed, it automatically shows in the trigger chooser as a possible choice.

- **Built-in events:** These events are system-wide events such as a resource coming online or a changing resource health condition. Built-in events are always available as a trigger choice. The Job Scheduler has eight possible built-in event triggers:
  - AGENT_VERSION_MISMATCH
  - RESOURCE_ONLINE
  - REPOSITORY_HEALTH
  - RESOURCE_HEALTH
  - SERVER_UP
  - USER_HEALTH
  - USER_ONLINE
  - VMHOST_HEALTH

  You can select any combination of these event triggers for a single schedule.

  The first trigger, AGENT_VERSION_MISMATCH, triggers the job when an Orchestration Agent of an incompatible version attempts to connect to this Orchestration Server. It can be used to initiate a configuration management tool for an agent software update or the job could email a message to an administrator to report the incompatible agent. The other seven available built-in event triggers are listed with accompanying descriptions in the dialog.
- **External events:** These events are fired by an outside process. These are not automatically shown as choices in the trigger chooser, but must be defined by the trigger editor.

An event trigger can be used in conjunction with a time trigger to allow flexibility in scheduling the job application for maximum effectiveness or convenience. Jobs triggered by events require that their job arguments contain a dictionary named context. For example, your event-triggered job should have this jobarg element in its policy:

```
<policy>
   <jobargs>
      <fact name="context" type="Dictionary"
            description="Dictionary containing the context for the event" />
   </jobargs>
</policy>
```

The key/values of the dictionary are dependent on the event type. For event objects, the jobargs.context dictionary contains the matching context of the triggered rule. For built-in events, the jobargs.context dictionary contains the key of the object type corresponding to the built-in event and the object ID that caused the event.

For example, if the USER_ONLINE event triggers because the user named foo logs in, the jobargs.context dictionary contains:

```
{ user : foo }
```

Likewise, if the RESOURCE_ONLINE event is triggered because the resource agent named “vmhost1” comes online, the jobargs.context dictionary contains:

```
{ resource : vmhost1 }
```

For the AGENT_VERSION_MISMATCH event, the jobargs.context dictionary contains more information, as shown in the following table:

Table 12-1 *Dictionary Information*

Key	Type
AgentBuild	Long
AgentIP	String
AgentId	String
AgentMajor	Integer
AgentMinor	Integer
AgentPoint	Integer
JavaMajor	Integer
JavaMinor	Integer
JavaPoint	Integer
JavaVendor	String
JavaVersion	String
OsMajor	Integer
OsMinor	Integer
OsName	String
OsPoint	Integer
OsVendor	String
OsVersion	String
SystemArch	String
UsingJRE	Boolean
resource	String

Time Triggers

A time trigger is the signal to the Job Scheduler to initiate, or “fire” a job when a prescheduled time occurs. A time trigger can be used in conjunction with an event trigger to allow flexibility in scheduling the job application for maximum effectiveness or convenience. No default time triggers are defined in the Job Scheduler. You need to create new time triggers by clicking Edit Triggers.

Edit Triggers

Click **Edit Triggers** to open the Triggers dialog.

The following controls and information are available in the dialog:

- **New:** Opens a secondary dialog where you can create a new time trigger name. When you create the trigger name, the attribute fields in the Triggers dialog are cleared and you can specify new attributes for the trigger. A new trigger must be given a unique trigger name.

- **Copy:** Lets you modify an existing time trigger by giving it a new name and attributes. This can be helpful if there are only slight differences in the new attributes. A copy of a trigger must be given a unique trigger name.
- **Deploy:** Opens a file chooser where you can choose an existing, stored trigger (that is, a `.trig` file) to deploy.
- **Delete:** Deletes a selected time trigger.

IMPORTANT: Deleted triggers are not recoverable. If the trigger is used by existing schedules, it is removed from all of those schedules when it is deleted.

- **Trigger Name:** Specifies the unique name of the trigger you are creating or modifying. This name is displayed in the Job Scheduler if you choose to associate this trigger with a schedule. After you create the trigger name, it cannot be modified.
- **Description:** Specifies a description for the time trigger you are creating or modifying. The description is optional and can be as detailed as you want.

  If the number of characters in the description string exceeds the space in the **Description** field, a button is enabled that opens a string editor when clicked.
- **Save:** Clicking this button saves the defined time trigger and its attributes.
- **Fire Starting In:** Displays multiple fields specifying the time increment and frequency to be used by the trigger to fire the job. If you select this type of time trigger, the **Fire using CRON Expression** button becomes inactive.

NOTE: You can use the **Fire Starting In** control to create either a “one-shot” time trigger or a “reoccurring” time trigger.

A one-shot time trigger fires just once after a specified period of time. To specify a one-shot trigger, click **Fire Starting in**, specify the amount of time before firing, then specify 0 as the time to **Repeat Indefinitely**.

A reoccurring time trigger fires after a specified period and then either fires repeatedly for an indefinite number of times or it fires for a specified number of times. To specify a reoccurring, indefinite trigger, click **Fire Starting in**, specify the amount of time before firing, then select **Repeat Indefinitely**. To specify a reoccurring but finite trigger, click **Fire Starting in**, specify the amount of time before firing, select **Repeat Range**, then specify the number of times you want the trigger to fire.

- **Fire using CRON Expression:** Specifies the cron expression that enables the job to fire automatically at a specified time or date. You need to be familiar with cron to use this field.

  The **Examples** list box of selected cron expressions and their associated descriptions is located just below this button. You can use a listed expression as is, or use it as a template to modify the expression to meet your needs.

  If you select this type of time trigger, the **Fire Starting In** and the **Fire Using Event** buttons become inactive.

  For an example of how a cron expression can be implemented in a trigger, see “Creating and Assigning a Time Trigger for the New Schedule” in the *Cloud Manager Procedures Guide*. For detailed information about cron syntax, see “Understanding Cron Syntax in the Job Scheduler” on page 157.

- **Fire Using Event:** Specifies a deployed event or an external event that enables the job to fire when a specified event occurs. Deployed events are defined using an XML syntax. You can specify a deployed event from **Events** (that is, listed in the **Events** drop down list) or you can enter the name of an external event.

  If the number of characters in the **Fire Using Event** description string exceeds the space in the field, a ... button is enabled that opens a string editor when clicked.

12.2.9 Job Arguments

This tab displays an area in the lower left corner of the Job Schedule Editor where possible job arguments are listed. If you select an existing schedule in the Job Schedules Table, any optional job arguments (jobargs) for the associated job are displayed in this area.

Figure 12-8 *The Job Arguments Area of the Job Scheduler View*



The jobargs are defined by the deployed job. Some jobs might already have a default value displayed, but others must have values specified in order for the job to be able to run.

IMPORTANT: Job arguments displayed in blue are required. You must supply data in the accompanying fields.

A job argument defines the values that can be passed in when a job is invoked. These values let you statically define and control job behavior. To learn more about each job argument, mouse over each jobarg line to display tooltip text.

The Job Scheduler uses the values you enter into the fields of this area to build a jobargs namespace in the policy for this job.

Each job argument has an accompanying **Lock** check box. When **Lock** is not selected, the accompanying job argument uses the default value specified in the job's policy. When **Lock** is selected, the value specified in the field is locked down and overrides the default value in the policy. A locked value continues to be used even if the policy value is modified.

You can click **Restore Jobargs** to restore job arguments to the values specified in the job policy. This function removes any changes you might have specified in the Job Scheduler and deselects all **Lock** check boxes.

12.2.10 User Environment

This tab displays an area in the lower left corner of the Job Schedule Editor that includes the **Pass User Environment** check box. Select this check box if you want to pass the assigned user's environment variables to the job when it runs. When environment variables are recorded on the user account, selecting the **Pass User Environment** check box makes those environment variables available to the job and joblet.

A user's environment is recorded under the `user.env` fact on his or her account. This fact can be set when a user logs in to the Orchestration Server and is persisted until changed. A user's environment variables can be uploaded with the zos command line tool at login time in one of two variations:

- `zos login --user=`*`foo`*` --env`

  This command uploads the entire environment to the Job Scheduler. The upload can also be seen on the User object in the Orchestration Console.

- `zos login --user=`*`foo`*` --env=PATH`

  When the user logs in, he or she can specify one or more environment variables to use at login. The example above would result in just the PATH environment variable being uploaded.

  Multiple environment variables can be specified by delimiting with a comma, as in the following example:

  ```
  --env=PATH,LD_PATH,ID
  ```

NOTE: The user's environment variables can also be passed to the server when the user implements the zos command line tool when running a job (as opposed to logging in). The command passes the environment variable only for that particular job run.

`zos run `*`jobname`*` --env=`*`environment_variable`*

12.2.11 Constraints

This tab displays a constraint editor that you can use to create additional constraints for the job being scheduled. Typically, additional resource constraints (such as "start") are useful to delay the start of a job when it is triggered.

Any XML constraints listed in this tab needs to have a top-level XML `<constraints>` element, as in the following example:

```
<constraints>
  <constraint type="resource">
    <contains fact="resource.groups" value="myResourceGroup" />
  </constraint>
</constraints>
```

12.3 Understanding Cron Syntax in the Job Scheduler

The cron triggers you can configure in the Orchestration Server Job Scheduler use a Quartz crontrigger class for deciding when to invoke job execution. This is based on the standard Quartz format that you can find further described on the KickJava (http://kickjava.com/src/org/quartz/CronTrigger.java.htm) website.

This section includes the following information:

- Section 12.3.1, “Format,” on page 157
- Section 12.3.2, “Special Characters,” on page 157
- Section 12.3.3, “Examples of Cron Syntax,” on page 159
- Section 12.3.4, “Cron Scheduling Precautions,” on page 160

12.3.1 Format

A cron expression is a string comprised of 6 or 7 fields separated by white space. Fields can contain any of the allowed values, along with various combinations of the allowed special characters for that field. The fields are explained in the following table:

Table 12-2 *Fields in a Cron Expression*

Field Name	Mandatory?	Allowed Values	Allowed special Characters
Seconds	Yes	0-59	, - * /
Minutes	Yes	0-59	, - * /
Hours	Yes	0-23	, - * /
Day of the Month	Yes	1-31	, - * ? / L W
Month	Yes	1-12 or JAN-DEC	, - * /
Day of the Week	Yes	1-7 OR SUN-SAT	, - * ? / L #
Year	No	EMPTY, 1970-2099	, - * /

Cron expressions can be as simple as this:

```
* * * * ? *
```

Or cron expressions can be more complex, like this:

```
0 0/5 14,18,3-39,52 ? JAN,MAR,SEP MON-FRI 2002-2011
```

12.3.2 Special Characters

Cron syntax incorporates logical operators, which are special characters that perform operations on the values provided in the cron fields.

Table 12-3 *Special Characters in the Orchestration Server Cron Syntax*

Operator	Purpose	Example
asterisk (*)	Specifies all possible values for a field	An asterisk in the hour time field is equivalent to “every hour.”
question mark (?)	A question mark (?) is allowed in the day-of-month and day-of-week fields. It is used to specify “no specific value,” which is useful when you need to specify something in one of these two fields, but not in the other.	If you want a trigger to fire on a particular day of the month (for example, the 10th), but you don't care what day of the week that is, enter 10 in the day-of-month field, and ? in the day-of-week field.
dash (-)	Specifies a range of values	2-5, which is equivalent to 2,3,4,5
comma (,)	Specifies a list of values	1,3,4,7,8
slash (/)	Used to skip a given number of values	*/3 in the hour time field is equivalent to 0,3,6,9,12,15,18,21. The asterisk (*) specifies “every hour,” but the /3 means only the first, fourth, seventh. You can use a number in front of the slash to set the initial value. For example, 2/3 means 2,5,8,11, and so on.
L (“last”)	The L character is allowed for the day-of-month and day-of-week fields. Specifies either the last day of the month, or the last *xxx* day of the month.	The value L in the day-of-month field means “the last day of the month,” which is day 31 for January, or day 28 for February in non-leap years. If you use L in the day-of-week field by itself, it simply means 7 or SAT. But if you use it in the day-of-week field after another value, it means “the last *xxx* day of the month.” For example, 6L means “the last Friday of the month.” **TIP:** When you use the L option, be careful not to specify lists or ranges of values. Doing so causes confusing results.
W (“weekday”)	The W character is allowed for the day-of-month field. Specifies the weekday (Monday-Friday) nearest the given day.	If you specify 15W as the value for the day-of-month field, the meaning is “the nearest weekday to the 15th of the month.” So if the 15th is a Saturday, the trigger fires on Friday the 14th. If the 15th is a Sunday, the trigger fires on Monday the 16th. If the 15th is a Tuesday, it fires on Tuesday the 15th. However, if you specify 1W as the value for day-of-month, and the 1st is a Saturday, the trigger fires on Monday the 3rd, because it does not “jump” over the boundary of a month’s days. The W character can only be specified when the day-of-month is a single day, not a range or list of days. **TIP:** You can combine the L and W characters for the day-of-month expression to yield LW, which translates to “last weekday of the month.”

Operator	Purpose	Example
pound sign (#)	The pound sign (#) character is allowed for the day-of-week field. This character is used to specify “the nth” *xxx* day of the month.	The value of 6#3 in the day-of-week field means the third Friday of the month (day 6 = Friday and #3 = the 3rd one in the month). **Other Examples:** 2#1 specifies the first Monday of the month and 4#5 specifies the fifth Wednesday of the month. However, if you specify #5 and there are fewer than 5 of the given day-of-week in the month, no firing occurs that month.

NOTE: The legal characters and the names of months and days of the week are not case sensitive. MON is the same as mon.

You can specify days in two fields: month day and weekday. If both are specified in an entry, they are cumulative, meaning that both of the entries are executed.

12.3.3 Examples of Cron Syntax

The following table shows examples of full cron expressions and their respective meanings.

Table 12-4 *Results of Altered Cron Syntax on Execution Times*

Cron Expression Example	Description
0 0 12 * * ?	Fire at 12:00 p.m. (noon) every day
0 15 10 ? * *	Fire at 10:15 a.m. every day
0 15 10 * * ?	Fire at 10:15 a.m. every day
0 15 10 * * ? *	Fire at 10:15 a.m. every day
0 15 10 * * ? 2012	Fire at 10:15 a.m. every day during the year 2012
0 * 14 * * ?	Fire every minute starting at 2:00 p.m. and ending at 2:59.p.m., every day
0 0/5 14 * * ?	Fire every five minutes starting at 2:00 p.m. and ending at 2:55 p.m., every day
0 0/5 14,18 * * ?	Fire every five minutes starting at 2:00 p.m. and ending at 2:55 p.m., **and** fire every five minutes starting at 6:00 p.m. and ending at 6:55 p.m., every day
0 0-5 14 * * ?	Fire every minute starting at 2:00 p.m. and ending at 2:05.p.m., every day
0 10,44 14 ? 3 WED	Fire at 2:10 p.m. and at 2:44 p.m. every Wednesday in the month of March
0 15 10 ? * MON-FRI	Fire at 10:15 a.m. every Monday, Tuesday, Wednesday, Thursday and Friday
0 15 10 15 * ?	Fire at 10:15 a.m. on the 15th day of every month
0 15 10 15 * ?	Fire at 10:15 a.m. on the last day of every month
0 15 10 ? * 6L	Fire at 10:15 a.m. on the last Friday of every month

Cron Expression Example	Description
`0 15 10 ? * 6L 2011-2014`	Fire at 10:15 a.m. on every last Friday of every month during the years 2011, 20012, 2014, and 2014
`0 15 10 ? * 6#3`	Fire at 10:15 a.m. on the third Friday of every month
`0 0 12 1/5 * ?`	Fire at 12:00 p.m. (noon) every five days every month, starting on the first day of the month
`0 11 11 11 11 ?`	Fire every November 11th at 11:11 a.m.

12.3.4 Cron Scheduling Precautions

You should remember the following items when you use cron scheduling:

- Always check the effect of adding the ? and * characters in the day-of-week and day-of-month fields to ensure that the expected behavior fires correctly.
- Support for specifying both a day-of-week and a day-of-month value is not complete. You must currently use the ? character in one of these fields.
- Be careful when setting fire times to occur between 12:00 a.m. and 1:00 a.m. Changing to or from daylight saving time can cause a skip or a repeat in the schedule firing, depending on whether the clock moves backward or forward.

13 The Policy Debugger

The Policy Debugger is a tabbed page available in several views of the Cloud Manager Orchestration Console. This tool helps you to determine the reasons for the current state of a job. The following figure shows the **Policy Debugger** tab opened in the Jobs view of the Orchestration Console.

Figure 13-1 *The Jobs View of the Orchestration Console with the Policy Debugger Page Open*

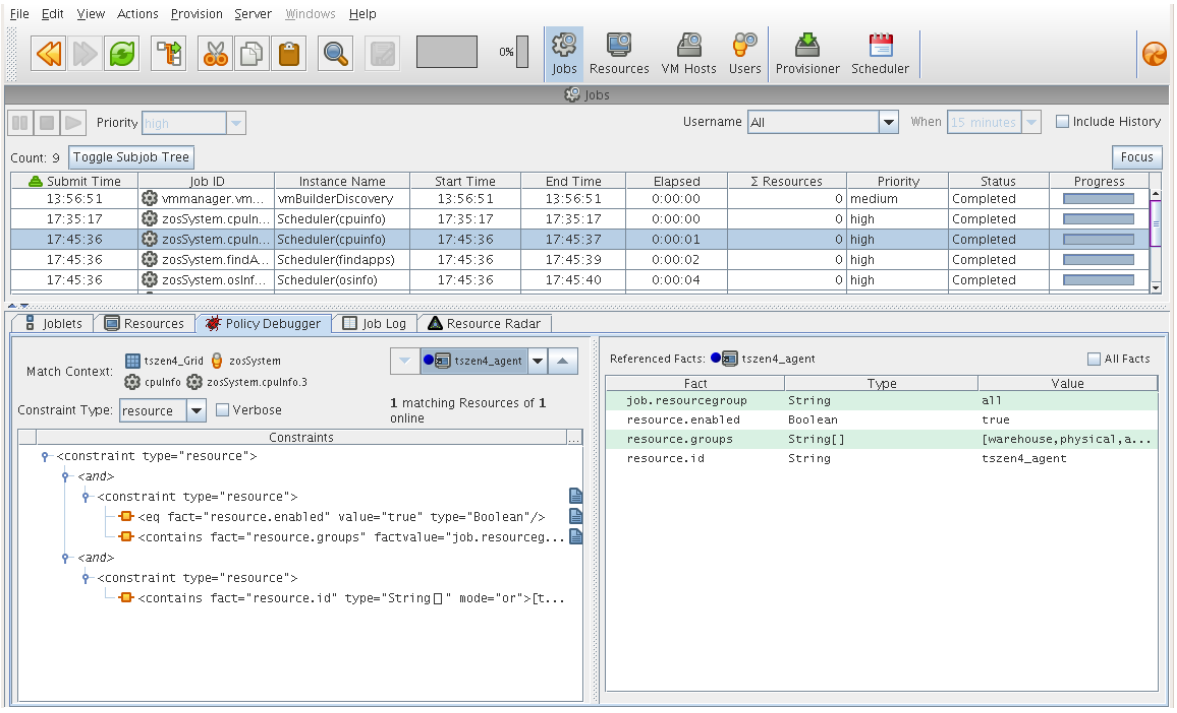

The **Policy Debugger** tab is also available in the VM Hosts view and in the Provisioner view of the Orchestration Console.

- Section 13.1, “Constraints Table View,” on page 161
- Section 13.2, “Facts Table View,” on page 166
- Section 13.3, “Policy Debugger Use Cases,” on page 167

13.1 Constraints Table View

The left side of the Policy Debugger page is referred to as the Constraints Table view.

Figure 13-2 *The Constraints Table View*



The appearance of this view can change, depending on the constraint type you select in the drop-down menu. For a description of these types, see Section 13.1.2, “Constraint Type List,” on page 163.

The Constraints Table View is composed of several parts:

- Section 13.1.1, “Match Context Area,” on page 162
- Section 13.1.2, “Constraint Type List,” on page 163
- Section 13.1.3, “Verbose Check Box,” on page 163
- Section 13.1.4, “Capable Resources Summary,” on page 164
- Section 13.1.5, “Constraints Column of the Constraints Table View,” on page 164
- Section 13.1.6, “Policy Column of the Constraints Table,” on page 165

13.1.1 Match Context Area

The policy debugger provides the general identification of a job instance in the **Match Context** area of the Constraints Table view. The Match Context defines everything associated with running a job on a particular resource because it references facts, which are also referenced in policies. The policies define how, when, and where the job runs.

Figure 13-3 *The Match Context Area of the Constraints Table View in the Policy Debugger*



That identifying facts in the Match Context include:

Matrix: The icon and a text string identifies the machine that matches the grid name given to the Orchestration Server where this job is running.

User: The icon and a text string identifies the User object that matches the user who is running the job.

Job: The icon and a text string identifies the deployed job that matches the one that is running on the grid.

Job Instance: The icon and a fully qualified text string identifies the specific job instance that matches the deployed job running on the grid.

Resource: The Resource drop-down list shows all resources. The list appears in the Match Context if the **resource** constraint type is selected. The resources in the list that are currently offline display with a dimmed icon. If available, a listed resource has a colored dot by its side. The color of the dot (blue or gray) and the resource type it accompanies has significance:

- A blue dot with the **All Resources** label indicates that at least one resource matches the constraints and is capable of servicing the job.
- A gray dot with the **All Resources** label indicates that no resources match the constraints.
- A blue dot with a named, selected resource indicates that its constraints match and it is capable of servicing the job.
- A gray dot by a named, selected resource indicates that its constraints do not match and that it is not capable of servicing the job.

The following figure shows a list of eight resources. Four of those resources (**lab.a**, **lab.b**, **lab.c** and **lab.d**) are online. Their constraints match, so they are capable of servicing the job.

Figure 13-4 *Resource Drop Down List Showing Online and Offline Resources*

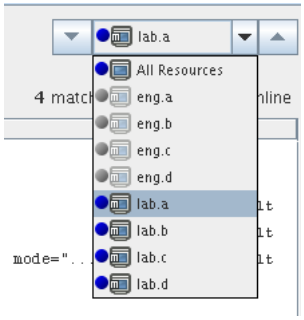

13.1.2 Constraint Type List

Select one of the constraint types in the drop down list to specify a policy context. Constraint types in the list are disabled (dimmed) if they do not apply to the job that you are debugging.

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

13.1.3 Verbose Check Box

When you select the Verbose check box, the `reason` string specified in the policy is displayed in the **Constraints** tree. For more information, see Section 13.1.5, “Constraints Column of the Constraints Table View,” on page 164.

13.1.4 Capable Resources Summary

Directly under the Resource List in the constraint view, a populated string summarizes the resources that are capable of servicing the job. For example, `4 matching Resource of 10 online` indicates that four of the ten total online resources are capable of servicing the job.

13.1.5 Constraints Column of the Constraints Table View

The Constraints column of the constraints table view shows the logical hierarchy (that is, a “tree” format) of the constraints that are defined by the policies associated with the job. You can identify the status of the listed constraints by the icons that might be displayed in the far left column of the table:

No icon: The constraint passes the match. It is a “true” match. The figure below shows that the resource `lab.a` is available to run the job because all of its constraints match. No red icons are displayed next to any listed constraint.

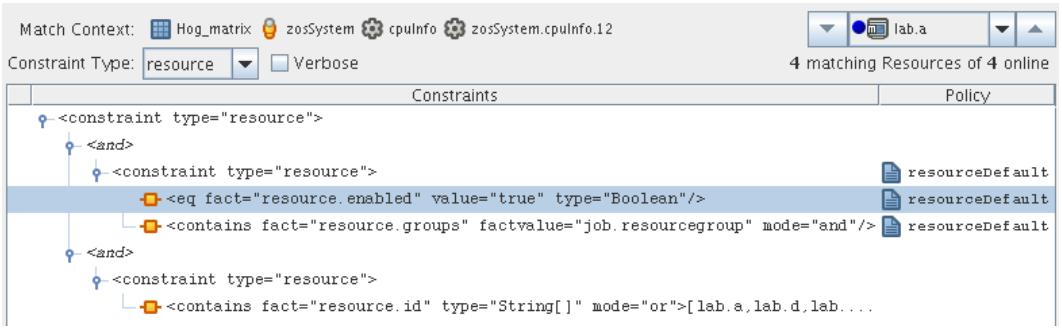

Red dot icon: The constraint does not pass the match. The figure below shows that the resource `eng.a` cannot run the job because its constraints do not match.

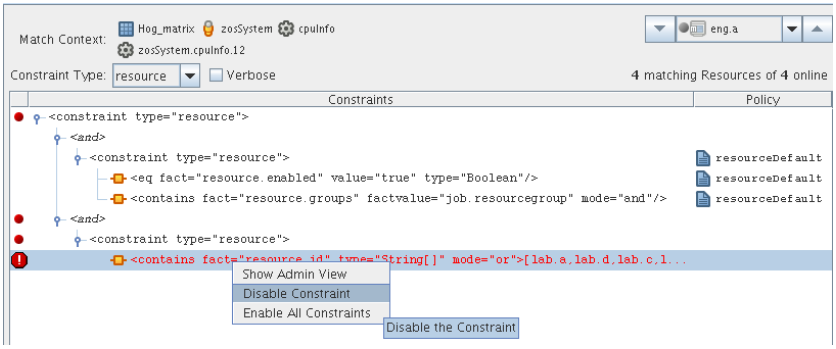

Red octagonal icon: The constraint does not pass the match and is blocking the job. The figure above also shows the blocking constraint (red octagon).

Green dot icon: A blocking constraint has been disabled so that it behaves like a match. The figure below shows the green dot icon next to that the constraint that was formerly blocked and can now behave as a match.

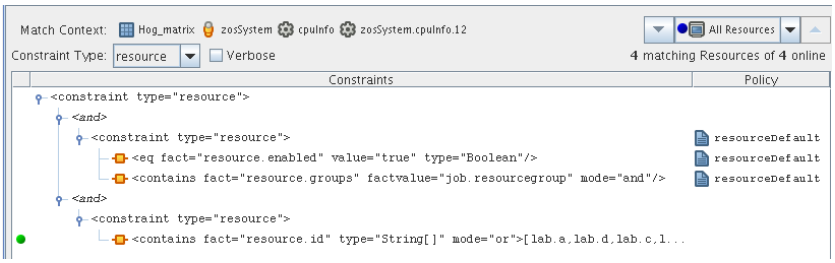

If you right-click a constraint in the table, a popup menu with three options is displayed. This menu lets you change the status of the constraint.

- **Show Admin View:** Select this option to open the Admin View for the specific resource selected.
- **Disable Constraint:** Select this option to disable (attach a green dot icon to) a constraint. Disabling a constraint with this function effectively makes it match, which is a condition that can be useful if you want to perform a “what if” test without actually changing a policy.
- **Enable All Constraints:** Select this option if you have disabled one or more constraints during testing and you want to restore them to the enabled state.

Cached Constraints in the Constraints Column

When you change the constraint type in the Constraints Type List, the background of the table changes to green for some types. These are “cached” constraints that are saved with the job after it is complete. Their purpose is to help you debug the policy.

Figure 13-5 *Cached Constraints Displayed in the Constraints Table View*

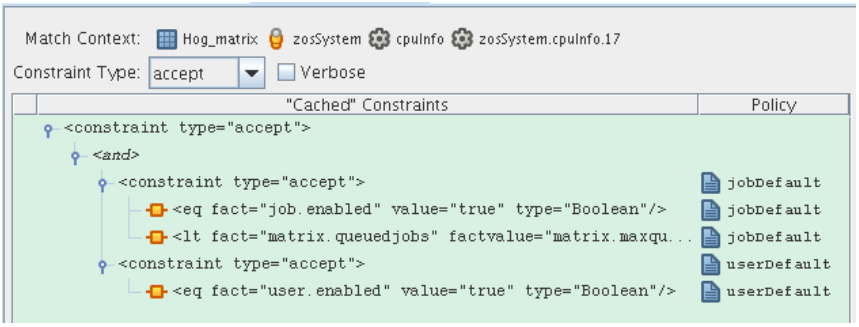

13.1.6 Policy Column of the Constraints Table

The Policy column of the Constraints Table displays the policy name that contributed the constraint. Right-click a policy name to open a popup menu offering the option to open the policy editor for the specified policy. The menu also includes constraint enabling or disabling options, just as the pop-up menu for constraint column does.

Figure 13-6 *The Pop-up Menu Launched from the Policy Column*

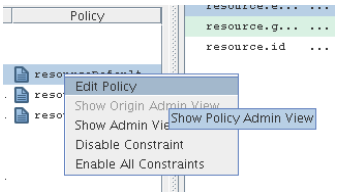

13.2 Facts Table View

The Facts Table view displays the facts referenced in the Constraint Tree view for a specified Resource. Selecting a fact in the Constraint Tree automatically selects that fact in the table.

Figure 13-7 *The Constraints Table View and the Accompanying Facts Table View*

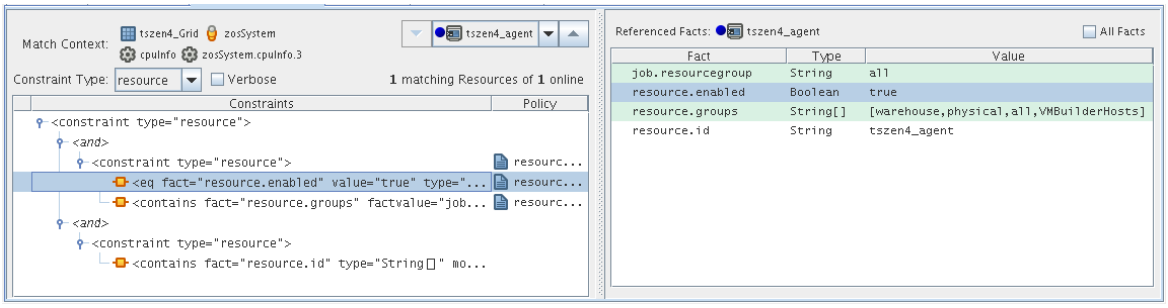

If you right-click a column head in this table, a menu is launched where you can select the columns that you want to display.

Figure 13-8 *Menu Used to Select the Columns Displayed in the Facts Table View of the Policy Debugger*

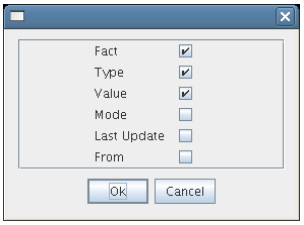

13.2.1 All Facts Check Box

If you select the **All Facts** check box at the top of the Facts Table view, all of the facts (including matrix, user, job, jobargs, jobinstance, and resource facts) associated with the Match Context are listed.

If you select **All Resources** in the Match Context (see Section 13.1.1, “Match Context Area,” on page 162) and you also select the **All Facts** check box, the Facts Table view displays all the facts for all resources for the specified Match Context.

Figure 13-9 *All Facts Check Box Selected with All Resources in Match Context*

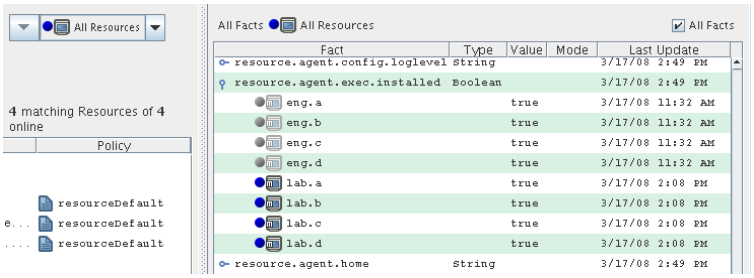

13.3 Policy Debugger Use Cases

The objective of this use case is to run a job that remmians in the waiting state and to use the policy debugger to identify why it is in the state and to make the necessary changes to get the job to run. The `quickie.job` is used along with a simple policy that specifies that the resources that are to be used must be in a Resource group called `debugger`.

Use the following steps to re-create the use case.

1. In the Orchestration Console, create a user named `debugger`.
2. Deploy `quickie.job` from the `/examples` directory.
3. In the Orchestration Console, create a schedule named `quickie`, specifying the `quickie` job and the `debugger` user.
4. In the Orchestration Console, create a policy and name it `debuggerExample`. The policy needs to specify that the resource used belongs to the group called `debugger`.

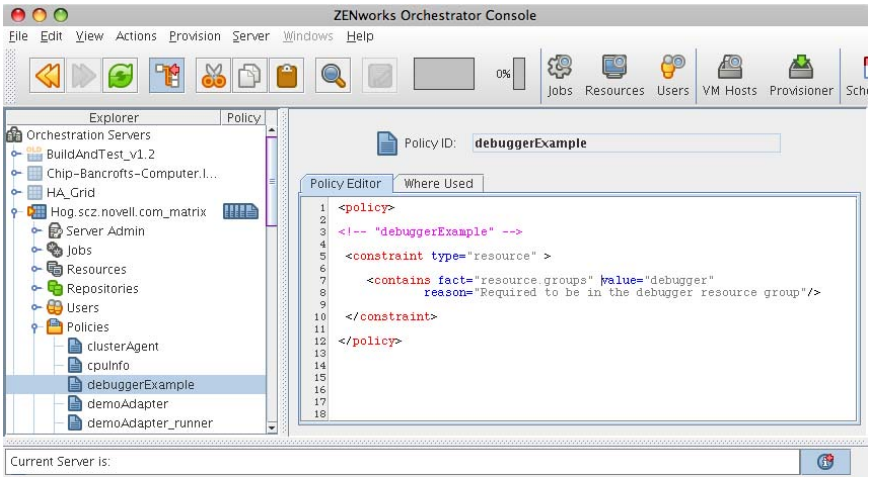

5. In the Orchestration Console, associate the `debuggerExample` policy to the `quickie` job.
6. In the Job Scheduler view of the Orchestration Console, select the `quickie` schedule, then click **Run Now** to run the `quickie` schedule.
7. In the Job Monitor view of the Orchestration Console, select the **Policy Debugger** tab and verify that the job is in the waiting state.
8. In the Constraints Table view, open the **Constraint Type** drop-down list, then select **Allocation**.
9. In the Match Context area of the Constraints Table view, open the **Resource** drop-down list, then select any resource to refresh the Constraints Table and Facts Table views.

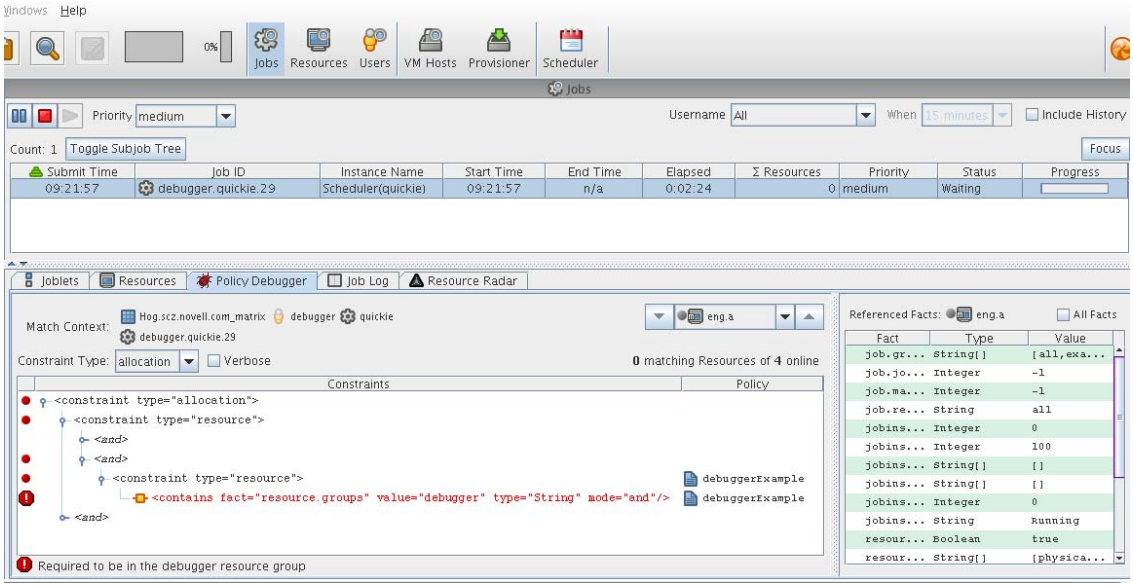

The Policy Debugger displays a red icon near the constraints that fail to match. The larger, red octagonal icon shows the particular constraint that is “blocking” and preventing the job from running on the resource. This is the constraint that is causing the job to be in a waiting state. The Constraints Table also displays the policy name (`debuggerExample`) that is contributing the constraint that is causing problems.

There are several ways to get the job to run:

- Create a Resource group called `debugger`, then place a resource in that group to satisfy the constraint specified in the policy.
- Disassociate the policy (`debuggerExample`) from the job (`quickie`).
- In the Constraints Table, right-click the blocking constraint and select **Disable Constraint**.

14 Grid Object Health Monitoring

This section includes the following information:

- Section 14.1, “Health Facts,” on page 169
- Section 14.2, “Health Events,” on page 171
- Section 14.3, “Health Debugger,” on page 171

14.1 Health Facts

The Resource grid object, the VM host grid object, the User grid object, and the Repository grid object each have an attribute or Fact that denotes the health of the object.

- `resource.health`
- `vmhost.health`
- `user.health`
- `repository.health`

Empirically, object health is a simple Boolean value, with `True` indicating that the object is healthy. This value can be controlled in a number of ways. An unhealthy object is displayed in the Cloud Manager Orchestration Console with a red “plus” symbol to signal the object’s condition.

Figure 14-1 *Tree View of Repository Grid Objects in the “all” Group, Some Objects Unhealthy*

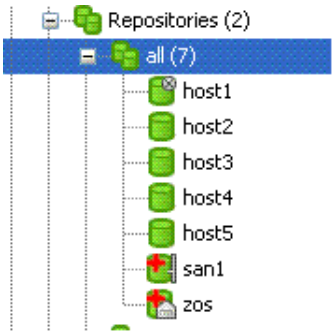

You can define what constitutes the health or non-health of the grid object by setting this health fact. The health fact can be set or cleared in several ways:

14.1.1 Explicitly Set or Cleared by the Administrator

The administrator can use tools in the Orchestration Console to explicitly set or clear the fact.

1 Select any grid object in the Orchestration Console, then click the **Info/Groups** tab in the Workspace view. This is the “info” attribute editor. The attributes on this page let you edit facts.

The object information panel of the page has a **Healthy** check box that you can select or deselect to set the health of the object.



2 On the Constraint/Fact fact page of a grid object, right click the *xxx*.health fact name, then click **Edit/View Fact** to open the Edit Fact dialog.

You can set the health of the object by selecting or deselecting the health check box. Changing the value in the Orchestration Console in this way has an immediate effect unless the value is overridden by an attached policy (this follows the normal rules of policy inheritance).

14.1.2 Set by Using a Discovery Job

A discovery job is a job that is periodically scheduled to run on resources and to explicitly set the health fact, much like it sets other discovered facts). In this case, the discovery job performs a setFact (*xxx*.health) from JDL code.

14.1.3 Set by Using a Policy

This method has little practical use except for locking the value immediately to override the setting (that is, the typical policy behavior) on the grid object:

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

14.1.4 Set by Using a Computed Fact

Set by using a computed fact. This method can be used to monitor the health according to a computed value. One applied scenario for this method might be a computed fact that performs a statistical analysis of historical load data, perhaps provided by the Metrics facility.

14.1.5 Set Automatically by Using a Health Constraint

Set automatically by using a health constraint. This is the most practical use and is best illustrated with examples.

Example 1: Define resources as “unhealthy” if their 10-minute load average is greater than 5.

```
<policy>
   <constraint type="health">
       <lt fact="resource.metrics.loadaverage.history.10_min" value="5.0" />
   </constraint>
</policy>
```

You could attach this policy directly to the Resource grid object or to a Resource group. Attaching it to a Resource group is more practical.

Example 2: Define a user as unhealthy if he or she has no money in an account.

```
<policy>
   <constraint type="health">
       <ge fact="user.account.balance" value="0" />
   </constraint>
</policy>
```

You could attach this policy directly to the User grid object, or to a User group. Attaching it to a User group is more practical.

You can aggregate (that is, group together with “and” or “or”) health constraints by using normal rules of policy aggregation.

By default, the Orchestration Server runs health constraints every 30 seconds. To alter this interval, you must contact NetIQ Support.

14.2 Health Events

Each time the value of a health fact changes, an event is generated. This event can be subscribed to by long-running Jobs or the event can be used to trigger Jobs in the Job Scheduler (see Chapter 12, “The Orchestration Server Job Scheduler,” on page 145). The event names are different for each object type. They are listed in the following table.

Table 14-1 *Event Names for Grid Objects*

Object	Event Name
User	`USER_HEALTH`
Resource	`RESOURCE_HEALTH`
Repository	`REPOSITORY_HEALTH`
VM host	`VMHOST_HEALTH`

14.3 Health Debugger

Several objects modeled by the Orchestration Console have a health fact that can be used to visually show the health of the object or to trigger a job (see Chapter 12, “The Orchestration Server Job Scheduler,” on page 145) upon a state change. This fact can optionally be manually set, or more usually automatically set through the periodic execution of the health constraint placed on that object. When such a health constraint is active, you might need to debug to discover the reasons for the changed state. The Health Debugger is a useful debugging tool.

The Health Debugger is a tabbed page (sometimes called an “admin view”) available in several views of the Orchestration Console and functions much like the more generic “Policy Debugger.” This tool helps you to determine the reasons for the current Health of a Grid object. The following figure shows the Health Debugger tab opened in the Resource object view of the Orchestration Console.

Figure 14-2 *Orchestration Console Jobs View with the Health Debugger Page Open*

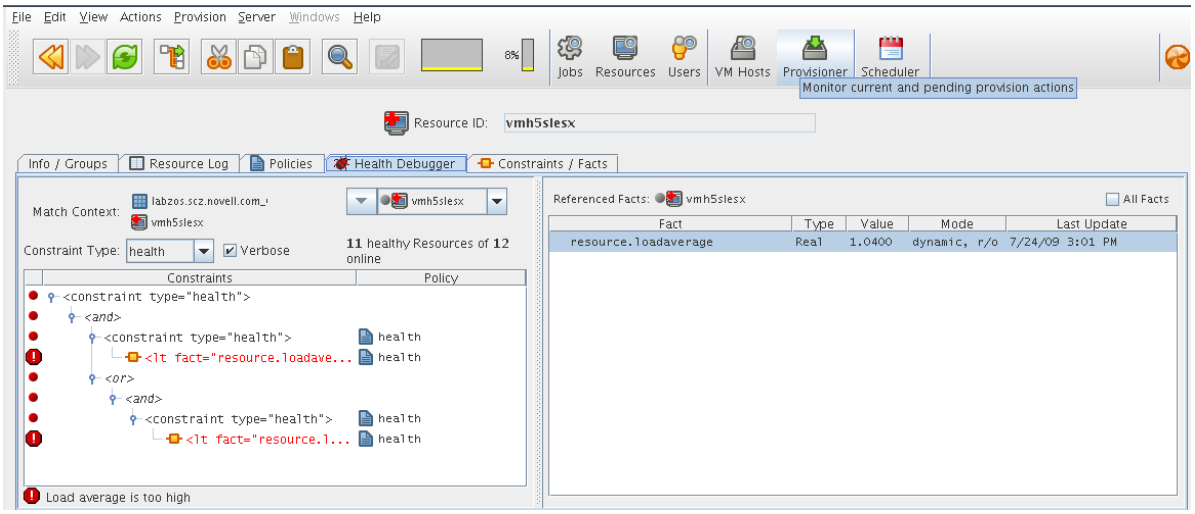

The Health Debugger tab is also available in the VM Host object view, the Repository object view and in the User object view of the Orchestration Console.

- Section 14.3.1, “Constraints Table Panel,” on page 172
- Section 14.3.2, “Facts Table View,” on page 175

14.3.1 Constraints Table Panel

The left side of the Health Debugger page is referred to as the Constraints Table view.

Figure 14-3 *The Constraints Table View*



The appearance of this view can change, depending on the Constraint Type you select in the Constraint drop down list. In effect, if you change from the health selection, you will be debugging other constraints. For a description of these constraint types, see Section 13.1.2, “Constraint Type List,” on page 163. Different objects selected for the view change the **Match Context** in which the constraint is executed, which is useful for comparing how the constraint evaluates other objects.

The Constraints Table view is composed of several parts:

- “Match Context Area” on page 173
- “Verbose Check Box” on page 173
- “Capable Objects Summary” on page 173
- “Constraints Column of the Constraints Table View” on page 174

Match Context Area

The Health Debugger provides the general identification of a Grid object in the **Match Context** area of the Constraints Table view. The Match Context defines every object that is available to the constraint you are viewing.

Figure 14-4 *The Match Context Area of the Constraints Table View in the Health Debugger*

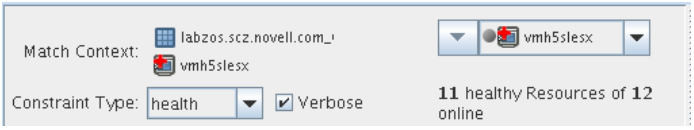

In this example, the identifying Facts in the Match Context include:

Matrix: The icon and a text string identifies the machine that matches the grid name given to the Orchestration Server where this object exists.

Object icon: The object icon and a text string identifies the object (VM host, Repository, or that matches the user who is running the job. If the object icon has a red cross overlaid, it is unhealthy.

Object list: The object drop down list shows all named objects of the type selected in the Explorer Tree. The objects in the list that are currently offline display with a dimmed icon. If available, a listed object has a colored dot by its side. The color of the dot (blue or gray) and the object type it accompanies has significance:

- A blue dot with the **All <Object Type>** label indicates that at least one object in the list matches the constraints and is healthy.
- A gray dot with the **All <Object Type>** label indicates that no objects of this type match the constraints.
- A blue dot with a named, selected object indicates that its constraints match and it is healthy.
- A gray dot by a named, selected object indicates that its constraints do not match and that it is not healthy.

Verbose Check Box

When you select the Verbose check box, the `reason` string specified in the policy is displayed in the **Constraints** tree. For more information, see Section 13.1.5, “Constraints Column of the Constraints Table View,” on page 164.

Capable Objects Summary

Directly under the Object list in the constraint view, a populated string summarizes the resources that are capable of servicing the job. For example, `11 healthy Resources of 12 online` indicates that 11 of the 12 total online Resources are healthy.

Constraints Column of the Constraints Table View

The Constraints column of the constraints table view shows the logical hierarchy (that is, a “tree” format) of the constraints that are defined by the Policies associated with the Job. You can identify the status of the listed constraints by the icons that might be displayed in the far left column of the table:

No icon: no icon: The constraint passes the match. It is a “true” match. No red icons are displayed next to any listed constraint.

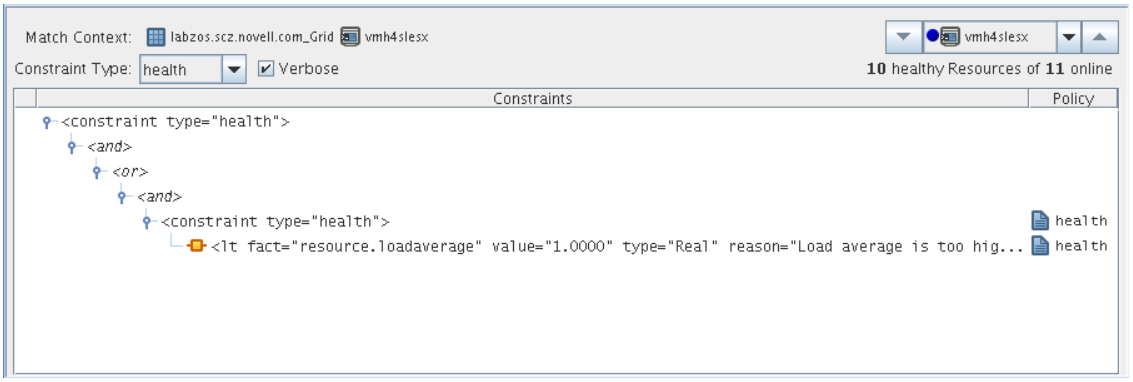

Red dot icon: The constraint does not pass the match. The figure below shows that the resource `vmh5slesx` cannot run the job because its health constraints do not match.

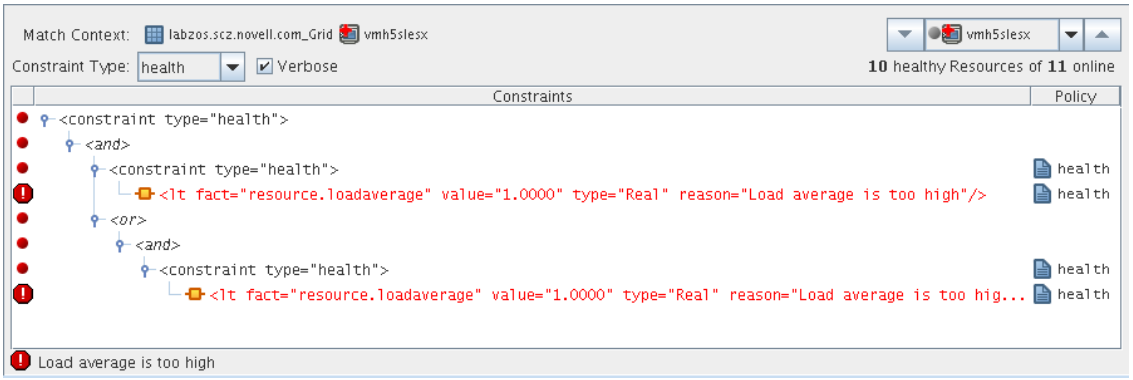

Red octagonal icon: For convenience, only one of the constraints is identified as the blocking constraint. This is the constraint that the system has determined is responsible for the constraint as a whole to fail (note that individual constraint lines can fail without causing the entire constraint to fail).

Green dot icon: green dot icon: A failing constraint has been disabled so that it behaves like a match (pass). The figure below shows the green dot icon next to that the constraint that was formerly failing and can now be forced to behave as a match.

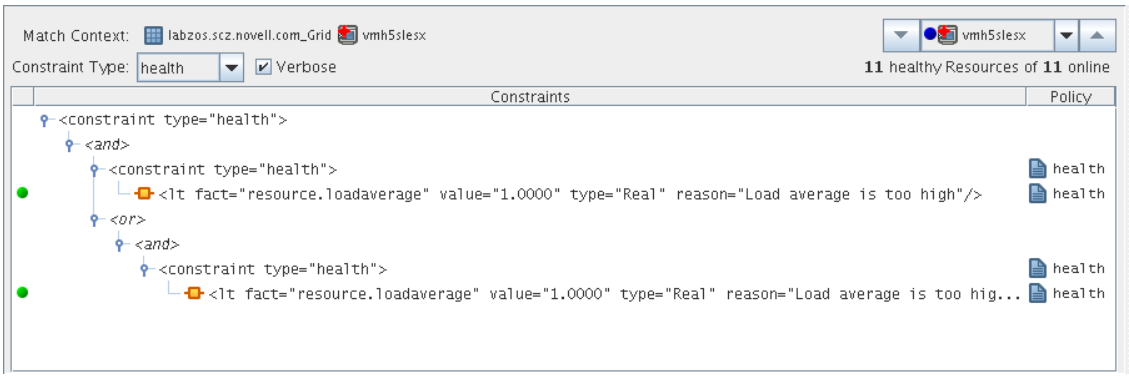

If you right-click a constraint in the table, a popup menu with three options is displayed. This menu lets you change the status of the constraint. Disabling a constraint is useful if you want to temporarily relax a condition without editing or redeploying the entire policy and potentially affecting other objects that share the policy. A disabled constraint can be re-enabled later.

The constraint status change options include the following:

- **Show Admin View:** Select this option to open the Admin view for the specific object selected.
- **Disable Constraint:** Select this option to disable (attach a green dot icon to) a constraint. Disabling a constraint with this function effectively makes it match, a condition that can prove useful if you want to perform a “what if” test without actually changing a policy. Green
- **Enable All Constraints:** Select this option if you have disabled one or more constraints during testing and you want to restore them to the enabled state.

NOTE: Health constraints are always re-evaluated in the debugger. The last system execution (cached constraint) is not available for health constraints.

The Policy column of the constraints table displays the policy name that contributed the constraint. Right-click a policy name to open a popup menu offering the option to open the policy editor for the specified policy. The menu also includes constraint enabling or disabling options, just as the pop-up menu for constraint column does.

Figure 14-5 *The Pop-up Menu Launched from the Policy Column*

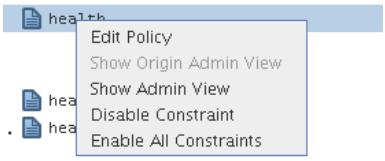

14.3.2 Facts Table View

The Facts Table view displays the facts referenced in the Constraint Tree view for a specified object. Selecting a fact in the Constraint tree automatically selects that fact in the table.

Figure 14-6 *The Constraints Table View and the Accompanying Facts Table View*

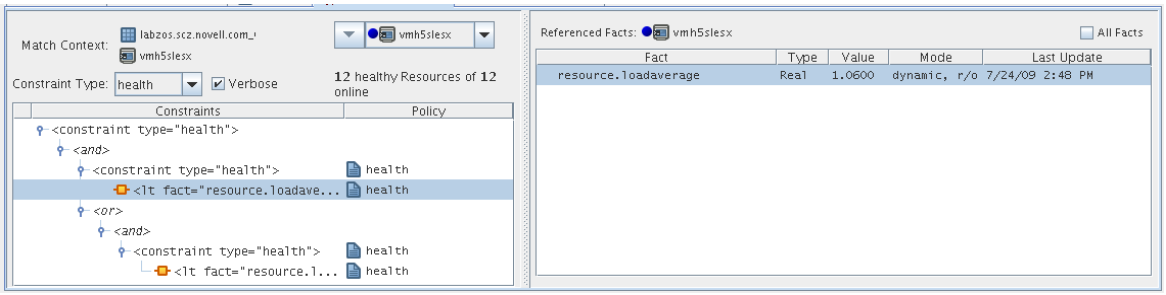

If you right-click a column head in this table, a menu is launched where you can select the columns that you want to display.

Figure 14-7 *Menu Used to Select the Columns Displayed in the Facts Table View of the Policy Debugger*

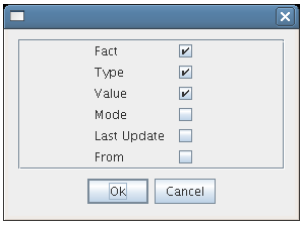

All Facts Check Box

If you select the **All Facts** check box at the top of the Facts Table view, all of the facts (including matrix, and <`Object type`> facts) associated with the Match Context are listed.

If you select **All <Object Type>** in the Match Context (see Section 13.1.1, “Match Context Area,” on page 162) and you also select the **All Facts** check box, the Facts Table view displays all the <object type> facts for the specified Match Context of the selected object.

Figure 14-8 *All Facts Check Box Selected with All VM Hosts in Match Context*

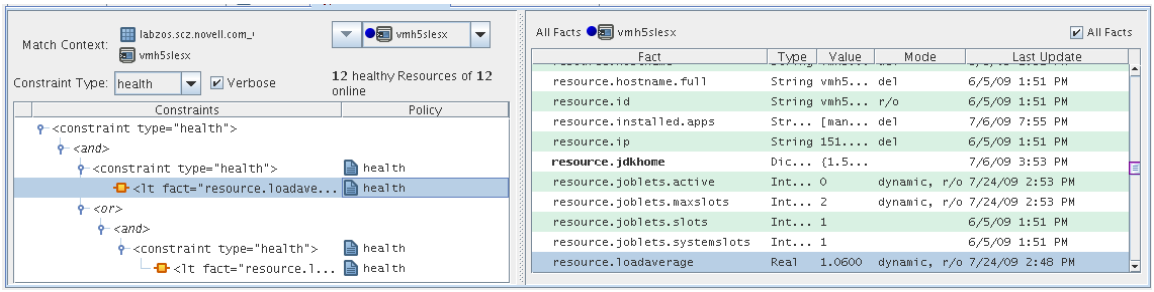

15 Events

This section contains the following information:

- Section 15.1, “Understanding the Orchestration Server Events System,” on page 177
- Section 15.2, “Event Object Visualization and Management in the Orchestration Console,” on page 180
- Section 15.3, “Event Debugger,” on page 182

15.1 Understanding the Orchestration Server Events System

The the Orchestration Server Event System integrates with the Job Scheduler. Event notifications can start jobs and can also invoke Event handler methods in long-running jobs. In turn, a job can react to the Event by starting other server actions, by modifying object attributes, or by executing another external process.

For example, an Event notification can occur when a VM Host has exceeded its configured load limits. This Event can start a job that migrates VMs off of the loaded VM Host or VM Hosts.

The Orchestration Server supports two Event types:

- Built-in Events, such as change of status of the health of a resource's or a change in the online status of the resource.
- Rule-based Events that are triggered when the attributes of an object satisfy the rules (constraint conditions) defining the Event.

This section includes the following information:

- Section 15.1.1, “Event Notification,” on page 177
- Section 15.1.2, “Built-in Events,” on page 177
- Section 15.1.3, “Rule-Based Events,” on page 178

15.1.1 Event Notification

An Event notifies two other Cloud Manager Orchestration services, the Job Scheduler and the Job Broker. The Job Scheduler starts jobs that are awaiting an Event to trigger them. The Job Broker invokes a callback on any long-running job that has registered for notification of an Event.

See Chapter 12, “The Orchestration Server Job Scheduler,” on page 145 for more information about setting up a Job Schedule.

15.1.2 Built-in Events

Built-in events occur when a managed object comes online/offline or has a health status change.

The Orchestration Server uses the following built-in Events to keep managed objects synchronized.

Table 15-1 *The Orchestration Server Built-in Events*

Event Name	Description
AGENT_VERSION_MISMATCH	Resource Agent version mismatch (agent needs upgrade)
REPOSITORY_HEALTH	Repository health status has changed
RESOURCE_HEALTH	Resource health status has changed
RESOURCE OFFLINE	Resource Agent has logged out of the server
RESOURCE_ONLINE	Resource Agent has logged in to the server
SERVER_UP	Server has fully started
USER_HEALTH	User health status has changed
USER_ONLINE	User has logged in to the server
VMHOST_ADDED	VM Host has been added
VMHOST_HEALTH	VM Host health status has changed
VMHOST_NOT_AVAILABLE	No VM Host is available
VMHOST_REMOVED	VM Host has been removed

For example, when a resource comes online (that is, the agent connects to the server), the `RESOURCE_ONLINE` Event is fired and both scheduled jobs with a trigger for that Event and long-running jobs with Event handlers are notified.

The `RESOURCE_ONLINE` built-in Event is used by the embedded discovery jobs, such as for discovering operating system and CPU information (`osInfo` and `cpuInfo` jobs). Both `osInfo` and `cpuInfo` job archives (`.job`) include a schedule file (`.sched`) specifying a trigger (`.trig`) that allows these jobs be started when notification of the `RESOURCE_ONLINE` Event occurs.

15.1.3 Rule-Based Events

Rule-based Events are defined in an XML document. They are deployed to the Orchestration Server and managed through the Orchestration Console. Rules can be a simple object attribute (fact) equivalency check or they can use AND,OR, IF, ELSE logic, among other things, in an Event ruleset.

The rules follow the same syntax as the constraints that are defined in XML policy files for all Grid Objects, such as Jobs, VM Hosts, etc.

The Orchestration Server Event Service evaluates the rules; if the rules pass, an Event notification occurs.

The XML Schema document specification can be found in `<install dir>/doc/xsds/event_1_0_0.xsd`.

The Event XML specification is composed of three sections.

- `<context>`
- `<trigger>`
- `<reset>`

NOTE: Both the `<context>` and `<trigger>` sections are required.

<context> section

The <context> section defines the context in which the Event rules are evaluated. With Events, you specify what objects are in the Event rule context in this section. The available objects are Job, Jobinstance, Resource, Repository, User, and VMHost. From these objects, you can specify one object set to iterate over and optionally a single instance of the object.

<trigger> section

The <trigger> section defines the rules for when an Event notification occurs. The <trigger> format is the same syntax as <constraints> used in policies.

<reset> section

The optional <reset> section defines the rules for when an Event is reset. If the <reset> rule is not used, an Event is reset based on a timeout. The <reset> format is also the same syntax as in <constraints> used in policies.

The resetInterval attribute is set on the <event> XML element. If "resetInterval" and <reset> are not used, the default timeout for resetting is 10 minutes.

The following example, taken from the "vmhost.event" in <install dir>/examples/events), defines that a notification occurs when a VM host becomes overloaded.

```
1<event>
2
3   <context>
4      <vmhost />
5   </context>
6
7   <trigger>
8      <gt fact="vmhost.vm.count" value="0" />
9      <gt fact="vmhost.resource.loadaverage" value="2" />
10   </trigger>
11
12   <reset>
13      <lt fact="vmhost.resource.loadaverage" value=".5" />
14   </reset>
15
16</event>
```

Lines 3-5: This section defines the context for the Event's rule evaluation.

Line 4: The context specifies all VM host objects, so the Event Service iterates over all VM hosts. On each VM host, the <trigger> rule will be evaluated, so in this case, the Event context is composed of one or more VM hosts.

Lines 7-12: This section defines the Trigger rule to determine if this Event is to fire notifications or not. If the trigger rule does not pass, no Event notifications occur.

Line 8: Consider only VM hosts that have at least one VM instance running.

Line 9: Check the running average of the VM host's load average if it exceeds a threshold value. In this case, run the check if the average is greater than 2.

Lines 12-14: This section defines the Reset rule to determine if a previously triggered VM host can be reset and triggered again.

Line 13: Only reset if the running average of the VM host's load average drops below a threshold.

When a VM host passes the trigger rule, the VM host does not pass the trigger rule again until the reset rule (load average drops below threshold) passes.

See the `repository.event` example (`<install dir>/examples/events/repository.event`) for an Event with a rule that evaluates the `freespace` fact on all repository objects.

15.2 Event Object Visualization and Management in the Orchestration Console

The Events folder is displayed in the Explorer tree between the Computed Facts and Metrics folders.

Figure 15-1 *The Events Folder in the Orchestration Console Explorer View*

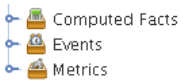

Although the Cloud Manager Orchestration Server includes several built-in Events (see Section 15.1.2, “Built-in Events,” on page 177), these Events are not displayed in the Explorer view. Only custom Events (defined in XML by the administrator and then deployed on the server) are displayed in the tree.

When an Event object is deployed, its icon is displayed in the tree in the Events folder.

Figure 15-2 *An Event Object in the Events Folder*

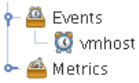

The icon in the Explorer tree might be overlaid with a write symbol  to indicate that its XML content has changed and needs to be saved. For more information about changing the XML content, see Section 15.2.4, “Event Editor,” on page 181.

This section includes the following information about how the Event object is managed in the Orchestration Console:

- Section 15.2.1, “Deploying a New Rule-Based Event,” on page 180
- Section 15.2.2, “Deploying a Pre-written Rule-Based Event,” on page 181
- Section 15.2.3, “Undeploying an Event,” on page 181
- Section 15.2.4, “Event Editor,” on page 181

15.2.1 Deploying a New Rule-Based Event

Use the following steps to create a new event.

1. In the Explorer view, right-click the Events folder, then select **New Event** to open the Create a New Event dialog.
2. Specify the name for the new Event, then click **OK** to create the new Event object.

   The Orchestration Server then deploys the new Event object on the server, where it can be managed. The Orchestration Console opens the Event Editor, where you can edit the XML definition of this Event. For more information, see Section 15.2.4, “Event Editor,” on page 181 and Section 15.1.3, “Rule-Based Events,” on page 178.

15.2.2 Deploying a Pre-written Rule-Based Event

Use the following steps to deploy a pre-written Event (an XML `.event` file).

1. Right-click the Events container, then select **Deploy Event** to open the Select the Component File to Deploy dialog.
2. In the dialog, navigate to the file system location of the Event file you previously created, or to an example `.event` file from `/opt/novell/zenworks/zos/server/examples/events`, then click **OK** to deploy the pre-written Event.

   When you deploy the rule-based Event, the Orchestration Console opens the Event Editor, where you can edit the XML definition of this Event. For more information, see Section 15.2.4, “Event Editor,” on page 181 and Section 15.1.3, “Rule-Based Events,” on page 178.

15.2.3 Undeploying an Event

When an Event has been deployed, it can be undeployed. Undeploying deletes the Event object within the server, but it does not delete the source `.event` file, which still exists and can be redeployed.

To undeploy an Event, right-click the Event object in the Explorer tree, then select **Undeploy**. You can also simply select the object and press Delete.

15.2.4 Event Editor

The Event Editor opens when you select a deployed Event in the Explorer tree of the Orchestration Console.

Figure 15-3 *The Event Editor*

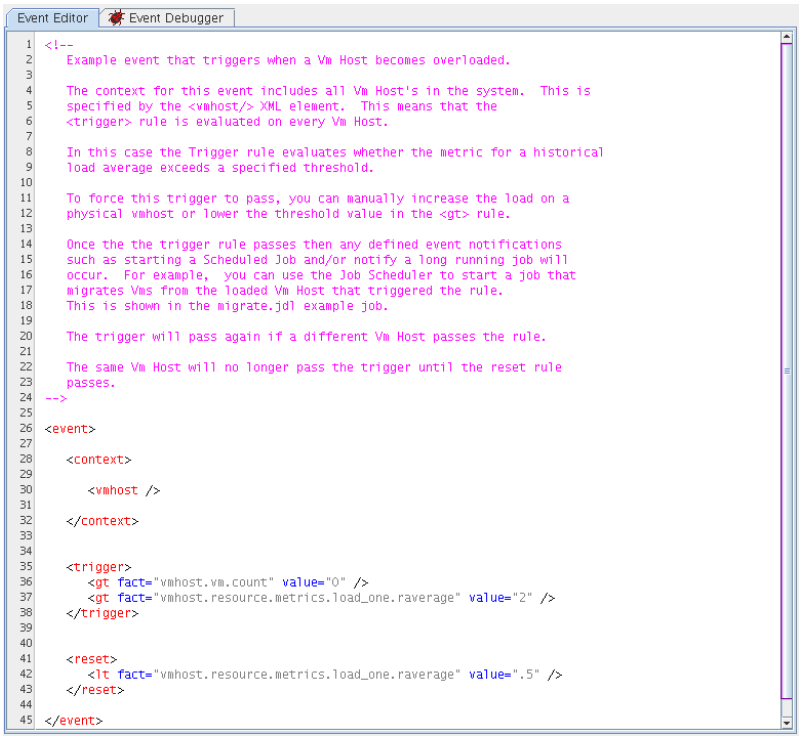

Inside this editor, you can make changes to the XML content of the Event. Example Events contain comments that explain how you can use them and the behavior you can expect to see as a result of deploying them. For the changes you make to be effective, you need to click the Save tool.

For more information about the allowed XML syntax within an Event, see Section 15.1.3, “Rule-Based Events,” on page 178.

15.3 Event Debugger

The Event Debugger is a tabbed page available from the Explorer view of an Event object when you select an event and then click the Event Debugger tab. This tool helps you to determine the reasons for the current state (`triggered` or `reset`) of an Event. The following figure shows the Event Debugger view:

Figure 15-4 *The Event Debugger*

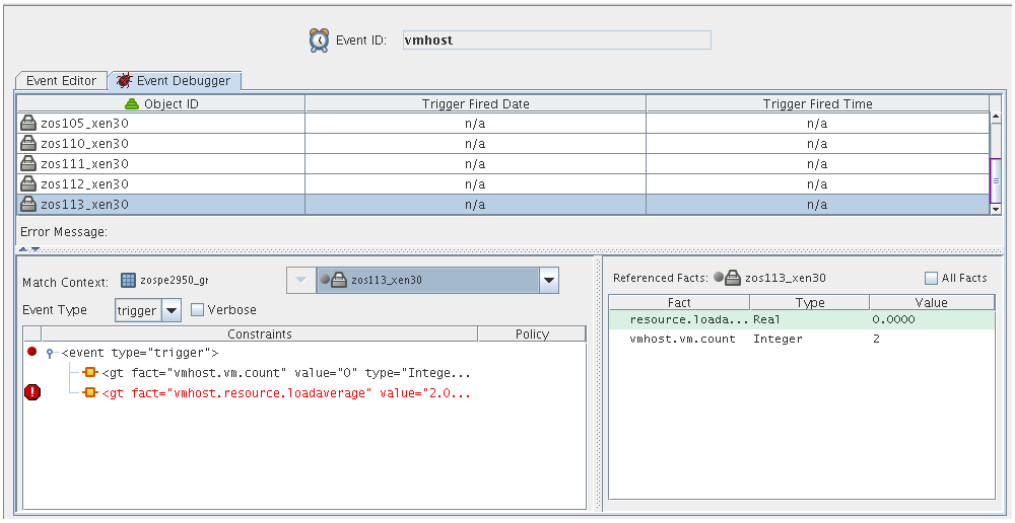

The information in this section describes the various parts of the debugger.

- Section 15.3.1, “Constraints Table,” on page 182
- Section 15.3.2, “The Facts Table,” on page 184

15.3.1 Constraints Table

The left side of the Event Debugger panel includes the Constraints Table. The rules that define the `<trigger>` and `<reset>` of an Event are defined using the same XML constraint syntax used in policies.

Figure 15-5 *The Constraints Table Area of the Events Debugger*

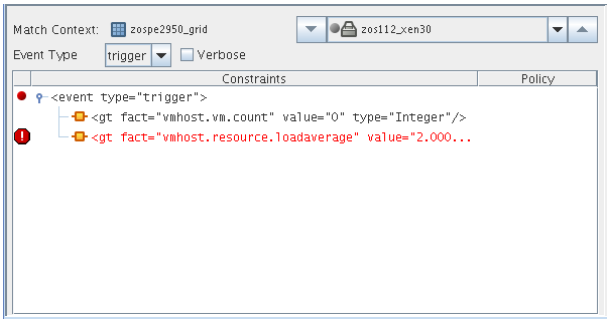

The Constraints Table has several parts:

- “Match Context” on page 183
- “Event Type List” on page 183
- “Verbose Check Box” on page 183
- “Constraints List (Tree)” on page 183

Match Context

Depending on the Event, the debugger identifies the Grid objects (`Job`, `Jobinstance`, `Resource`, `Repository`, `User`, or `VMHost`) that define the context of the trigger or reset rules specified in the Event XML.

Figure 15-6 *The Match Context Area of the Constraints Table View in the Event Debugger*



Event Type List

Select one of the Event types in the drop-down list to debug how either the `<trigger>` or `<reset>` rules are being applied. Constraint types in the list are disabled (dimmed) if they do not apply to the Event that you are debugging.

- **trigger:** The rules defining the conditions (through a constraint) in which an Event is generated.
- **reset:** The rules defining the conditions (through a constraint) in which an Event is reset (that is, able to be triggered again).

Verbose Check Box

When you select the Verbose check box, additional constraint information is displayed.

Constraints List (Tree)

The Constraints tree, whichh is a column in the constraints table, lists the constraints that constitute a particular rule in a hierarchical view.

Each constraint is flagged with an icon to signify whether it “passes” or not. A constraint flagged with an exclamation point indicates a constraint causing the rule to not “pass.”

Right-click a constraint to display a menu where you can perform one of the following actions:

Show Admin View: Selects the currently evaluated Grid object in the Explorer Tree and displays its Info/Groups administration information view.

Disable Constraint: Passes the constraint, regardless of how it evaluates.

Enable All Constraints: Re-enables any disabled constraints.

NOTE: The right-click menu is available only when you select specific constraints.

15.3.2 The Facts Table

The Facts Table view displays the facts referenced in the Constraint Tree view for a specified Event. Selecting a rule containing a particular fact in the Constraint tree automatically selects that fact and its current value in the table.

Figure 15-7 *The Facts Table View*

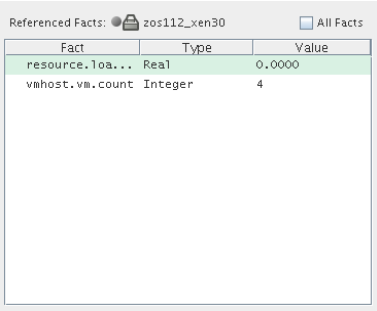

If you right-click a column head in this table, a menu is launched where you can select the columns that you want to display.

Figure 15-8 *Menu Used to Select the Columns Displayed in the Facts Table View of the Policy Debugger*

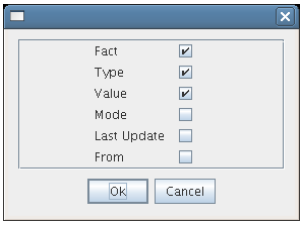

All Facts Check Box: If you select the **All Facts** check box at the top of the Facts Table view, all facts for the selected Grid object, as well as those for the Server itself (`matrix.* facts`) are displayed.

If you right-click a fact, you have the option of adding a new fact, deleting the selected fact, or viewing/editing the selected fact, if the fact is editable or can be deleted.

16 The Metrics Facility

The the Orchestration Server Metrics Facility collects, aggregates, and allows simple fact-based retrieval of metric values by jobs and computed facts (via JDL), policy constraints, and Event triggers on a per-resource basis. This provides aggregated metrics generated by gmond without the need for the gmetad Ganglia service. Note that gmetad can still be used in parallel for aggregating gmond reported metrics for visualization purposes.

- Section 16.1, “Metrics Facility Functionality,” on page 185
- Section 16.2, “Ganglia Metrics,” on page 185
- Section 16.3, “How Does the Metrics Facility Impact Orchestration Server Performance?,” on page 186
- Section 16.4, “RRD Definition Using Deployable .metric Files,” on page 187
- Section 16.5, “Query of Aggregated Metric Values,” on page 189
- Section 16.6, “MetricsManager MBean API,” on page 190
- Section 16.7, “Using the Metrics Facility in the Orchestration Console,” on page 191

16.1 Metrics Facility Functionality

The Metrics Facility provides the following functionality:

- Collection of gmond provided metrics using the Orchestration Agent.
- Retrieval of instantaneous metric values via `resource.metrics.<METRIC_NAME>` fact space, where `<METRIC_NAME>` is the name of the metric.
- Deployable Round Robin Database (RRD) (data aggregation) definition using XML `.metric` files, which allows flexible definition of aggregation periods. For example, using `resource.metrics.<METRIC_NAME>.10_minute.average` as a 10-minute aggregation period separate from `resource.metrics.<METRIC_NAME>.1_hour.average`.
- Retrieval of an array of aggregated metric values using `resource.metrics.<METRIC_NAME>.xxx.values`.
- Zero-configuration for core Ganglia metrics. The Orchestration Agent automatically discovers if gmond is running on a resource, and the Orchestration Server collects and, if a `.metric` file is configured, aggregates those metrics.
- Persistence of collected RRD data across server restart and high availability fail-over conditions.

16.2 Ganglia Metrics

Not all Ganglia metrics are suitable for aggregation, such as those of “String” type. By default, only the 24 metrics listed in the following table are supported (either of type “Real” or type “Integer”):

Table 16-1 Supported Ganglia Metrics

Metric	Type
bytes_in	Real
bytes_out	Real
cpu_aidle	Real
cpu_idle	Real
cpu_nice	Real
cpu_system	Real
cpu_user	Real
cpu_wio	Real
disk_free	Integer
disk_total	Integer
load_fifteen	Real
load_five	Real
load_one	Real
mem_buffers	Integer
mem_cached	Integer
mem_free	Integer
mem_shared	Integer
mem_total	Integer
pkts_in	Real
pkts_out	Real
proc_run	Integer
proc_total	Integer
swap_free	Integer
swap_total	Integer

16.3 How Does the Metrics Facility Impact Orchestration Server Performance?

The Metrics Facility balances flexibility and minimizing the impact on overall Orchestration Server performance. There is no continuous parsing of XML on the server. Instead, parsing of gmond-generated XML is performed by each managed resource.

This section includes the following information:

- Section 16.3.1, “I/O Contention,” on page 187
- Section 16.3.2, “Too Many Open Files,” on page 187

16.3.1 I/O Contention

By default, RRD-based data aggregation is file-based. Because of the frequency of updates and queries of RRD files, this poses a significant performance issue. The Metrics Facility minimizes I/O contention by using in-memory caching and batched write operations to avoid I/O contention and the resulting performance degradation.

16.3.2 Too Many Open Files

The the Orchestration Server Monitoring Server uses a “one RRD per resource” approach, where a RRD contains the AVERAGE, MIN, and MAX RRAs for multiple metrics (DS). In contrast, the Metrics Facility takes an “inside out” approach, which results in “one RRD per metric.” For 1,000 agents reporting 24 Ganglia metrics, this reduces the number of files dramatically (from 24,000 in the “one RRD per metric, per resource” case, and 1,000 in the “one RRD per resource” case, to 24 in the “one RRD per metric” case). This approach avoids a “too many open files” condition.

NOTE: A “too many open files” condition occurs when the default maximum file descriptors available to a process launched from the Linux shell is exceeded.

16.4 RRD Definition Using Deployable .metric Files

Definition of the “aggregation” functions performed by the Metrics Facility’s internal RRD data structures are customizable using deployable XML `.metric` definition files. This accommodates a flexible configuration of the following:

- The Ganglia-reported metrics to include in aggregated data structures
- The data aggregation periods that are of specific interest

The deployable definition files, one per metric to be aggregated, consist of the following:

- The name of the “instantaneous” metric to be aggregated, for example, `load_one`
- An optional description of the metric to be aggregated.
- A “heartbeat” value that governs the updates to the contained RRAs.
- One or more named Periods (corresponding to the RRAs to be created) with an optional description.
- For each Period, the number of data points to aggregate (steps), the number of aggregated data points to archive (rows), and the “xff' (x-files factor) allowed for the metric.

  NOTE: The xff determines how many of the samples can be NaN for the consolidated sample to be considered NaN. Usually, this is set to 0.5, or 50%.).

When an RRD is defined through deployment of its definition file, three RRAs are created for each Period: AVERAGE, MAX, and MIN. A new DS (datasource) is added to the RRD for each resource reporting the metric to be aggregated. This requires the RRD file to be re-created each time a new resource begins reporting a given metric and the previously aggregated values copied from the old RRD to the new one. This approach enhances performance and flexibility, but the RRD file is not of fixed size: Over time, the RRD grows or shrinks as new resources are added to the system or are deleted from the Orchestration Server model.

NOTE: The RRD is actually re-created with a new DS added for each new resource and the “old” RRA’s data copied into it.

Deleting a Resource Grid object removes its DS from the RRD file (actually, from all RRDs with metrics reported by that resource).

One optimization you can implement for storing the smallest Period (consisting of a single step) is to create only a single RRA (vs. three), because the average of a single datapoint is equal to the maximum and minimum of a single datapoint.

NOTE: The RRD files created by the Java rrd4j library are not binary compatible with RRD files generated by the rrdtool used by gmetad. They are however portable across operating system architectures (e.g., 32-bit bigendian vs. 64-bit little-endian) which is not possible with traditional RRD files created using rrdtool.

16.4.1 XML Format for Deployable .metric Definitions

An example of the format of the deployable RRD definition is shown below.

```
<metric name="load_one" heartbeat="120"
        description="Ganglia oneminute load average">


    <period name="1_minute" steps="1" rows="60" xff=".5"
        description="1 hour worth of 1 minute (raw) data"/>


    <period name="5_minute" steps="5" rows="12" xff=".5"
        description="1 hour worth of 5 minute aggregations"/>


    <period name="10_minute" steps="10" rows="72" xff=".5"
        description="12 hours worth of 10 minute aggregations"/>


    <period name="1_hour" steps="60" rows="24" xff=".5"
        description="1 days worth of 1 hour aggregations"/>


</metric>
```

This example creates an RRD for the `load_one` metric, with four aggregation periods (RRAs) called `1_minute`, `5_minute`, `10_minute`, and `1_hour`. The default sample (RRD update) time is one minute, so the `10_minute` aggregation period has 10 steps. 72 “rows” of aggregated datapoints are retained in the RRA before the oldest is dropped off, representing 12 hours (12 * 60 / 10) worth of data.

NOTE: The `1_minute` period is not a true aggregation because the default sample (RRD update) time is also one minute. In this case, the “raw” datapoints are stored for historical reference.

16.5 Query of Aggregated Metric Values

Aggregated metric values can be queried similar to the instantaneous values either from JDL, from within a policy/event constraint, or from an array of multiple metrics values. This section includes the following information:

- Section 16.5.1, “Example of a JDL Query for Aggregated Metric Values,” on page 189
- Section 16.5.2, “Example of a Policy Constraint or Event Constraint Using Aggregated Metric Values,” on page 189
- Section 16.5.3, “Example of Using Non-aggregated (“Raw”) Historical Metric Values,” on page 189

16.5.1 Example of a JDL Query for Aggregated Metric Values

```
r = getMatrix().getResource("local_vmhost")

print "resource.id: %s" % (r.getFact("resource.id"))

print "load_one.10_minute.average: %s" %
( r.getFact("resource.metrics.load_one.10_minute.average") )
```

16.5.2 Example of a Policy Constraint or Event Constraint Using Aggregated Metric Values

```
   <event>
      <context>
         <vmhost />
      </context>
      <trigger>
         <gt fact="vmhost.vm.count" value="0" />
        <gt fact="vmhost.resource.metrics.load_one.10_minute.average" value="2" />
      </trigger>
      <reset>
        <lt fact="vmhost.resource.metrics.load_one.10_minute.average" value=".5" /
>
      </reset>
   </event>
```

The consolidation functions (AVERAGE,MAX, and MIN) are supported for each defined aggregation period. In RRD terminology, this means that for each metric, there are three RRAs defined for each “period” element in the `.metric` deployable definition.

16.5.3 Example of Using Non-aggregated (“Raw”) Historical Metric Values

You can query for an array of “raw” values that constitute the aggregated datapoints for a given RRA within the RRD data structure by appending `.values` to the factname representing the metric (and period) of interest. For example, to print all the MIN values collected for the `10_minute` aggregation period, the JDL is

```
r = getMatrix().getResource("local_vmhost")

print "resource.id: %s" % (r.getFact("resource.id"))
print "load_one.10_minute.min.values: %s" %

( r.getFact("resource.metrics.load_one.10_minute.min.values") )
```

Similarly, to print the array of AVERAGE and MAX values collected, the JDL is

```
print "load_one.10_minute.average.values: %s" %
( r.getFact("resource.metrics.load_one.10_minute.average.values") )

print "load_one.10_minute.max.values: %s" %
( r.getFact("resource.metrics.load_one.10_minute.max.values") )
```

16.6 MetricsManager MBean API

The MetricsManager facility exposes a small number of methods for disabling and enabling a RRD file creation and update, and for query of both instantaneous and aggregated metric values. The facility also allows a UI to query for information needed for populating a pull-down list that could include, for example, valid metric names for a specified resource or valid aggregation periods for such a metric. The API also currently provides a way to fetch a “running average” of the raw datapoints for a metric (a cached value that does not require a fetch operation from an RRD) and server-side generation of a simple graph of metric data (requires a fetch from an RRD file).

This section includes the following information:

- Section 16.6.1, “MBean Methods Exposed by the MetricsManager Facility,” on page 190
- Section 16.6.2, “The MetricsDeployer Facility,” on page 191

16.6.1 MBean Methods Exposed by the MetricsManager Facility

The following table lists the MBean methods that are exposed by the MetricsManager Facility.

Table 16-2 *MBean Methods Exposed by the MetricsManager Facility*

MBean Method	Description
`setMetricsEnabled()`	A parameter of false disables the creation or update of RRD files.
`getObjectNames()`	Returns a list of current deployed metric definitions
`update()`	Updates RRDs with instantaneous values collected by the Metrics Facility.
`getFact()`	Given the fully qualified fact name, fetches the value of a specified resource’s fact.
`getMetricNames()`	Fetches the metric names relevant for a specified resource .
`getPeriodNames()`	Fetches the names of periods (RRAs) defined for a specified metric.
`getRunningAverage()`	Fetches the running average of a metric for a specified resource.
`getRrdGraphic()`	Used by the Orchestration Console to fetch a server-generated ImageIconInfo object.

16.6.2 The MetricsDeployer Facility

The MetricsDeployer deployment facility parses the `.metric` definition files and creates the associated grid objects that are used to maintain the metadata related to the RRD used for aggregation.

16.7 Using the Metrics Facility in the Orchestration Console

The Orchestration Console supports the Metrics Facility in the following ways:

- **Metrics object:** A Metrics object is deployed in the Explorer tree. Use the right-click menu to display the “deploy” and “undeploy” actions. Pre-defined `.metric` files are located in the `/opt/novell/zenworks/zos/server/components/metrics` folder, or you can create a new `.metric` file and paste in the XML as shown in Section 16.4.1, “XML Format for Deployable .metric Definitions,” on page 188.

  Metrics objects are listed by their deployment name, which may or may not be the same as the name of the actual metric. This potentially allows multiple, separately deployable, RRD definitions for a single “instantaneous” metric, with different aggregation periods defined.
- **Metrics Editor:** Selecting the Metrics object opens a free-form XML editor in the admin view of the client. The editor is similar to the Event” editor for viewing RRD definition XML.

II Orchestration Server Tools

This section includes details about the tools you can use to administer the Cloud Manager Orchestration Server.

- Chapter 17, “Orchestration Server zos Command Line Tool,” on page 195
- Chapter 18, “Orchestration Server zosadmin Command Line Tool,” on page 243

17 Orchestration Server zos Command Line Tool

The zos command line tool can be used by logged-in users. To use the command line tool, each user requires a login name and password that are created by the administrator for the Orchestration Server. The zos command line tool can be used to log in to the Cloud Manager Orchestration Server, to submit jobs, monitor jobs, and perform other user functions.

This section includes information about the following:

- Section 17.1, “List of zos Commands,” on page 195
- Section 17.2, “Getting Started with the zos Command,” on page 196
- Section 17.3, “Details, Usage, and Syntax Examples of zos Commands,” on page 199

17.1 List of zos Commands

The following table includes a list of zos commands and a description for each command.

Table 17-1 *Available zos Commands and Their Descriptions*

Command	Description
cancel	Cancel a running job
cat	Displays the contents of a datagrid file
copy	Copy files and directories to and from the datagrid
delete	Delete files and directories in the datagrid
dir	List files and directories in the datagrid
event	Send an event to a running job
fail	Directs the job to end in failure
head	Displays the first part of a datagrid file
help	Displays help for any of the commands in this list
info	Display information about a grid object
jobinfo	Display important information for a job
joblist	List all of the available jobs and arguments
jobs	List the running queued or historical job
listen	Listen to events generated by a running job
log	Displays the log for the specified job
login	Log in to the grid

Command	Description
logout	Log out of the grid
mkdir	Make a new directory in the datagrid
move	Move files and directories in the datagrid
password	Change user password on current grid
pause	Pause a running job
priority	Change the priority of a running job
resume	Resume a running job
run	Run a job
search	Perform grid object search
status	Display the status of a running or previously run job
tail	Displays the end of a datagrid file

17.2 Getting Started with the zos Command

This section is for advanced users or those who prefer a command line interface (CLI) to manage jobs.

All zos commands begin with `zos` on the command line. The general format for a zos command is `zos` followed by the command name, followed by command line parameters, if needed:

`zos` *`command`* [*`command_parameters`*]

Before you use the zos CLI, ensure that your path is correctly pointing to the Orchestration Server tools directory. On Windows platforms, the path is automatically set when you launch the Orchestration command prompt window. If the path is set, you then need to log in to the Orchestration Server before the commands work.

This section includes the following information:

- Section 17.2.1, “Logging In to the Orchestration Server,” on page 196
- Section 17.2.2, “Running a Job,” on page 197
- Section 17.2.3, “Monitoring a Job,” on page 197

17.2.1 Logging In to the Orchestration Server

The administrator for the Orchestration Server provides a user name, a password, and the name for the Orchestration Server. The following shows a sample login sequence. Italics are used to indicate variable names.

```
>zos login -user=username Orchestration_server_name
Please enter current password for 'Orchestration_server_name': *******
Logged into Orchestration_server_name_zos> as username
```

The login should now be complete.

Login information is stored in the home directory, so further `zos` commands on this server use the saved login information. To operate on a different Orchestration Server, run `zos login` again with the new server information.

17.2.2 Running a Job

Use the `run` command to start a job. For example, if you want to run the “quickie” job, you would enter the following command:

```
>zos run quickie
```

If the job submission was successful, a JobID is displayed:

```
JobID: tester.1.quickie
```

To specify parameters when you start a job, add the parameters to the command line. For example, if you wanted to run the quickie job, specifying the `numJoblets` parameter with a value of 100, you would enter the following command:

```
>zos run quickie numJoblets=100
JobID: tester.2.quickie
```

The Orchestration Server keeps track of your last started job so that further zos commands can default to the last started job. To operate on a different job, specify the Job ID as a command line parameter.

17.2.3 Monitoring a Job

A job can be monitored or managed when it is started or after it is started. To view the status of the last started job, use the `status` command:

```
>zos status
Running
```

The server returns the status immediately. In the preceding example, the job is `Running`, which means the job is still being processed. Possible status messages include the following:

```
Submitted
Queued
Started
Running
Paused
Completing
Cancelling
Failing
Completed
Cancelled
Failed
```

You can retrieve detailed status information about a job by using the `-e` option in the command line, as shown in the following example for the last started job:

```
>zos status -e
Job Status for tester.quickie.2
------------------------------
             State: Running (waiting for resources)
    Resource Count: 0              (0 this job)
  Percent Complete: 1%
         Queue Pos: n/a
   Child Job Count: 0              (0 this job)

     Instance Name: quickie
          Job Type: quickie
              Memo:
          Priority: medium
         Arguments: numJoblets=100

       Submit Time: 5/01/2012 12:20:19
     Delayed Start: n/a
        Start Time: 5/01/2012 12:20:19
          End Time: n/a  (estimated)
      Elapsed Time: 0:00:05
        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: $0.0000        ($0.0000 this job)
         Burn Rate: $0.0000/hr     ($0.0000/hr this job)

       Cancel Type: <none>
         Job Error: <none>
   Job Error Count: 0              (0 this job)
Joblet Retry Count: 0              (0 this job)
  Node Error Count: 0              (0 this job)
```

17.3 Details, Usage, and Syntax Examples of zos Commands

This section includes a detailed list of the `zos` commands you can use. It also includes examples for using these commands and shows the syntax of typical commands.

NOTE: Items shown in brackets [ ] are optional. Items shown in *italics* are contextual examples.

The section is organized alphabetically by the command names:

- “cancel” on page 200
- “cat” on page 202
- “copy” on page 203
- “delete” on page 205
- “dir” on page 206
- “event” on page 207
- “fail” on page 208
- “head” on page 209
- “info” on page 211
- “jobinfo” on page 213
- “joblist” on page 215
- “jobs” on page 216
- “listen” on page 218
- “log” on page 219
- “login” on page 221
- “logout” on page 223
- “mkdir” on page 224
- “move” on page 226
- “password” on page 227
- “pause” on page 228
- “priority” on page 230
- “resume” on page 232
- “run” on page 233
- “search” on page 235
- “status” on page 237
- “tail” on page 239

cancel

This command cancels a running job.

Syntax

```
zos cancel [jobid] [--reason=]
```

Arguments

jobid

The user must enter the jobid that he or she is interested in. The default for this argument is the last job run.

Options

-m, --reason=<value>

Specifies the reason to log for this action. The user must enter a value with this option.

-C, --cols=<value>

Specifies the terminal width for formatting. The user must enter a value with this option.

-h, --help

Displays a help message for this operation.

-d, --debug

Turns on debug output during this operation.

-v, --verbose

Turns on verbose output during this operation.

-z, --tls=<value>

Force TLS/SSL encryption [on|off]. The user must enter a value with this option.

-Z, --tlsPort=<value>

Port number for TLS/SSL connections. The user must enter a value with this option.

-y, --tlsYesNew

Force acceptance of new TLS server certificates.

-Y, --tlsYesAll

Force acceptance of all TLS server certificates.

Examples

Example 1

To cancel the current job for the login user, use the following command:

```
zos cancel
```

Example 2

To cancel the named job instance <jobid> for the login user, use the following command:

```
zos cancel <jobid>
```

cat

This command displays the contents of a datagrid file.

Syntax

`zos cat` *`path`*

Arguments

path

The path on the datagrid.

Options

-C, --cols=<value>

Specifies the terminal width for formatting. The user must enter a value with this option.

-h, --help

Displays a help message for this operation.

-d, --debug

Turns on debug output during this operation.

-v, --verbose

Turns on verbose output during this operation.

-z, --tls=<value>

Force TLS/SSL encryption [on|off]. The user must enter a value with this option.

-Z, --tlsPort=<value>

Port number for TLS/SSL connections. The user must enter a value with this option.

-y, --tlsYesNew

Force acceptance of new TLS server certificates.

-Y, --tlsYesAll

Force acceptance of all TLS server certificates.

Examples

Example 1

To display the JDL file for the `quickie` job, use the following command:

```
zos cat grid:///!quickie/quickie.jdl
```

Example 2

To display the `job.log` file for the active or last run job, use the following command:

```
zos cat grid:///^/job.log
```

copy

This command copies files and directories to and from the data zos.

Syntax

```
zos copy source dest [--recursive]
```

Arguments

source

The datagrid source from which the files or directories are to be copied.

dest

The datagrid destination to which the files or directories are to be copied.

Options

-r, --recursive

Processes directories recursively.

-C, --cols=<value>

Specifies the terminal width for formatting. The user must enter a value with this option.

-h, --help

Displays a help message for this operation.

-d, --debug

Turns on debug output during this operation.

-v, --verbose

Turns on verbose output during this operation.

-z, --tls=<value>

Force TLS/SSL encryption [on|off]. The user must enter a value with this option.

-Z, --tlsPort=<value>

Port number for TLS/SSL connections. The user must enter a value with this option.

-y, --tlsYesNew

Force acceptance of new TLS server certificates.

-Y, --tlsYesAll

Force acceptance of all TLS server certificates.

Examples

Example 1

To copy a local file into a previously created datagrid subdirectory, use the following command:

```
zos copy local.txt grid:///images
```

Example 2

To copy a file from the current job instance directory to a local file, use the following command:

```
zos copy grid:///^/results.txt local.txt
```

Example 3

To copy a file from a named job instance directory to a local file, use the following command:

```
zos copy grid:///^user.myjob.1024/results.txt local.txt
```

Example 4

To copy a local file into the job directory for job *myjob*, use the following command:

```
zos copy local.txt grid:///!myjob
```

Example 5

To copy a local file into the current user's home directory, use the following command:

```
zos copy local.txt grid:///~/local.txt
```

Example 6

To copy a local file into another user's (barney) home directory, use the following command:

```
zos copy local.txt grid:///~barney/local.txt
```

Example 7

To copy a local file into a subdirectory of the job directory, use the following command:

```
zos copy local.txt grid:///!myjob/subdir
```

Example 8

To copy a local file into the current job instance directory, use the following command:

```
zos copy local.txt grid:///^/
```

Example 9

To recursively copy a local directory to the datagrid root, creating a new directory, use the following command:

```
zos copy /home/tester/myDir grid:/// -r
```

delete

This command deletes files and directories in the datagrid.

Syntax

```
zos delete path
```

Arguments

path

The path on the datagrid.

Options

-C, --cols=<value>

Specifies the terminal width for formatting. The user must enter a value with this option.

-h, --help

Displays a help message for this operation.

-d, --debug

Turns on debug output during this operation.

-v, --verbose

Turns on verbose output during this operation.

-z, --tls=<value>

Force TLS/SSL encryption [on|off]. The user must enter a value with this option.

-Z, --tlsPort=<value>

Port number for TLS/SSL connections. The user must enter a value with this option.

-y, --tlsYesNew

Force acceptance of new TLS server certificates.

-Y, --tlsYesAll

Force acceptance of all TLS server certificates.

Examples

Example 1

To delete a file under the datagrid directory images, use the following command:

```
zos delete grid:///images/myFile
```

dir

This command lists files and directories in the datagrid.

Syntax

```
zos dir path
```

Arguments

path

The path on the datagrid.

Options

-C, --cols=<value>

Specifies the terminal width for formatting. The user must enter a value with this option.

-h, --help

Displays a help message for this operation.

-d, --debug

Turns on debug output during this operation.

-v, --verbose

Turns on verbose output during this operation.

-z, --tls=<value>

Force TLS/SSL encryption [on|off]. The user must enter a value with this option.

-Z, --tlsPort=<value>

Port number for TLS/SSL connections. The user must enter a value with this option.

-y, --tlsYesNew

Force acceptance of new TLS server certificates.

-Y, --tlsYesAll

Force acceptance of all TLS server certificates.

Examples

Example 1

To list files for the current job instance, use the following command:

```
zos dir grid:///^
```

Example 2

To list the deployed files for the job named *animation*, use the following command:

```
zos dir grid:///\!animation
```

event

This command sends an event to a running job.

Syntax

`zos event` *`jobid eventname`*

Arguments

jobid

The user must enter the jobid that he or she is interested in. The default for this argument is the last job run.

eventname

The user must enter the name of the event he or she wants to send to the job.

Options

-C, --cols=<value>

Specifies the terminal width for formatting. The user must enter a value with this option.

-h, --help

Displays a help message for this operation.

-d, --debug

Turns on debug output during this operation.

-v, --verbose

Turns on verbose output during this operation.

-z, --tls=<value>

Force TLS/SSL encryption [on|off]. The user must enter a value with this option.

-Z, --tlsPort=<value>

Port number for TLS/SSL connections. The user must enter a value with this option.

-y, --tlsYesNew

Force acceptance of new TLS server certificates.

-Y, --tlsYesAll

Force acceptance of all TLS server certificates.

Examples

Example 1

To send the event `MyEvent` to the current job with a simple parameter, use the following command:

```
zos event MyEvent myarg=12345
```

fail

This command directs the job to end in failure.

Syntax

`zos fail [jobid] [--reason=]`

Arguments

jobid

The user must enter the jobid that he or she is interested in. The default for this argument is the last job run.

Options

-m, --reason=<value>

Specifies the reason to log for this action. The user must enter a value with this option.

-C, --cols=<value>

Specifies the terminal width for formatting. The user must enter a value with this option.

-h, --help

Displays a help message for this operation.

-d, --debug

Turns on debug output during this operation.

-v, --verbose

Turns on verbose output during this operation.

-z, --tls=<value>

Force TLS/SSL encryption [on|off]. The user must enter a value with this option.

-Z, --tlsPort=<value>

Port number for TLS/SSL connections. The user must enter a value with this option.

-y, --tlsYesNew

Force acceptance of new TLS server certificates.

-Y, --tlsYesAll

Force acceptance of all TLS server certificates.

head

This command displays the first part of a datagrid file.

Syntax

```
zos head [--bytes=] [--lines=] [--quiet] path
```

Arguments

path

The path on the datagrid.

Options

-c, --bytes=<value>

Generate file output of the specified number of bytes only. The user must enter a value with this option.

-n, --lines=<value>

Generate file output of the specified number of lines only. The user must enter a value with this option.

-q, --quiet

Suppress all warning and informational messages.

-C, --cols=<value>

Specifies the terminal width for formatting. The user must enter a value with this option.

-h, --help

Displays a help message for this operation.

-d, --debug

Turns on debug output during this operation.

-v, --verbose

Turns on verbose output during this operation.

-z, --tls=<value>

Force TLS/SSL encryption [on|off]. The user must enter a value with this option.

-Z, --tlsPort=<value>

Port number for TLS/SSL connections. The user must enter a value with this option.

-y, --tlsYesNew

Force acceptance of new TLS server certificates.

-Y, --tlsYesAll

Force acceptance of all TLS server certificates.

Examples

Example 1

To display the first ten lines of the job log for the last job, use the following command:

```
zos head grid:///^/job.log
```

info

This command displays information about a Grid object.

Syntax

```
zos info objectName --type= [--detail] [--xml] [--python]
```

Arguments

objectName

The user must enter the Grid object ID or the Grid object display name that he or she is interested in.

Options

-t, --type=<value>

The user must specify one of the following Grid object types:

- `user`
- `resource`
- `job`
- `jobinstance`
- `vmhost`
- `grid`
- *`XXX`*`group`

-e, --detail

Displays detailed information for each Grid object, rather than just its ID.

-x, --xml

Displays information in formatted XML document instead of in human readable form.

-P, --python

Displays fact information as parseable python dictionary.

-C, --cols=<value>

Specifies the terminal width for formatting. The user must enter a value with this option.

-h, --help

Displays a help message for this operation.

-d, --debug

Turns on debug output during this operation.

-v, --verbose

Turns on verbose output during this operation.

-z, --tls=<value>

Force TLS/SSL encryption [on|off]. The user must enter a value with this option.

-Z, --tlsPort=<value>

Port number for TLS/SSL connections. The user must enter a value with this option.

-y, --tlsYesNew

Force acceptance of new TLS server certificates.

-Y, --tlsYesAll

Force acceptance of all TLS server certificates.

Examples

Example 1

To get detailed information on user *paul*, use the following command:

```
zos info paul -t user
```

Example 2

To get detailed information in XML format on resource `vmh1`, use the following command:

```
zos info vmh1 -t resource -e -x
```

Example 3

If the display name of vmh1 is *somethingelse,* you can also use the following command to give you the same results:

```
zos info somethingelse -t resource -e -x
```

jobinfo

This command displays important information about a given job for Cloud Manager Orchestration.

Syntax

```
zos jobinfo [--detail] [--xml] jobname
```

Arguments

jobname

The user must enter the name of the job for which he or she is seeking information.

Options

-e, --detail

Displays detailed job information for each job.

-x, --xml

Displays jobargs in XML format.

-C, --cols=<value>

Specifies the terminal width for formatting. The user must enter a value with this option.

-h, --help

Displays a help message for this operation.

-d, --debug

Turns on debug output during this operation.

-v, --verbose

Turns on verbose output during this operation.

-z, --tls=<value>

Force TLS/SSL encryption [on|off]. The user must enter a value with this option.

-Z, --tlsPort=<value>

Port number for TLS/SSL connections. The user must enter a value with this option.

-y, --tlsYesNew

Force acceptance of new TLS server certificates.

-Y, --tlsYesAll

Force acceptance of all TLS server certificates.

Examples

Example 1

To display a quick list of job parameters for the `quickie` sample job, the user would enter this command:

```
zos jobinfo quickie
```

joblist

This command lists all of the available jobs and arguments.

Syntax

```
zos joblist [--detail] [--description]
```

Options

-e, --detail

Displays detailed job information for each job.

-D, --description

Displays a description for each job.

-C, --cols=<value>

Specifies the terminal width for formatting. The user must enter a value with this option.

-h, --help

Displays a help message for this operation.

-d, --debug

Turns on debug output during this operation.

-v, --verbose

Turns on verbose output during this operation.

-z, --tls=<value>

Force TLS/SSL encryption [on|off]. The user must enter a value with this option.

-Z, --tlsPort=<value>

Port number for TLS/SSL connections. The user must enter a value with this option.

-y, --tlsYesNew

Force acceptance of new TLS server certificates.

-Y, --tlsYesAll

Force acceptance of all TLS server certificates.

Examples

Example 1

To generate a listing of all runnable jobs, you would use the following command:

```
zos joblist -e
```

jobs

This command lists the running queued or historical jobs.

Syntax

```
zos jobs [--all] [--detail] [--errors] [--submit=] [--hours=] [--active] [--count=]
[--parent=]
```

Options

-a, --all

Lists all users' jobs, not just the logged in user's jobs.

-e, --detail

Displays detailed job information for each job.

-E, --errors

Displays details of job and node errors.

-s, --submit=<value>

Specifies a time when the user wants to see historical jobs. The user must enter a value with this option.

-h, --hours=<value>

Specifies submit time in the form of last *xxx* hours. If specified, this option overrides the `submit` option. The user must enter a value with this option.

-A, --active

Lists only the currently active or running jobs.

-c, --count=<value>

Sets the maximum number of historical jobs retrieved. The user must enter a value with this option.

-p, --parent=<value>

Specifies the parent job id. If this option is set, only the child jobs of this parent are displayed. The user must enter a value with this option.

-C, --cols=<value>

Specifies the terminal width for formatting. The user must enter a value with this option.

-h, --help

Displays a help message for this operation.

-d, --debug

Turns on debug output during this operation.

-v, --verbose

Turns on verbose output during this operation.

-z, --tls=<value>

Force TLS/SSL encryption [on|off]. The user must enter a value with this option.

-Z, --tlsPort=<value>

Port number for TLS/SSL connections. The user must enter a value with this option.

-y, --tlsYesNew

Force acceptance of new TLS server certificates.

-Y, --tlsYesAll

Force acceptance of all TLS server certificates.

Examples

Example 1

To display a detailed list of running jobs for all users, the user would enter the following command:

```
zos jobs -a -e
```

Example 2

To display a list of queued, running, and recently run jobs, the user would enter the following command:

```
zos jobs
```

Example 3

To display a list of jobs he or she submitted since 2:20 p.m., the user would enter the following command:

```
zos jobs --submit="2:20 PM"
```

NOTE: The date and time values should be specified as shown in the example—in the format `MM/DD/YY hour:minute AM/PM` and enclosed in quotes—`"11/21/11 4:04 PM"`.

Example 4

To display a list of all active subjobs for the jobid `barney.load.677`, the user would enter the following command:

```
zos jobs -A --parent=barney.load.677 --all
```

Example 5

To display a list of all his or her jobs run in the past 24 hours, the user would enter the following command:

```
zos jobs -h 24
```

listen

This command starts a process to listen to events generated by a running job.

Syntax

```
zos listen [jobid] [--detail=]
```

Arguments

jobid

The user must enter the jobid that he or she is interested in. The default for this argument is the last job run.

Options

-e, --detail

Displays detailed job information for each job.

-C, --cols=<value>

Specifies the terminal width for formatting. The user must enter a value with this option.

-h, --help

Displays a help message for this operation.

-d, --debug

Turns on debug output during this operation.

-v, --verbose

Turns on verbose output during this operation.

-z, --tls=<value>

Force TLS/SSL encryption [on|off]. The user must enter a value with this option.

-Z, --tlsPort=<value>

Port number for TLS/SSL connections. The user must enter a value with this option.

-y, --tlsYesNew

Force acceptance of new TLS server certificates.

-Y, --tlsYesAll

Force acceptance of all TLS server certificates.

log

This command displays the log for the specified job.

Syntax

```
zos log [jobid]
```

Arguments

jobid

The user must enter the jobid that he or she is interested in. The default for this argument is the last job run.

Options

-C, --cols=<value>

Specifies the terminal width for formatting. The user must enter a value with this option.

-h, --help

Displays a help message for this operation.

-d, --debug

Turns on debug output during this operation.

-v, --verbose

Turns on verbose output during this operation.

-z, --tls=<value>

Force TLS/SSL encryption [on|off]. The user must enter a value with this option.

-Z, --tlsPort=<value>

Port number for TLS/SSL connections. The user must enter a value with this option.

-y, --tlsYesNew

Force acceptance of new TLS server certificates.

-Y, --tlsYesAll

Force acceptance of all TLS server certificates.

Examples

Example 1

To display the log for the last job, use the following command:

```
zos log
```

Example 2

To display the log for job ID *itsajob*, use the following command

```
zos log itsajob
```

login

This command logs the user into the Orchestration grid.

Syntax

```
zos login [servername] [--check] [--user=] [--passwd=] [--port=] [--env=]
```

Arguments

servername

The user must enter the hostname of the Cloud Manager Orchestration Server he or she is logging into. The default is `localhost`.

Options

-c, --check

Check and report existing login. Ignores other options.

-u, --user=<value>

Username known by the Orchestration Server. The user must enter a value with this option. The default is `login`.

-p, --passwd=<value>

User password for authentication. The user must enter a value with this option.

-P, --port=<value>

The port number used for communication between the Orchestration Server Portal and the Orchestration Server. The user must enter a value with this option.

-V, --env=<value>

Upload the local environment and store on the server. The user must enter a value with this option.

-C, --cols=<value>

The specification for the terminal width for formatting. The user must enter a value with this option.

-h, --help

Displays a help message for this operation.

-d, --debug

Turns on debug output during the login operation.

-v, --verbose

Turns on verbose output during the login operation.

-z, --tls=<value>

Force TLS/SSL encryption [on|off]. The user must enter a value with this option.

-Z, --tlsPort=<value>

Port number for TLS/SSL connections. The user must enter a value with this option.

-y, --tlsYesNew

Force acceptance of new TLS server certificates.

-Y, --tlsYesAll

Force acceptance of all TLS server certificates.

Examples

Example 1

To log in to the server *Eng* as *tester20* you would use the following command:

```
zos login Eng --user=tester20
```

Example 2

To check the current login, you would use the following command:

```
zos login --check
```

logout

This command logs the user out of the current Orchestration grid.

Syntax

```
zos logout
```

Options

-C, --cols=<value>

Specifies the terminal width for formatting. The user must enter a value with this option.

-h, --help

Displays a help message for this operation.

-d, --debug

Turns on debug output during the logout operation.

-v, --verbose

Turns on verbose output during the logout operation.

-z, --tls=<value>

Force TLS/SSL encryption [on|off]. The user must enter a value with this option.

-Z, --tlsPort=<value>

Port number for TLS/SSL connections. The user must enter a value with this option.

-y, --tlsYesNew

Force acceptance of new TLS server certificates.

-Y, --tlsYesAll

Force acceptance of all TLS server certificates.

mkdir

This command makes a new directory in the datagrid.

Syntax

```
zos mkdir path [--parents]
```

Arguments

path

The path on the datagrid.

Options

-p, --parents

Creates parent directories if needed.

-C, --cols=<value>

Specifies the terminal width for formatting. The user must enter a value with this option.

-h, --help

Displays a help message for this operation.

-d, --debug

Turns on debug output during this operation.

-v, --verbose

Turns on verbose output during this operation.

-z, --tls=<value>

Force TLS/SSL encryption [on|off]. The user must enter a value with this option.

-Z, --tlsPort=<value>

Port number for TLS/SSL connections. The user must enter a value with this option.

-y, --tlsYesNew

Force acceptance of new TLS server certificates.

-Y, --tlsYesAll

Force acceptance of all TLS server certificates.

Examples

Example 1

To create a new directory name *images* under the DataGrid root, you would use the following command:

```
zos mkdir grid:///images
```

Example 2

To create a new directory in the job deployment area for the job *myjob*, use the following command:

```
zos mkdir grid:///!myjob/newdir
```

move

This command moves files and directories in the datagrid.

Syntax

```
zos move source dest
```

Arguments

source

The datagrid source from which the files or directories are to be moved.

dest

The datagrid destination where the files or directories are to be moved.

Options

-C, --cols=<value>

Specifies the terminal width for formatting. The user must enter a value with this option.

-h, --help

Displays a help message for this operation.

-d, --debug

Turns on debug output during this operation.

-v, --verbose

Turns on verbose output during this operation.

-z, --tls=<value>

Force TLS/SSL encryption [on|off]. The user must enter a value with this option.

-Z, --tlsPort=<value>

Port number for TLS/SSL connections. The user must enter a value with this option.

-y, --tlsYesNew

Force acceptance of new TLS server certificates.

-Y, --tlsYesAll

Force acceptance of all TLS server certificates.

Examples

Example 1

To move a file from one datagrid directory to another, use the following command:

```
zos move grid:///linux/myFile grid:///windows/myFile
```

password

This command lets the user change his or her password on the current Orchestration Server.

Syntax

```
zos password [--passwd=] [--newpasswd=]
```

Options

-p, --passwd=<value>

Current password for authentication. The user must enter a value with this option.

-n, --newpasswd=<value>

New password for authentication. The user must enter a value with this option.

-C, --cols=<value>

Specifies the terminal width for formatting. The user must enter a value with this option.

-h, --help

Displays a help message for this operation.

-d, --debug

Turns on debug output during the logout operation.

-v, --verbose

Turns on verbose output during the logout operation.

-z, --tls=<value>

Force TLS/SSL encryption [on|off]. The user must enter a value with this option.

-Z, --tlsPort=<value>

Port number for TLS/SSL connections. The user must enter a value with this option.

-y, --tlsYesNew

Force acceptance of new TLS server certificates.

-Y, --tlsYesAll

Force acceptance of all TLS server certificates.

pause

This command directs the running job to pause.

Syntax

```
zos pause [jobid]
```

Arguments

jobid

The user must enter the jobid that he or she is interested in. The default for this argument is the last job run.

Options

-C, --cols=<value>

Specifies the terminal width for formatting. The user must enter a value with this option.

-h, --help

Displays a help message for this operation.

-d, --debug

Turns on debug output during this operation.

-v, --verbose

Turns on verbose output during this operation.

-z, --tls=<value>

Force TLS/SSL encryption [on|off]. The user must enter a value with this option.

-Z, --tlsPort=<value>

Port number for TLS/SSL connections. The user must enter a value with this option.

-y, --tlsYesNew

Force acceptance of new TLS server certificates.

-Y, --tlsYesAll

Force acceptance of all TLS server certificates.

Examples

Example 1

To pause the currently running job for the logged-in user, you would use the following command:

```
zos pause
```

Example 2

To pause the named running job instance, use the following command:

```
zos pause griduser.myjob.1234
```

priority

This command changes the priority of a running job.

Syntax

```
zos priority [jobid] [--priority=]
```

Arguments

jobid

The user must enter the jobid that he or she is interested in. The default for this argument is the last job run.

Options

-r, --priority=<value>

Specifies the priority level for job. The user must enter a value with this option.

-C, --cols=<value>

Specifies the terminal width for formatting. The user must enter a value with this option.

-h, --help

Displays a help message for this operation.

-d, --debug

Turns on debug output during this operation.

-v, --verbose

Turns on verbose output during this operation.

-z, --tls=<value>

Force TLS/SSL encryption [on|off]. The user must enter a value with this option.

-Z, --tlsPort=<value>

Port number for TLS/SSL connections. The user must enter a value with this option.

-y, --tlsYesNew

Force acceptance of new TLS server certificates.

-Y, --tlsYesAll

Force acceptance of all TLS server certificates.

Examples

Example 1

To change the current job's priority to `lowest`, the user would enter the following command:

```
zos priority --priority=lowest
```

Example 2

To reset priority back to the user default (*paul*), the user would enter the following command:

```
zos priority paul.quickie.23
```

resume

This command directs the running job to resume.

Syntax

```
zos resume [jobid]
```

Arguments

jobid

The user must enter the jobid that he or she is interested in. The default for this argument is the last job run.

Options

-C, --cols=<value>

Specifies the terminal width for formatting. The user must enter a value with this option.

-h, --help

Displays a help message for this operation.

-d, --debug

Turns on debug output during this operation.

-v, --verbose

Turns on verbose output during this operation.

-z, --tls=<value>

Force TLS/SSL encryption [on|off]. The user must enter a value with this option.

-Z, --tlsPort=<value>

Port number for TLS/SSL connections. The user must enter a value with this option.

-y, --tlsYesNew

Force acceptance of new TLS server certificates.

-Y, --tlsYesAll

Force acceptance of all TLS server certificates.

run

This command runs a job.

Syntax

```
zos run jobname [--priority=] [--as=] [--watch] [--env=] [--start=] [--listen] [--
detail] [--policy=] [--policyfile] [--trace]
```

Arguments

jobname

The user must enter the name of the job he or she wants to run.

Options

-r, --priority=<value>

Specifies the priority level for a job. The user must enter a value with this option.

-a, --as=<value>

Specifies the name of the job instance. The user must enter a value with this option.

-W, --watch

Lets user watch job log messages in real time.

-V, --env=[<value>]

Lets user use the local environment for running a job or to use an optional list. The user must enter a value with this option.

-s, --start=<value>

Specifies the start time for a job. The user must enter a value with this option.

-l, --listen

Lets user listen for job messages until the job is completed.

-e, --detail

Lets user display detailed message information.

-P, --policy=<value>

Lets the user specify the policy to be associated with this job. The user must enter a value with this option. The job must be deployed.

-F, --policyfile=<value>

Lets the user specify the policy file to be associated with this job. The user must enter a value with this option.

-t, --trace

Lets the user turn on event tracing for this job.

-C, --cols=<value>

Specifies the terminal width for formatting. The user must enter a value with this option.

-h, --help

Displays a help message for this operation.

-d, --debug

Turns on debug output during this operation.

-v, --verbose

Turns on verbose output during this operation.

-z, --tls=<value>

Force TLS/SSL encryption [on|off]. The user must enter a value with this option.

-Z, --tlsPort=<value>

Port number for TLS/SSL connections. The user must enter a value with this option.

-y, --tlsYesNew

Force acceptance of new TLS server certificates.

-Y, --tlsYesAll

Force acceptance of all TLS server certificates.

Examples

Example 1

To run the job called *unittest* with the *highest* priority and another parameter, the user would use the following command:

```
zos run unittest --priority=highest testno=test57
```

Example 2

To run the job called *unittest* passing the PATH environment variable only, the user would use the following command:

```
zos run unittest --env=PATH
```

Example 3

To run the job called *unittest* passing a local policy file, the user would use the following command:

```
zos run unittest --policyfile=/tmp/mypolicy.policy
```

Example 4

To run the job called *unittest* at *5:01 p.m.*, the user would use the following command:

```
zos run unittest --start="5:01 PM"
```

NOTE: The date and time values should be specified as shown in the example—in the format `MM/DD/YY hour:minute AM/PM` and enclosed in quotes—`"11/21/11 5:01 PM"`.

---

search

This command performs a grid object search.

Syntax

```
zos search [--type=] [--detail] [--facts] [--constraint=] [--order=] [--name]
```

Options

-t, --type=<value>

Specify the grid object type (user, resource, job, jobinstance, vmhost, grid. *XXX*group). The user must enter a value with this option.

-e, --detail

Displays detailed information for each grid object rather than just a name.

-f, --facts

Displays facts for each grid object.

-c, --constraint=<value>

Specifies a constraint string to restrict search. The command defaults to "" to match everything. The user must enter a value with this option.

-o, --order=<value>

Specifies returned order (*fact*/a|d, *fact*/a|d,..).

-n, --name

Displays object name instead of object ID.

-C, --cols=<value>

Specifies the terminal width for formatting. The user must enter a value with this option.

-h, --help

Displays a help message for this operation.

-d, --debug

Turns on debug output during this operation.

-v, --verbose

Turns on verbose output during this operation.

-z, --tls=<value>

Force TLS/SSL encryption [on|off]. The user must enter a value with this option.

-Z, --tlsPort=<value>

Port number for TLS/SSL connections. The user must enter a value with this option.

-y, --tlsYesNew

Force acceptance of new TLS server certificates.

-Y, --tlsYesAll

Force acceptance of all TLS server certificates.

Examples

Example 1

To generate a detailed listing of all users in group *eng*, the user would enter the following command:

```
zos search -t user -c '<contains fact="user.groups" value="eng"/>' -o
"user.location/d"
```

status

This command displays the status of a running or previously run job.

Syntax

```
zos status [jobid] [--detail] [--errors] [--recurse] [--joblets]
```

Arguments

jobid

The user must enter the job ID that he or she is interested in. The default for this argument is the last job run.

Options

-e, --detail

Displays joblet information for the job.

-E, --errors

Displays job and node error details.

-r, --recurse

Shows the full job hierarchy. This option can be used with or without the `--detail` option.

-J, --joblets

Retrieves individual joblet details. This option can be used only with the `--detail` option.

-C, --cols=<value>

Specifies the terminal width for formatting. The user must enter a value with this option.

-h, --help

Displays a help message for this operation.

-d, --debug

Turns on debug output during this operation.

-v, --verbose

Turns on verbose output during this operation.

-z, --tls=<value>

Force TLS/SSL encryption [on|off]. The user must enter a value with this option.

-Z, --tlsPort=<value>

Port number for TLS/SSL connections. The user must enter a value with this option.

-y, --tlsYesNew

Force acceptance of new TLS server certificates.

-Y, --tlsYesAll

Force acceptance of all TLS server certificates.

Examples

Example 1

To display a detailed status for the last job run by the login user, the user would enter the following command:

```
zos status -e
```

tail

This command displays the end of a datagrid file.

Syntax

```
zos tail [--retry] [--bytes=] [--follow] [--by-name] [--follow-retry] [--lines=] [--max-unchanged=] [--jobid=] [--sleep-interval=] [--refresh=] [--count=] [--timeout=] [--quiet] path
```

Arguments

path

The path on the datagrid.

Options

-r, --retry

Continues trying to open the file if it is not yet accessible.

-c, --bytes=<value>

Generate file output of the specified number of bytes only. The user must enter a value with this option.

-f, --follow

Follow the file, watching for newly appended data.

-N, --by-name

If following a file, periodically reopen the file.

-F, --follow-retry

The same as `--follow --by-name --retry`

-n, --lines=<value>

Generate file output of the specified number of lines only. The user must enter a value with this option.

-U, --max-unchanged=<value>

The number of size checks between reopening a followed file (used with `--by-name` only). The user must enter a value with this option.

-J, --jobid=<value>

If following a file, follow until the specified job terminates. The user must enter a value with this option.

-s, --sleep-interval=<value>

The frequency of file size checks in seconds when following a file. The user must enter a value with this option.

-R, refresh=<value>

Server connection refresh rate. The user must enter a value with this option.

-X, --count=<value>

Maximum number of lines/bytes to display.

-t, --timeout=<value>

The maximum number of seconds to follow a file.

-q, --quiet

Suppress all warning and informational messages.

-C, --cols=<value>

Specifies the terminal width for formatting. The user must enter a value with this option.

-h, --help

Displays a help message for this operation.

-d, --debug

Turns on debug output during this operation.

-v, --verbose

Turns on verbose output during this operation.

-z, --tls=<value>

Force TLS/SSL encryption [on|off]. The user must enter a value with this option.

-Z, --tlsPort=<value>

Port number for TLS/SSL connections. The user must enter a value with this option.

-y, --tlsYesNew

Force acceptance of new TLS server certificates.

-Y, --tlsYesAll

Force acceptance of all TLS server certificates.

Examples

Example 1

To display the last 5 lines of `grid:///!benchmark.jdl`, use the following command:

```
zos tail -n 5 grid:///!benchmark/benchmark.jdl
```

Example 2

To display the last 5 lines of `grid:///!benchmark.jdl` in an alternative way, use the following command:

```
zos tail -5 grid:///!benchmark/benchmark.jdl
```

Example 3

To follow a file in (approximately) real time, use the following command:

```
zos tail -f grid:///^myjob/job.log
```

Example 4

To continually watch the job log of the latest job, use the following command:

```
zos tail --follow --by-name --retry grid:///!/lastjob/job.log
```

Example 5

To watch the job log of job *myname.myjob.123* while it runs, use the following command:

```
zos tail -F -J myname.myjob.123 grid:///^myname.myjob.123/job.log
```

18 Orchestration Server zosadmin Command Line Tool

The `zosadmin` command line tool is used by administrators of the Cloud Manager Orchestration Server to log in to the server, add or remove server components such as jobs and policies, to report on the status of nodes, users, and the audit database, and to perform other administrative functions.

This section includes information about the following:

- Section 18.1, “List of zosadmin Commands,” on page 243
- Section 18.2, “Getting Started with the zosadmin Command,” on page 244
- Section 18.3, “Details, Usage, and Syntax Examples of zosadmin Commands,” on page 246

18.1 List of zosadmin Commands

The following table includes a list of zosadmin commands in order of common usage, and a description for each command.

Table 18-1 *Available zosadmin Commands and Their Descriptions*

Command	Description
auditclean	Clean the audit database by removing old data
auditcount	Count the number of jobs in the audit database
auditreport	Generate an audit report
cancelalljobs	Cancel all running jobs
create	Create a new server instance
deploy	Deploy a new component onto a server
disconnect	Disconnect and/or revoke user or node sessions
dump	Dump contents of the namespace (advanced diagnostics)
facts	Retrieve descriptions of all facts
get	Retrieve an attribute
help	Displays help for any of the commands in this list
init	Restore a server configuration to initial state
invoke	Invoke an MBean method (advanced diagnostics)
list	List the running servers
login	Log in to the server
logout	Log out of the server

Command	Description
nodes	Retrieve the list of active or inactive nodes
password	Change admin password
redeploy	Redeploy a component on a server
rotatelogs	Back up and rotate the log files
sessions	Display session information
set	Set an attribute
start	Start a local server
status	Collect the status of a server
stop	Shut down the server
undeploy	Undeploy a component from the server
upgrade	Upgrade old server snapshot to current version
users	Retrieve the list of active or inactive users
verify	Verify a component for syntax

18.2 Getting Started with the zosadmin Command

All zosadmin commands begin with zosadmin on the command line. The general format for a zosadmin command is `zosadmin` followed by the command name, followed by command line parameters, if needed:

`zosadmin [`*`standard_options`*`] `*`command`*` [`*`command_options_and_arguments`*`]`

Before you use the zosadmin CLI, ensure that your path is correctly pointing to the Orchestration `/bin` directory.

This section includes the following information:

- Section 18.2.1, “Logging In,” on page 244
- Section 18.2.2, “Checking Login Status,” on page 245
- Section 18.2.3, “Logging Out,” on page 245

18.2.1 Logging In

Login is required to operate on a running server. The commands `start`, `list`, `init`, and `create` do not require a login. Use the following syntax to log in:

```
>zosadmin login -user=username Orhestration_Server_name
Please enter current password for 'username': *******
Logged into Orchestration_grid_name> on Orchestration_Server_name
```

The login should now be complete.

Login information is stored in the `/home` directory, so further `zosadmin` commands use the saved login information. To operate on a different Orchestration Server, run `zosadmin login` to log in to the new Orchestration Server.

18.2.2 Checking Login Status

Enter the following command and parameter to retrieve the status of the current login:

```
>zosadmin login -c
Currently logged into testgrid on server 'tszen5'
```

18.2.3 Logging Out

Enter the following command to log out of the Orchestration Server:

```
>zosadmin logout
Logged out from testgrid
```

18.3 Details, Usage, and Syntax Examples of zosadmin Commands

This section includes a detailed list of the `zosadmin` commands you can use. It also includes examples for using these commands and shows the syntax of typical commands.

NOTE: Items shown in brackets [ ] are optional. Items shown in *italics* are contextual examples.

The section is organized alphabetically by the command names:

- “auditclean” on page 247
- “auditcount” on page 248
- “auditreport” on page 249
- “cancelalljobs” on page 251
- “create” on page 252
- “deploy” on page 255
- “disconnect” on page 256
- “dump” on page 257
- “facts” on page 258
- “get” on page 259
- “init” on page 260
- “invoke” on page 261
- “list” on page 262
- “login” on page 263
- “logout” on page 265
- “nodes” on page 266
- “password” on page 267
- “redeploy” on page 268
- “rotatelogs” on page 269
- “sessions” on page 270
- “set” on page 272
- “start” on page 273
- “status” on page 275
- “stop” on page 276
- “undeploy” on page 277
- “upgrade” on page 278
- “users” on page 280
- “verify” on page 281

auditclean

This command cleans the audit database by removing old data.

Syntax

Enter the command using the syntax shown below. Available options are shown in brackets.

```
zosadmin auditclean --dayskept= [--grid=] [--table=]
```

Options

When you add an option, use the syntax in the command line as shown above (where options are shown in their long form). Short and long options for this command are listed below.

-d, --dayskept <value>

Specify the number of days of history kept in the database (0 removes all).

-g, --grid <value>

Specify a different grid (by ID). Use '*' to match all. Include the single quotes only in a shell environment.

-t, --table <value>

Specify the table (workflow, actions). Default is workflow.

-v, --verbose

Turn on verbose output during this operation.

-V, --debug

Turn on debug output during this operation.

-h, --help

Display a help message for this operation.

Examples

Example 1

To clean up all records older than 7 days for the logged in grid, use the following command:

```
zosadmin auditclean --dayskept=7
```

Example 2

To clean up all records older than 7 days from all servers stored in the audit database, use the following command:

```
zosadmin auditclean --dayskept=7 --gridId=*
```

auditcount

This command retrieves the count of the number of jobs or actions in the audit database.

Syntax

Enter the command using the syntax shown below. Available options are shown in brackets.

```
zosadmin auditcount [--root] [--grid=] [--table=]
```

Options

When you add an option, use the syntax in the command line as shown above (where options are shown in their long form). Short and long options for this command are listed below.

-R, --root

Display the count for root jobs only (default counts all jobs).

-g, --grid <value>

Specify a different grid (by ID). Use '*' to match all. Include the single quotes only in a shell environment.

-t, --table <value>

Specify the table (workflow, actions). The default is "workflow."

-v, --verbose

Turns on verbose output during this operation.

-V, --debug

Turns on debug output during this operation.

-h, --help

Displays a help message for this operation.

Examples

Example 1

To get the count of job records for the logged in grid, use the following command:

```
zosadmin auditcount
```

auditreport

This command displays an audit report generated from the audit database.

Syntax

Enter the command using the syntax shown below. Available options are shown in brackets.

```
zosadmin auditreport [--username=] [--from=] [--to=] [--limit=] [--grid=] [--
childjobs] [--table=]
```

Options

When you add an option, use the syntax in the command line as shown above (where options are shown in their long form). Short and long options for this command are listed below.

-m, --username <value>

Specify the username of the user who ran the job (default is all users).

-F, --from <value>

Select jobs submitted on or after this date. For example, `Fri, 15 Oct 2010` or `Fri, 15 Oct 2010 13:30:00`. The default day is the current day.

-T, --to <value>

Select jobs submitted on or before this date. For example, `Fri, 15 Oct 2010` or `Fri, 15 Oct 2010 13:30:00`. The default day is the current time.

-L, --limit <value>

Limit the number of jobs reported (the default is 500).

-g, --grid <value>

Specify a different grid (by ID). Use '*' to match all. Include the single quotes only in a shell environment.

-I, --childjobs

Include child jobs (the default is root jobs only).

-t, --table <value>

Specify the table (workflow, actions). The default is "workflow."

-V, --debug

Turn on debug output during this operation.

-v, --verbose

Turn on verbose output during this operation.

-h, --help

Display a help message for this operation.

Examples

Example 1

To display a report of job records for the logged in grid, use the following command:

```
zosadmin auditreport
```

Example 2

To display a report of action status records for the logged in grid, use the following command:

```
zosadmin auditreport --table=actions
```

cancelalljobs

This command cancels all running jobs.

Syntax

Enter the command using the syntax shown below. Available options are shown in brackets.

```
zosadmin cancelalljobs
```

Options

When you add an option, use the syntax in the command line as shown above (where options are shown in their long form). Short and long options for this command are listed below.

-v, --verbose

Turns on verbose output during this operation.

-V, --debug

Turns on debug output during this operation.

-h, --help

Displays a help message for this operation.

create

This command creates a new server instance.

Syntax

Enter the command using the syntax shown below. Available options are shown in brackets.

```
zosadmin create [--newdir=] [--upgrade=] [--profile=] [--id=] [--grid=] [--
clusterhost=] [--user=] [--passwd=] [--rmiport=] [--rmihost=] [--httpport=] [--
adminport=] [--agentport=] [--auditurl=] [--audituser=] [--auditpasswd=] [--
auditdriver=] [--cert=] [--key=] [<hostname>]
```

Arguments

hostname

Specify the hostname to use for this server. The default is the system hostname.

Options

When you add an option, use the syntax in the command line as shown above (where options are shown in their long form). Short and long options for this command are listed below.

-N, --newdir <value>

The location where the new Orchestration Server should be installed.

-G, --upgrade

Switch the upgrade behavior for snapshot handling.

Specifically, this command removes everything from the server instance directory except the `logs/`, `dataGrid/`, and `snapshot/` sub-directories. It then invokes `zosadmin upgrade` to apply XSL transforms on the `/snapshot` directory. This last action might be a problem if you want more control over where your snapshot is stored: the `--upgrade` option assumes that the snapshot exists in the `/var/opt/novell/zenworks/zos/server/<snapshot_default_location>`.

If you want more control or flexibility over where you store the snapshot, use the `zosadmin upgrade` command and specify the location of the `/snapshot` directory.

When you use either this command or the `zosadmin upgrade` command, you need to follow it with the `zosadmin start --upgrade` command, optionally specifying the location of the upgraded `/snapshot` directory to complete the upgrade.

NOTE: We recommend that the `zosadmin create --upgrade` command be used only by advanced users. For most upgrade scenarios, you should use the `config` or `guiconfig` scripts for upgrading. For more information, see "Upgrading Orchestration Server Components" in the *Cloud Manager Installation and Upgrade Guide*.

-Q, --agentUpgrade

Quietly accept automatic upgrade of existing agents.

If this parameter appears on the command line, it causes the server to automatically accept agent upgrade requests from old agents. This flag is set by the config and guiconfig scripts if the person doing the upgrade/install checks the **Automatically Upgrade Agents** check box while using the configuration tools.

-e, --profile <value>

Create a new instance based on named profile (the default is `server`).

-i, --id <value>

Select by kernel ID.

-g, --grid <value>

Select by grid name.

-C, --clusterhost <value>

Hostname of virtual cluster IP resource (for High Availability configuration).

-u, --user <value>

Username used in accessing secure remote sites.

-p, --passwd <value>

Password used in accessing secure remote sites.

-P, --rmiport <value>

Select by RMI port.

-H, --rmihost <value>

Select or specify the RMI host (can be different from *server host*).

-W, --httpport <value>

Specify the http port used for the Server Portal (the default is 80).

-M, --adminport <value>

Specify the http port used for Administrative Information Portal (default 8001).

-A, --agentport <value>

Specify the communication port for agent connections (default 8100).

-T, --agenttlsport <value>

Specify the communication port for agent TLS connections (default 8101).

-X, --auditurl <value>

Specify the connection URL to a PostgreSQL database (for example, `jdbc:postgresql://localhost/zos_db`).

-Y, --audituser <value>

Specify the audit database user.

-Z, --auditpasswd <value>

Specify the audit database password.

-D, --auditdriver <value>

Specify the audit database driver.

-k, --cert <value>

Specify the `.pem` file containing the new server's TLS certificate.

-K, --key <value>

Specify the `.pem` file containing the new server's TLS key.

-v, --verbose

Turns on verbose output during this operation.

-V, --debug

Turns on debug output during this operation.

-h, --help

Displays a help message for this operation.

Examples

Example 1

To create a new server instance in the install directory (default), use the following command:

```
zosadmin create
```

Example 2

To create a new server instance using host name *myhost.mydomain.com*, use the following command:

```
zosadmin create myhost.mydomain.com
```

Example 3

To create a new server instance in the `/zos/myserver/` directory, use the following command:

```
zosadmin create --newdir=/zos/myserver
```

Example 4

To create a new upgraded server instance in an existing directory (`/zos/myserver/`), use the following command:

```
zosadmin create --newdir=/zos/myserver --upgrade
```

IMPORTANT: This command removes the previous server instance data from the directory.

Example 5

To create an upgraded server and reuse saved tls certificates, use the following command:

```
zosadmin create --newdir=/zos/myserver --upgrade --cert=/tls/cert.pem
     --key=/tls/private.pem
```

deploy

This command deploys a new component onto a server.

Syntax

Enter the command using the syntax shown below. Available options are shown in brackets.

```
zosadmin deploy [--port=] file|dir
```

Arguments

file|dir

Specify the deployable file or a directory containing deployable files.

Options

When you add an option, use the syntax in the command line as shown above (where options are shown in their long form). Short and long options for this command are listed below.

-v, --verbose

Turn on verbose output during this operation.

-V, --debug

Turn on debug output during this operation.

-h, --help

Display a help message for this operation.

Examples

Example 1

To deploy the service file `myservice.sar`, use the following command:

```
zosadmin deploy myservice.sar
```

Example 2

To deploy the job file `quickie.job`, use the following command:

```
zosadmin deploy quickie.job
```

Example 3

To deploy a job file in the `/test/myjob.job` directory, use the following command:

```
zosadmin deploy /test/myjob.job
```

disconnect

This command allows forceful disconnect or revocation of user or node sessions.

Syntax

Enter the command using the syntax shown below. Available options are shown in brackets.

```
zosadmin disconnect [--id=] [--node=] [--user=] [--all] [--allUsers] [--allNodes]
[--revoke]
```

Options

When you add an option, use the syntax in the command line as shown above (where options are shown in their long form). Short and long options for this command are listed below.

-i, --id

Identity of a specific user session.

-n, --node <value>

Disconnect a node by name.

-u, --user <value>

Disconnect all of a user’s sessions by name.

-a, --all

Disconnect all user and node sessions.

-U, --allUsers

Disconnect all user sessions.

-N, --allNodes

Disconnect all node sessions.

-r, --revoke

Revoke the session or sessions in addition to disconnecting.

-V, --debug

Turn on debug output during this operation.

-v, --verbose

Turn on verbose output during this operation.

-h, --help

Display a help message for this operation.

dump

This command dumps the contents of namespace.

Syntax

Enter the command using the syntax shown below. Available options are shown in brackets.

```
zosadmin dump     [--dir=] --jndi=
```

Options

When you add an option, use the syntax in the command line as shown above (where options are shown in their long form). Short and long options for this command are listed below.

-d, --dir <value>

Specify the working directory for a Orchestration Server installation.

-j, --jndi <value>

Specify the JNDI path to be displayed.

-v, --verbose

Turn on verbose output during this operation.

-V, --debug

Turn on debug output during this operation.

-h, --help

Display a help message for this operation.

Examples

Example 1

To dump the configuration for the naming facility, use the following command:

```
zosadmin dump --jndi=/facility/naming/config
```

facts

This command retrieves descriptions for all Orchestration facts.

Syntax

Enter the command using the syntax shown below. Available options are shown in brackets.

```
zosadmin facts
```

Options

When you add an option, use the syntax in the command line as shown above (where options are shown in their long form). Short and long options for this command are listed below.

-v, --verbose

Turn on verbose output during this operation.

-V, --debug

Turn on debug output during this operation.

-h, --help

Display a help message for this operation.

Examples

Example 1

To get a listing of all facts and their descriptions, use the following command:

```
zosadmin facts
```

get

This command retrieves an attribute.

Syntax

Enter the command using the syntax shown below. Available options are shown in brackets.

```
zosadmin get --mbean= --attr=
```

Options

When you add an option, use the syntax in the command line as shown above (where options are shown in their long form). Short and long options for this command are listed below.

-m, --mbean <value>

Specifies the name of the MBean to view.

-a, --attr <value>

Specifies the attribute name of the MBean to view.

-v, --verbose

Turn on verbose output during this operation.

-V, --debug

Turn on debug output during this operation.

-h, --help

Display a help message for this operation.

Examples

Example 1

To get the node autoregistration setting, use the following command:

```
zosadmin  get --mbean=local:facility=nodeManager --attr=Autoregister
```

init

This command restores a server configuration to its initial state.

Syntax

Enter the command using the syntax shown below. Available options are shown in brackets.

```
zosadmin init [--dir=]
```

Options

When you add an option, use the syntax in the command line as shown above (where options are shown in their long form). Short and long options for this command are listed below.

-d, --dir <value>

The working directory for a Orchestration installation.

-v, --verbose

Turn on verbose output during this operation.

-V, --debug

Turn on debug output during this operation.

-h, --help

Display a help message for this operation.

invoke

This command invokes an MBean method.

Syntax

Enter the command using the syntax shown below. Available options are shown in brackets.

```
zosadmin invoke --mbean= --method= [--sig=]
```

Options

When you add an option, use the syntax in the command line as shown above (where options are shown in their long form). Short and long options for this command are listed below.

-m, --mbean <value>

Specify the name of the MBean to view.

-i, --method <value>

Specify the method name of the MBean to invoke.

-s, --sig <value>

Specify the method argument signature(s) of the MBean to invoke..

-v, --verbose

Turn on verbose output during this operation.

-V, --debug

Turn on debug output during this operation.

-h, --help

Display a help message for this operation.

list

This command lists all of the running servers.

Syntax

Enter the command using the syntax shown below. Available options are shown in brackets.

```
zosadmin list [--grid=] [--id=] [--server=] [--rmiport] [--rmipath] [--rmihost=] [-
--rmiurl=] [--timeout=]
```

Options

When you add an option, use the syntax in the command line as shown above (where options are shown in their long form). Short and long options for this command are listed below.

-g, --grid <value>

Select by grid name.

-i, --id <value>

Select by kernel ID.

-s, --server <value>

Select by host name.

-P, --rmiport <value>

Select by RMI port.

-J, --rmipath <value>

Select or specify the RMI bind path.

-H, --rmihost <value>

Select or specify the RMI host (can be different from <server host>).

-r, --rmiurl <value>

Select by full RMI URL.

-t, --timeout <value>

Sets the query timeout (in seconds).

-V, --debug

Turn on debug output during this operation.

-v, --verbose

Turn on verbose output during this operation.

-h, --help

Display a help message for this operation.

login

This command logs into the Orchestration Server.

Syntax

Enter the command using the syntax shown below. Available options are shown in brackets.

```
zosadmin login [--grid=] [--id=] [--rmiport=] [--rmipath=] [--rmihost=] [--rmiurl=]
[--timeout=] [--user=] [--passwd=] [--check=] [<server host>]
```

Arguments

<server host>

Enter the server host name that you are logging into. The command polls for a server if none is supplied.

Options

When you add an option, use the syntax in the command line as shown above (where options are shown in their long form). Short and long options for this command are listed below.

-g, --grid <value>

Select by grid name.

-i, --id <value>

Select by kernel ID.

-P, --rmiport <value>

Select by RMI port.

-J, --rmipath <value>

Select or specify the RMI bind path.

-H, --rmihost <value>

Select/specify RMI host (can be different from <server host>).

-r, --rmiurl <value>

Select by full RMI URL.

-t, --timeout <value>

Sets the query timeout (in seconds).

-u, --user <value>

Username used in accessing secure server management.

-p, --passwd <value>

Password used in accessing secure server management.

-c, --check

Check and report existing login. Ignores other options.

-V, --debug

Turns on debug output during the login operation.

-v, --verbose

Turns on verbose output during the login operation.

-h, --help

Display a help message for this operation.

Examples

Example 1

To login to server *Eng*, use the following command:

```
zosadmin login Eng
```

Example 2

To check the current login, use the following command:

```
zosadmin login --check
```

logout

This command logs out of the Orchestration Server.

Syntax

Enter the command using the syntax shown below. Available options are shown in brackets.

```
zosadmin logout
```

Options

When you add an option, use the syntax in the command line as shown above (where options are shown in their long form). Short and long options for this command are listed below.

-v, --verbose

Specifies the terminal width for formatting. The user must enter a value with this option.

-V, --debug

Turns on debug output during the logout operation.

-h, --help

Displays a help message for this operation.

nodes

This command retrieves a list of all, online, offline, or mismatched nodes.

Syntax

Enter the command using the syntax shown below. Available options are shown in brackets.

```
zosadmin nodes [--offline] [--mismatch] [--name] [--all]
```

Options

When you add an option, use the syntax in the command line as shown above (where options are shown in their long form). Short and long options for this command are listed below.

-o, --offline

Retrieve a list of offline users and nodes.

-m, --mismatch

Retrieve a list of nodes requiring update because of a version mismatch.

-a, --all

Retrieve a list of all users and nodes.

-n, --name

Display the Grid object names rather than their object IDs.

-v, --verbose

Turn on verbose output during this operation.

-V, --debug

Turn on debug output during this operation.

-h, --help

Display a help message for this operation.

password

This command lets the administrator or developer change his or her password on the Orchestration Server.

Syntax

Enter the command using the syntax shown below. Available options are shown in brackets.

```
zosadmin password [--passwd=] [--newpasswd=]
```

Options

When you add an option, use the syntax in the command line as shown above (where options are shown in their long form). Short and long options for this command are listed below.

-p, --passwd <value>

Password used for secure server management.

-l, --newpasswd <value>

New password to use for secure remote access.

-v, --verbose

Turns on verbose output during this operation.

-h, --help

Displays a help message for this operation.

redeploy

This command redeploys a component onto a server.

Syntax

Enter the command using the syntax shown below. Available options are shown in brackets.

```
zosadmin redeploy [--session=] file|dir
```

Arguments

file|dir

Specify the deployable file or a directory containing deployable files.

Options

When you add an option, use the syntax in the command line as shown above (where options are shown in their long form). Short and long options for this command are listed below.

-i, --session <value>

Specifies the deployment session ID.

-v, --verbose

Turn on verbose output during this operation.

-V, --debug

Turn on debug output during this operation.

-h, --help

Display a help message for this operation.

Examples

Example 1

To redeploy the service file `myservice.sar`, use the following command:

```
zosadmin redeploy myservice.sar
```

Example 2

To redeploy the job file `quickie.job`, use the following command:

```
zosadmin redeploy quickie.job
```

Example 3

To redeploy the job files in directory `/test/myjob.job`, use the following command:

```
zosadmin redeploy /test/myjob.job
```

rotatelogs

This command backs up and rotates the log files.

Syntax

Enter the command using the syntax shown below. Available options are shown in brackets.

```
zosadmin rotatelogs
```

Options

When you add an option, use the syntax in the command line as shown above (where options are shown in their long form). Short and long options for this command are listed below.

-v, --verbose

Turn on verbose output during this operation.

-V, --debug

Turn on debug output during this operation.

-h, --help

Display a help message for this operation.

sessions

This command displays session information.

Syntax

Enter the command using the syntax shown below. Available options are shown in brackets.

```
zosadmin sessions [id=] [--node=] [--user=] [--all] [--allNodes] [--allUsers] [--
subSessions] [--inactive] [--full] [--connaddr] [--noDns]
```

Options

When you add an option, use the syntax in the command line as shown above (where options are shown in their long form). Short and long options for this command are listed below.

-i, --id <value>

Display this session for this session ID.

-n, --node <value>

Display sessions for the specified node.

-u, --user <value>

Display sessions for the specified user.

-a, --all

Display sessions for all clients.

-N, --allNodes

Display sessions for all nodes.

-U, --allUsers

Display sessions for all users.

-s, --subSessions

Include subsession entries.

-I, --inactive

Include inactive entries.

-f, --full

Include full session information.

-c, --connaddr

Include local and remote addresses.

-d, --noDns

Do not look up host names in DNS.

-V, --debug

Turn on debug output during this operation.

-v, --verbose

Turn on verbose output during this operation.

-h, --help

Display a help message for this operation.

set

This command sets an attribute on an MBean.

Syntax

Enter the command using the syntax shown below. Available options are shown in brackets.

```
zosadmin set --mbean= --attr= --value= --type=
```

Options

When you add an option, use the syntax in the command line as shown above (where options are shown in their long form). Short and long options for this command are listed below.

-m, --mbean <value>

Specify the name of the MBean to view.

-a, --attr <value>

Specify the attribute name of the MBean to view.

-o, --value <value>

Specify the attribute value of the MBean to set in string form.

-t, --type <value>

Specify the attribute value type to convert string form into (String/Boolean/Integer).

-v, --verbose

Turn on verbose output during this operation.

-V, --debug

Turn on debug output during this operation.

-h, --help

Display a help message for this operation.

Examples

Example 1

To change the node autoregistration to true, use the following command:

```
zosadmin set --mbean=local:facility=authenticationManager,properties=auth --
attr=AutoregisterUnauthAgent --value=true --type=Boolean
```

start

This command starts a local server.

Syntax

Enter the command using the syntax shown below. Available options are shown in brackets.

```
zosadmin start [--dir=] [--jvmargs=] [--javaargs=] [--timeout=] [--upgrade] [--
snapshot=]
```

Options

When you add an option, use the syntax in the command line as shown above (where options are shown in their long form). Short and long options for this command are listed below.

-d, --dir <value>

Specify the working directory for an Orchestration components installation.

-b, --jvmargs <value>

Pass the value as an argument to an invoked JVM process.

-B, --javaargs <value>

Pass the value as an argument to an invoked Java program.

-t, --timeout <value>

Specify the query timeout (in seconds).

-G, --upgrade

Switch upgrade behavior for snapshot handling.

-S, --snapshot=

Upgrade a server using the snapshot directory of an existing server.

-D, --heapDump=

Switch off heap dump generation on an out of memory condition

-V, --debug

Turn on debug output during this operation.

-v, --verbose

Turn on verbose output during this operation.

-h, --help

Display a help message for this operation.

Examples

Example 1

To start a server instance using the install directory (default), use the following command:

```
zosadmin start
```

Example 2

To start a server instance using directory `/zos/server/`, use the following command:

```
zosadmin start --dir=/zos/server
```

Example 3

To start a server instance and upgrade from an existing server snapshot, use the following command:

```
zosadmin start --upgrade --snapshot=/oldzos/server/snapshot
```

NOTE: Before you use this option, you must be using the “zosadmin upgrade” command to manually upgrade the server configuration snapshot.

Example 4

To start a server instance setting the JVM maximum memory to be 4 Gb:

```
zosadmin start --jvmargs=-Xmx4g
```

status

This command displays the status of a server (including information on managed components)

Syntax

Enter the command using the syntax shown below. Available options are shown in brackets.

```
zosadmin status [--mbeans] [--sessions] [--facilities]
```

Options

When you add an option, use the syntax in the command line as shown above (where options are shown in their long form). Short and long options for this command are listed below.

-m, --mbeans

Retrieve the list of all MBeans.

-s, --sessions

Display deployment sessions ID along with component name.

-f, --facilities

Retrieve information on all facilities.

-v, --verbose

Turn on verbose output during this operation.

-V, --debug

Turn on debug output during this operation.

-h, --help

Display a help message for this operation.

stop

This command shuts down the Orchestration Server.

Syntax

Enter the command using the syntax shown below. Available options are shown in brackets.

```
zosadmin stop [--force] [--snap]
```

Options

When you add an option, use the syntax in the command line as shown above (where options are shown in their long form). Short and long options for this command are listed below.

-f, --force

Do not prompt for server shutdown confirmation.

-s, --snap

Create a snapshot of server state for use in a later upgrade.

-v, --verbose

Turn on verbose output during this operation.

-V, --debug

Turn on debug output during this operation.

-h, --help

Display a help message for this operation.

Examples

Example 1

To stop a server instance without prompting for confirmation, use the following command:

```
zosadmin stop --force
```

Example 2

To stop a server instance and create a snapshot, use the following command:

```
zosadmin stop --snap
```

undeploy

This command undeploys a component from the server.

Syntax

Enter the command using the syntax shown below. Available options are shown in brackets.

`zosadmin undeploy [--session=] component|dir`

Arguments

component|dir

Specify the deployed file or a directory containing the deployed files.

Options

When you add an option, use the syntax in the command line as shown above (where options are shown in their long form). Short and long options for this command are listed below.

-i, --session <value>

Specify the deployment session ID.

-v, --verbose

Turn on verbose output during this operation.

-V, --debug

Turn on debug output during this operation.

-h, --help

Display a help message for this operation.

Examples

Example 1

To undeploy the service `myservice.sar`, use the following command:

```
zosadmin undeploy myservice.sar
```

Example 2

To undeploy the job file `quickie.job` from the server, use the following command:

```
zosadmin undeploy quickie.job
```

Example 3

To undeploy the job files in directory `/test/myjob.job`, use the following command:

```
zosadmin undeploy /test/myjob.job
```

upgrade

This command upgrades an older server snapshot to a current version.

Syntax

Enter the command using the syntax shown below. Available options are shown in brackets.

`zosadmin upgrade snapshotdir`

Arguments

snapshotdir

Specify the directory containing a server configuration snapshot.

Options

When you add an option, use the syntax in the command line as shown above (where options are shown in their long form). Short and long options for this command are listed below.

-v, --verbose

Turn on verbose output during this operation.

-V, --debug

Turn on debug output during this operation.

-h, --help

Display a help message for this operation.

The `zosadmin upgrade` command does not actually upgrade a server instance. Use the `zosadmin create` and `zosadmin start` commands for manually upgrading an instance using a transformed configuration snapshot.

The `zosadmin create` command with the `--upgrade` option removes everything from the server instance directory except the `logs/`, `dataGrid/`, and `snapshot/` sub-directories. It then invokes `zosadmin upgrade` to apply XSL transforms on the `/snapshot` directory. This last action might be a problem if you want more control over where your snapshot is stored: the `--upgrade` option assumes that the snapshot exists in the `/var/opt/novell/zenworks/zos/server/<snapshot_default_location>`.

If you want more control or flexibility over where you store the snapshot, use the `zosadmin upgrade` command and specify the location of the `/snapshot` directory.

NOTE: We recommend that this command be used only by advanced users. For most upgrade scenarios, you should use the `config` or `guiconfig` scripts for upgrading. For more information, see "Upgrading Orchestration Server Components" in the *Cloud Manager Installation and Upgrade Guide*.

When you use either this command or the `zosadmin create --upgrade` command, you need to follow it with the `zosadmin start --upgrade` command, optionally specifying the location of the upgraded `/snapshot` directory to complete the upgrade.

Examples

Example 1

To upgrade the server snapshot in the current directory, use the following command:

```
zosadmin upgrade ./snapshot
```

users

This command retrieves a list of all active or inactive users.

Syntax

Enter the command using the syntax shown below. Available options are shown in brackets.

```
zosadmin users [--offline] [--all]
```

Options

When you add an option, use the syntax in the command line as shown above (where options are shown in their long form). Short and long options for this command are listed below.

-o, --offline

Retrieve a list of offline users/nodes.

-a, --all

Retrieve a list of all users/nodes.

-v, --verbose

Turn on verbose output during this operation.

-V, --debug

Turn on debug output during this operation.

-h, --help

Display a help message for this operation.

verify

This command verifies a component for syntax.

Syntax

Enter the command using the syntax shown below. Available options are shown in brackets.

```
zosadmin verify file
```

Arguments

file

Specify the deployable file to verify.

Options

When you add an option, use the syntax in the command line as shown above (where options are shown in their long form). Short and long options for this command are listed below.

-v, --verbose

Turn on verbose output during this operation.

-V, --debug

Turn on debug output during this operation.

-h, --help

Display a help message for this operation.

Examples

Example 1

To verify the policy file `myjob.policy`, use the following command:

```
zosadmin verify myjob.policy
```

Example 2

To verify the schedule file `mysched.sched` and dump reconstituted XML, use the following command:

```
zosadmin verify mysched.sched --debug
```

III Orchestration Provisioning Adapters

This section includes additional reference information about the NetIQ Cloud Manager Orchestration provisioning adapters used to discover VM hosts and VMs in various supported hypervisor environments:

- Chapter 19, “The VMware vSphere Provisioning Adapter,” on page 285

19 The VMware vSphere Provisioning Adapter

Because you probably want to get your Cloud Manager Orchestration environment configured for use in your vCenter hypervisor environment at installation time, detailed configuration information you need for the vSphere provisioning adapter is included in the "Configuring the vSphere Provisioning Adapter" in the *Cloud Manager Installation and Upgrade Guide*.

- Section 19.1, "Provisioning Actions Supported by the vSphere Provisioning Adapter," on page 285
- Section 19.2, "vSphere Server Changes Synchronized to Cloud Manager by vSphereUpdater," on page 286
- Section 19.3, "Apply Config and Save Config Actions on vSphere VMs," on page 288
- Section 19.4, "Discovering or Dynamically Creating vSphere Port Groups," on page 291
- Section 19.5, "Enabling VNC Access to vSphere 5 VM Guest Consoles," on page 293
- Section 19.6, "Known Configuration Limitations for vSphere VMs," on page 296

19.1 Provisioning Actions Supported by the vSphere Provisioning Adapter

The following table lists the VM provisioning actions supported by the Orchestration Console for the vsphere provisioning adapter job.

NOTE: For a list of changes this provisioning adapter detects in the vSphere environment and synchronizes to Cloud Manager, see "vSphere Server Changes Synchronized to Cloud Manager by vSphereUpdater" on page 286.

Table 19-1 *Provisioning Actions Supported by the vSphere Provisioning Adapter on Guest Operating Systems*

Orchestration Server Managed VM Action	vSphere Provisioning Adapter on Guest OS
Provision	Yes
Clone	Yes
Shutdown	Yes
Destroy	Yes
Suspend	Yes
Pause	No
Resume	No
Create Template	Yes
Move Disk Image[1]	Yes

Orchestration Server Managed VM Action	vSphere Provisioning Adapter on Guest OS
Hot Migrate[2]	Yes
Warm Migrate[3]	Yes
Checkpoint	Yes
Restore	Yes
Install Orchestration Agent	Yes
Make Standalone	Yes
Check Status	Yes
Save Config	Yes
Cancel Action	Yes
Check Host Assignment	Yes
Launch Remote Desktop	Yes

NOTE: Host operating systems are dependent on the VMware vSphere support matrix

[1] A “move” is the relocation of VM disk images between two storage devices when the VM is in a not running state (including VMs that are suspended with a checkpoint file). This action does not require shared storage; the move is between separate repositories.

[2] A “hot migrate” (also called a “live migrate”) is the migration of a running VM to another host and starting it there with minimal resulting downtime (measured in milliseconds). This action requires shared storage.

[3] A “warm migrate” is the migration of a suspended VM to another host and starting it there with brief resulting downtime (measured in seconds). This action requires shared storage.

19.2 vSphere Server Changes Synchronized to Cloud Manager by vSphereUpdater

Typically, you perform all cloud management from the Cloud Manager interface. However, if you do make a change by using the VMware vCenter tool, it is vSphereUpdater that brings many of those changes into your Cloud Manager system.

The Cloud Manager vSphere Provisioning Adapter detects changes that occur in the vSphere Server environment, and synchronizes many of them to the Cloud Manager Orchestration Server, which then communicates them to the Cloud Manager Application Server, and ultimately, to the Web Console.

The Cloud Manager Orchestration update infrastructure consists of two main components:

- A vSphereUpdate Client component, which is executed by the Orchestration Agent.
- The vSphereUpdate monitor job, which starts the vSphereUpdate Client component and ensures that it runs when necessary.

For information about configuring the vSphereUpdate components, see “Configuring vSphereUpdate Client and Monitoring Jobs” in the *Cloud Manager Installation and Upgrade Guide*.

The following sections detail the changes that are synchronized:

- Section 19.2.1, “VM Changes,” on page 287
- Section 19.2.2, “Repository Changes,” on page 287
- Section 19.2.3, “Network Changes,” on page 287
- Section 19.2.4, “Resource Pool Changes,” on page 287

19.2.1 VM Changes

When aspects of unique virtual machines are changed in the vSphere environment, Cloud Manager synchronizes these changes to its interface:

- A vDisk is created, modified, or deleted on the VM.
- A vNIC is created, modified, or deleted on theVM.
- A VM host and its datastore is migrated.
- A VM is moved to a new resource pool.
- A VM is renamed.
- The RAM or CPU is changed on a VM.

19.2.2 Repository Changes

When aspects of vSphere repositories are changed, Cloud Manager synchronizes these changes to its interface:

- A repository is created, renamed, or deleted.
- A repository on a cluster is created, renamed, or deleted.

19.2.3 Network Changes

When aspects of vSphere networks are changed, Cloud Manager synchronizes these changes to its interface:

- A network is created, modified, or deleted.

19.2.4 Resource Pool Changes

When aspects of vSphere Resource Pool are changed, Cloud Manager synchronizes these changes to its interface:

- A VM is moved into or out of the resource pool.
- The configuration of a resource pool is changed.

19.3 Apply Config and Save Config Actions on vSphere VMs

This section provides information to help you anticipate the behavior of vsphere VMs of various supported platforms when the Apply Config action or the Save Config action is applied.

- Section 19.3.1, “Expected Results of the Apply Config Action on Windows Server VMs,” on page 288
- Section 19.3.2, “Expected Results of the Save Config Action on Windows Server VMs,” on page 289
- Section 19.3.3, “Expected Results of the Apply Config Action on Linux Server VMs,” on page 290
- Section 19.3.4, “Expected Results of the Save Config Action on Linux Server VMs,” on page 291

19.3.1 Expected Results of the Apply Config Action on Windows Server VMs

The table below lists the supported Windows Server VM functionality when the **Apply Config** action is applied while the VM is running. A “Yes” in the table indicates support for the function.

Table 19-2 *Supported Disk, NIC, and Memory Setting Functionality on vSphere Windows Server VMs after the* **Apply Config** *Action Is Applied*

Function	Windows Server 2003 R2 VM	Windows Server 2008 R2 VM
Attach IDE disk	Yes Effective on reboot.	Yes Effective on reboot.
Detach IDE disk	Yes Effective on reboot.	Yes Effective on reboot.
Modify IDE disk	Yes Effective on reboot.	Yes Effective on reboot.
Insert DVD / CD-ROM media into IDE disk	Does not work.	Does not work.
Eject DVD / CD-ROM media from IDE disk	Does not work.	Does not work.
Attach USB disk	Not supported by the hypervisor.	Not supported by the hypervisor.
Detach USB disk	Not supported by the hypervisor.	Not supported by the hypervisor.
Attach SCSI disk.	Yes Effective on reboot.	Yes Effective on reboot.
Detach SCSI disk	Yes Effective on reboot.	Yes Effective on reboot.
Detach Fibre Channel disk	Yes Effective on reboot.	Yes Effective on reboot.
Attach Fibre Channel disk	Yes Effective on reboot.	Yes Effective on reboot.

Function	Windows Server 2003 R2 VM	Windows Server 2008 R2 VM
Attach Virtual NIC	Yes Effective on reboot.	The vNIC is displayed in the Orchestration Console, but does not operate correctly.
Detach Virtual NIC	Yes Effective on reboot.	Yes Effective on reboot.
Set required VM memory	Not supported by the hypervisor.	Not supported by the hypervisor.
Set host CPU number	Not supported by the hypervisor.	Not supported by the hypervisor.
Set VNC port	Yes	Yes

19.3.2 Expected Results of the Save Config Action on Windows Server VMs

The table below lists the supported Windows Server VM functionality when the **Save Config** action is applied while Windows VMs are shut down. A “Yes” in the table indicates support for the function.

Table 19-3 *Supported Disk, NIC, and Memory Setting Functionality on vSphere Windows Server VMs after the* **Save Config** *Action Is Applied*

Function	Windows Server VM
Attach IDE disk	Yes
Detach IDE disk	Yes
Modify IDE disk	Yes
Insert DVD / CD-ROM media into IDE disk	Does not work.
Eject DVD / CD-ROM media from IDE disk	Does not work.
Attach USB disk	Not supported by the hypervisor.
Detach USB disk	Not supported by the hypervisor.
Attach SCSI disk.	Yes
Detach SCSI disk	Yes
Detach Fibre Channel disk	Yes
Attach Fibre Channel disk	Yes
Attach Virtual NIC	Yes
Detach Virtual NIC	Yes
Set required VM memory	Yes
Set host CPU number	Yes
Set VNC port	Yes

19.3.3 Expected Results of the Apply Config Action on Linux Server VMs

Table 19-4 lists the supported VM functionality when the **Apply Config** action is applied on running Linux Server VMs. A “Yes” in the table indicates support for the function. Required drivers for each platform are noted below the table.

Table 19-4 *Supported Disk, NIC, and Memory Setting Functionality on vSphere Linux VMs after the* **Apply Config** *Action is Applied*

Function	Linux Server VM
Attach IDE disk	Yes Effective on reboot.
Detach IDE disk	Yes Effective on reboot.
Modify IDE disk	Yes Effective on reboot.
Insert DVD / CD-ROM media into IDE disk	Does not work.
Eject DVD / CD-ROM media from IDE disk	Does not work.
Attach USB disk	Not supported by the hypervisor.
Detach USB disk	Not supported by the hypervisor.
Attach SCSI disk.	Yes Effective on reboot.
Detach SCSI disk	Yes Effective on reboot.
Detach Fibre Channel disk	Yes Effective on reboot.
Attach Fibre Channel disk	Yes Effective on reboot.
Attach Virtual NIC	Yes Effective on reboot.
Detach Virtual NIC	Yes Effective on reboot.
Set required VM memory	Not supported by the hypervisor.
Set host CPU number	Not supported by the hypervisor.
Set VNC port	Yes

19.3.4 Expected Results of the Save Config Action on Linux Server VMs

The table below lists the supported VM functionality when the **Save Config** action is applied while Linux VMs are shut down. A “Yes” in the table indicates support for the function. Required drivers for each platform are noted below the table.

Table 19-5 *Supported Disk, NIC, and Memory Setting Functionality on vSphere Linux VMs after the* **Save Config** *Action Is Applied*

Function	Linux Server VM
Attach IDE disk	Yes
Detach IDE disk	Yes
Modify IDE disk	Yes
Insert DVD / CD-ROM media into IDE disk	Does not work.
Eject DVD / CD-ROM media from IDE disk	Does not work.
Attach USB disk	Not supported by the hypervisor.
Detach USB disk	Not supported by the hypervisor.
Attach SCSI disk.	Yes
Detach SCSI disk	Yes
Detach Fibre Channel disk	Yes
Attach Fibre Channel disk	Yes
Attach Virtual NIC	Yes
Detach Virtual NIC	Yes
Set required VM memory	Yes
Set host CPU number	Yes
Set VNC port	Yes

19.4 Discovering or Dynamically Creating vSphere Port Groups

If you manage a vSphere hypervisor environment, you can configure virtual switches (vSwitches) and distributed virtual switches (dvSwitches) with virtual ports to allow virtual machines on the same ESX Server host to communicate on the network. Your VMware environment also requires port groups, which, according to the VMware Infrastructure Library (http://pubs.vmware.com/vi3/serverconfig/wwhelp/wwhimpl/common/html/wwhelp.htm?context=serverconfig&file=sc_networking.5.3.html) “define how a connection is made through the vSwitch to the (physical) network construct.”

You can use the Orchestration Server to monitor the discovered VMs in your vSphere environment (provided that you have configured the vSwitches or dvSwitches on those VMs), dynamically deleting unused port groups from the virtual switches to free up resources on the vSwitch or dvSwitch so that it can be used elsewhere in the virtual network. Then, if you later decide to provision a VM that

requires the port group, the Orchestration Server can be configured to automatically re-create the required port group. This allows VMs to be moved or migrated from one host to another without reconfiguring the networking information on either the VM or the host.

This section includes the following information:

- Section 19.4.1, “How the Orchestration Server Models a vSphere Port Group and Exposes vSwitch Settings,” on page 292
- Section 19.4.2, “Dynamically Creating Port Groups,” on page 292

19.4.1 How the Orchestration Server Models a vSphere Port Group and Exposes vSwitch Settings

When the Orchestration Server discovers a vSphere VM, it maps port group configuration settings (including its associated vSwitch or dvSwitch configuration settings) to a virtual bridge (vBridge) grid object in the Orchestration Server. As the virtual networking settings (including the VLAN ID) of these resources are determined during VM and VM host discovery, the Orchestration Server also automatically creates a Network group object to group similarly associated vBridge objects. All necessary information required for re-creating a port group is preserved on the Network object and the vBridge object.

19.4.2 Dynamically Creating Port Groups

You can instruct the Orchestration Server to dynamically create a port group in the vSphere environment if a VM that you are provisioning requires it or you can delete and existing port group when the last VM utilizing that port group is shut down. Port groups are dynamically created on the host/dvSwitch only when a VM requiring a port group is provisioned. A port group is deleted only when there are no active VMs utilizing it and if it is not configured as a management network interface. In order for dynamic creation or deletion to function properly, you need to configure the dvSwitches/vSwitches in the vSphere environment.

The `vsphere` provisioning adapter utilizes information it gathers during discovery and stores on the Network objects to determine how to correctly create a required port group and to attach it to the correct dvSwitches or vSwitches. You can modify this information on the discovered Network objects, or you can add it to Network objects that you customize in the Orchestration Console.

This section includes the following information:

- “Network Configuration Facts” on page 292
- “Enabling Dynamic Creation of Port Groups” on page 293

Network Configuration Facts

The Orchestration Server uses the following facts when it creates a port group:

Table 19-6 *Network Configuration Facts for a Port Group*

Fact Name	Type	Description
`group.dvswitch`	Boolean	Specifies a Boolean fact that specifies if the port group is to be attached to a dvSwitch or to a standard host vSwitch. The value is set to true for a dvSwitch or false if set for a vSwitch.

Fact Name	Type	Description
group.switch	String	Specifies the name of the dvSwitch or vSwitch that the port group is to be attached to. This fact can be a regular expression to match against different dvSwitches or vSwitches on different hosts.
group.vsphere.portgroupsettings	String	Specifies an XML string that describes the settings on the port group (for example, Security, Traffic Shaping, or Teaming policies). Use the vSphere tools to configure these settings, so that the Orchestration Server can properly populate the fact upon discovery.
group.vlanid	Integer	Specifies the vLAN ID that is associated with the port group. A value of zero (0) is equivalent to no vLAN.

The Orchestration Server also uses the `group.vbridge.matchname`.fact. This is a regular expression used during discovery of Networks to correlate port groups that are functionally identical (and should therefore be grouped in the same Network object in the Orchestration Server) but are not named the same. So, if Host A contains a port group named HostA-InternetConnection and Host B contains a port group named HostB-InternetConnection, modifying the `group.vBridge.matchname` fact to Host?-InternetConnection" results in creating a vBridge object for each host under a single Network object.

Enabling Dynamic Creation of Port Groups

1. Assign the vsphere_ignoreNetwork.policy to all vSphere hosts where you want to enable port group creation. This lets the host be considered as a provisioning target, even if the port groups do not currently exist on the host's dvSwitches or vSwitches.
2. Associate the vsphere_ignoreNetwork.policy with all VMs where you want to enable port group creation. This enables the computed facts that are necessary to provide proper ordering of provisioning plans.
3. In the vsphere.policy, update the `vsphere.vcenters` fact by setting the **auto_portgroups_creation** setting to true.

   Modifying the **auto_portgroups_deletion** and **auto_portgroups_disconnection** settings is not necessary to enable dynamic creation of port groups, but is necessary for automatically reclaiming the port group resources when they are not in use.

19.5 Enabling VNC Access to vSphere 5 VM Guest Consoles

You can allow access through the firewall in a vSphere 5 environment, provided that the ESXi 5 Server and the vCenter Server are properly configured according to VMware documentation. The vSphere Client must also be installed and configured properly according to VMware documentation.

You can use one of two methods to open firewall ports for VNC access:

- Section 19.5.1, "Enabling VNC Access By Opening Multiple Firewall Ports," on page 294
- Section 19.5.2, "Enabling VNC Access by Creating a Special Configuration File," on page 294

19.5.1 Enabling VNC Access By Opening Multiple Firewall Ports

Use the following steps to enable the 59*xx* firewall ports for VNC access to vSphere 5 VM guest consoles:

1. In your vSphere environment, log in to the vSphere Client, then select **Home** > **Inventory** > **Hosts and Clusters**.
2. In the Hosts/Clusters tree view, select the ESXi host name that represents the server you want to open for VNC access.
3. Select the **Configuration** tab, locate and open the **Software** list box, then select **Security Profile**.
4. In the Firewall section, select the **Properties** link to display the Firewall Properties dialog.
5. In the dialog, scroll to and select **GDB Server**, then click **OK**.

   Your ESXi server now allows VNC access to Guest VM consoles through its firewall.

   NOTE: If you are using vSphere 4.*x* or earlier, select **VNC Server** in the list box for this step.

   ---

6. Repeat these steps for each ESXi host system.

The **GDB Server** setting covers the needed 59*xx* port range in its own port range, so when you enable firewall access for GDB Server, VNC services also become open. Extra open ports do not present a serious security problem, because a user rarely runs manual services that listen on those ports.

If you are concerned about these extra open ports, you can use the method for opening firewall ports explained in Section 19.5.2, “Enabling VNC Access by Creating a Special Configuration File,” on page 294.

19.5.2 Enabling VNC Access by Creating a Special Configuration File

If leaving the extra ports open is a security concern, you can manually add the VNC Server entry to the ESXi 5.1 firewall configuration and persist that entry across reboots of the server.

IMPORTANT: The preferred method to enable VNC Access to an ESXi 5.1 server is to use an existing, preconfigured GDB Server firewall entry, as described in Section 19.5.1, “Enabling VNC Access By Opening Multiple Firewall Ports,” on page 294.

If you use the method described in this section to enable VNC access, NetIQ strongly recommends that you have competent experience with command line Linux/Unix system administration. It is possible to make mistakes while performing these steps that might render your ESXi Server unbootable.

---

1. In your vSphere environment, log in to the vSphere Client, then select **Home** > **Inventory** > **Hosts and Clusters**.
2. In the Hosts/Clusters tree view, select the ESXi host name that represents the server you want to open for VNC access.
3. Select the **Configuration** tab, locate and open the **Software** list box, then select **Security Profile**.
4. In the Firewall section, select the **Properties** link to display the Firewall Properties dialog.
5. In the dialog, scroll to and select **SSH Server**, then click **OK**.
6. From a Linux console, ssh to the IP address of your ESXi host. Log in as `root` using that host's root password.

7 Using a Linux editor (such as vi), use one of the following methods to execute commands in ESXi during the boot process. This process differs depending on the version of ESXi being used.

- **ESXi 5.1 and higher:** Add the following shell script lines at the end above the line `exit 0` of the `local.sh` file located at the `/etc/rc.local.d/` directory. When you are done, save the `local.sh` file. Commands added after the `exit 0` line are not executed.
- **ESXi 5.0:** Add the following shell script lines to the end of the `/etc/rc.local` file. When you are done, save the `/etc/rc.local` file.

```
cat <<EOF > /etc/vmware/firewall/vncServer.xml
<ConfigRoot>
  <service>
    <id>vncServer</id>
    <rule id='0000'>
      <direction>inbound</direction>
      <protocol>tcp</protocol>
      <porttype>dst</porttype>
      <port>
        <begin>5900</begin>
        <end>5999</end>
      </port>
    </rule>
    <enabled>true</enabled>
    <required>false</required>
  </service>
</ConfigRoot>
EOF
esxcli network firewall refresh
```

IMPORTANT: Enter the code *exactly* as shown in the sample above. Use spaces to indicate indents in the code, *do not use* tab characters.

---

8 While still logged in, run the following command:

```
/sbin/auto-backup.sh
```

9 Log out from the SSH session.

10 From either the ESXi host's console or from the VMWare Client, reboot the ESXi host.

You should now see **VNC Server** as an available service in the Firewall Properties pane. The service should be enabled.

This process creates the `/etc/vmware/firewall/vncServer.xml` config file with the necessary settings to open the firewall ports.

Simply creating and editing this file does not work when the ESXi Server is rebooted because the root file system in ESXi 5 is a volatile RAM disk that is loaded from a master copy on each boot. Any changes made to this RAM disk are lost upon reboot.

A workaround to this rule relies on the fact that the ESXi Server uses the `auto-backup.sh` script to persist a select set of files every 10 minutes (or when changes are made by with the VMware Client or the VI-SDK facilities) from this file system to the master persistent copy. The `/etc/rc.local` file is one of these select files, so adding the shell script to the end of the file can add the needed firewall entry each time the ESXi server boots.

19.6 Known Configuration Limitations for vSphere VMs

When a vSphere VM is deleted by the vSphere client, it can leave behind a VM template in a “defined” state that can be subsequently discovered and registered as a resource in the Orchestration Server.

If you attempt to remove this VM template in the Orchestration Console, the Orchestration Server attempts to remove the VM’s files from the file system, but because the VM has no files to delete, the provisioning job fails with the following error:

```
Error: Delete : Failed delete of Managed Entity :
```

For identified vSphere VM issues, see Section 20.7, “Troubleshooting vSphere VM Provisioning Actions,” on page 308.

IV Orchestration Server and Provisioning Adapter Troubleshooting

Use the information in this section to troubleshoot Orchestration components of Cloud Manager.

- Chapter 20, “Troubleshooting Cloud Manager Orchestration Issues,” on page 299
- Chapter 21, “Increasing the Kernel ARP Threshold Value on the Orchestration Server,” on page 313
- Chapter 22, “Orchestration Server Logs,” on page 315

20 Troubleshooting Cloud Manager Orchestration Issues

This section provides information that identifies some ongoing known issues in Cloud Manager Orchestration components and the methods that you can use to address those issues.

- Section 20.1, “Troubleshooting Installation Issues,” on page 299
- Section 20.2, “Troubleshooting Orchestration Server Upgrade Issues,” on page 300
- Section 20.3, “Troubleshooting Orchestration Server Operation Issues,” on page 300
- Section 20.4, “Troubleshooting Orchestration Console Issues,” on page 302
- Section 20.5, “Troubleshooting Orchestration Agent Issues,” on page 303
- Section 20.6, “Troubleshooting General VM Management Issues,” on page 304
- Section 20.7, “Troubleshooting vSphere VM Provisioning Actions,” on page 308
- Section 20.8, “Multicast Troubleshooting,” on page 312

20.1 Troubleshooting Installation Issues

The following sections provide solution to the problems you might encounter while performing the installation or configuration of the product:

- “Configuration Programs Do Not Include a Way to Edit the Agent Configuration” on page 299

Configuration Programs Do Not Include a Way to Edit the Agent Configuration

Source: Orchestration Installation and Configuration Programs.

Explanation: Although the scenario is not supported in a production environment, it is common in demonstration or evaluation situations to install the Orchestration Agent and the Orchestration Server on the same machine.

An error might occur if you install the agent after the initial server installation or if you attempt to use the configuration programs (`config`, `guiconfig`) to change the agent configuration after it is installed. Because of port checking routine in the configuration program, the error alerts you that port 8100 is already in use.

Action: To correct the problem for a demonstration setup, stop the Orchestration Server, configure the agent with one of the configuration programs, then restart the server.

20.2 Troubleshooting Orchestration Server Upgrade Issues

The following sections provide solutions to the problems you might encounter when you upgrade the Orchestration Server:

- “The vSphereUpdate User Name Might Be Lost After An Upgrade” on page 300

The vSphereUpdate User Name Might Be Lost After An Upgrade

Source: The Orchestration Server.

Explanation: When you perform an upgrade from earlier versions of Cloud Manager, the vSphereUpdate user is deleted in the vSphereUpdater policy.

Action: Rather than using this policy for the user name, we recommend that you associate the user to the scheduled vSphereUpdater job to preserve the user name during upgrade.

20.3 Troubleshooting Orchestration Server Operation Issues

The following sections provide solutions to the problems you might encounter while using the Orchestration Server:

- “Orchestration Server Might Appear to Be Deadlocked When Provisioning Large Numbers of Jobs with Subjobs” on page 300
- “Orchestration Server Might Hang if the System Clock Is Changed Abruptly” on page 301
- “Authentication to an Active Directory Server Might Fail” on page 301
- “The Orchestration Server Must Have Sufficient RAM” on page 301
- “Calling terminate() from within a Job Class Allows the JDL Thread Execution to Continue” on page 301
- “Java programs That Use the JDL Exec Class Might Hang” on page 302

Orchestration Server Might Appear to Be Deadlocked When Provisioning Large Numbers of Jobs with Subjobs

Source: Cloud Manager Orchestration Server

Explanation: In some deployments where a large number of running jobs spawn subjobs, the running jobs might appear to stop, leaving jobs in the queue.

Possible Cause: This occurs because of job limits set in the Orchestration Server to avoid overload or “runaway” conditions.

Action: If this deadlock occurs, you can slowly adjust the job limits to tune them according to your deployment. For more information, see “Job Limits Panel” on page 25.

Orchestration Server Might Hang if the System Clock Is Changed Abruptly

Source: Cloud Manager Orchestration Server

Explanation: As with many applications, you should avoid abrupt changes in the system clock on the machine where the Orchestration Server is installed; otherwise, the agent might appear to hang, waiting for the clock to catch up.

This issue is not affected by changes in clock time occurring from daylight saving adjustments.

Action: We recommend that you use proper clock synchronization tools such as a Network Time Protocol (NTP) server in your network to avoid large stepping of the system clock.

Authentication to an Active Directory Server Might Fail

Source: Cloud Manager Orchestration Server

Explanation: A simplified Active Directory Server (ADS) setup might be insufficient because of a customized ADS install (for example, `namingContexts` entries that generate referrals when they are looked up).

Possible Cause: The checking logic in the current AuthLDAP auth provider assumes that if any `namingContext` entry is returned, it has found the domain and it stops searching.

Action: If you encounter this issue, you need to manually configure LDAP as a generic LDAP server, which offers many more configuration options.

The Orchestration Server Must Have Sufficient RAM

Source: Cloud Manager Orchestration Server

Explanation: If the Orchestration Server fails to start after installation and configuration, sufficient RAM might not be installed on your hardware or assigned to the VM you are attempting to use.

Possible Cause: The Orchestration Server requires 3 GB of RAM to function with the preset defaults.

Action: If the server does not start, increase your physical RAM size (or, for a VM, increase the setting for virtual RAM size). Alternatively, you can reduce the JVM heap size, as explained in "Validating and Optimizing the Orchestration Installation" in the *Cloud Manager Installation and Upgrade Guide*.

Calling terminate() from within a Job Class Allows the JDL Thread Execution to Continue

Source: Cloud Manager Orchestration Server

Explanation: Calling `terminate()` from within the Job class does not immediately terminate the JDL thread of that job; instead, it sends a message to the server requesting termination of the job.

Action: This can take time to occur (because subjobs need to be recursively terminated and joblets cancelled), so if the calling JDL thread needs to terminate immediately, immediately follow the invocation of this method with return.

Java programs That Use the JDL Exec Class Might Hang

Source: Cloud Manager Orchestration Server

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

Action: To work around this issue, use the following steps to turn off the enhanced ExecWrapper:

1. In the Explorer tree of the Orchestration Console, select the job that you want to change.
2. In the admin view of the job, select the JDL Editor tab to open the JDL Editor.
3. Paste the following code into the editor:

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

4. Save the changes.

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

For more information about the JDL Exec class, see the Cloud Manager 2 JDL documentation (http://www.novell.com/documentation/cloudmanager2/resources/jdljavadoc_2/com/novell/zos/jdl/Exec.html).

20.4 Troubleshooting Orchestration Console Issues

The following sections provide solution to the problems you might encounter while using the Orchestration Console:

- "The Server Console Displays Incorrect CPU Speed for SLES 11 SP1 Resources" on page 302
- "After Installing the Orchestration Agent on VM, the VM is Not Displayed as a Resource in the Orchestration Console" on page 303

The Server Console Displays Incorrect CPU Speed for SLES 11 SP1 Resources

Source: Orchestration Console

Explanation: The CPU speed displayed in the Orchestration Console (see the `resource.cpu.mhz` and `resource.metrics.cpu_speed` facts) for SLES 11 SP1 resources is incorrect. The invalid display results from powersave settings on the CPU. Until the CPU has been run at full speed, `/proc/cpuinfo` displays this incorrect value for CPU MHz, and the value in the Orchestration Server is also incorrect.

Possible Cause: The issue results from the CPU starting in powersave mode. This slows down the CPU until it is needed, so `/proc/cpuinfo` does not show the maximum potential speed of the CPU. Instead, it shows the maximum speed that the CPU has shown since boot time.

Action: To work around this issue, run the `powersave --performance-speed` command at the server command line.

This command forces the CPU to reach its maximum speed, so you should see the correct value displayed in `/proc/cpuinfo` and the Development Client should also display the correct speed. After you run this command, you can set the powersave mode to a normal state with either of the following commands:

```
powersave --powersave-speed
```

or

```
powersave --dynamic-speed
```

When the powersave mode is set to a normal state, `/proc/cpuinfo` retains the accurate value for the current CPU speed.

TIP: To see the contents of `/proc/cpuinfo`, run the `cat /proc/cpuinfo` command at the bash prompt of your SLES server.

After Installing the Orchestration Agent on VM, the VM is Not Displayed as a Resource in the Orchestration Console

Source: The Orchestration Console

Action: Do the following:

- Ensure that the Orchestration Agent is running on the VM.
- Ensure that no errors have been logged into the `agent.log` file.

  The log file is located in the `<Orchestration_Agent_installation_directory>\novell\zos\agent\node.default` directory on Windows and in the `/opt/novell/zos/agent/node.default` directory on Linux.
- Ensure that the Orchestration Server is registered to the DNS server.

20.5 Troubleshooting Orchestration Agent Issues

The following sections provide solution to the problems you might encounter while using the Orchestration Console:

- "Orchestration Agent Fails to Set the UID on Files Copied from the Datagrid" on page 303

Orchestration Agent Fails to Set the UID on Files Copied from the Datagrid

Source: Orchestration Agent

Explanation: If Network File System (NFS) is used to mount a shared volume across nodes that are running the Orchestration Agent, the agent cannot properly set the UID on files copied from the datagrid to the managed nodes by using the default NFS configuration on most systems.

Action: To address this problem, disable root squashing in NFS so that the agent has the necessary privileges to change the owner of the files it copies.

For example, on a Red Hat Enterprise Linux (RHEL) NFS server or on a SUSE Linux Enterprise Server (SLES) NFS server, the NFS configuration is set in `/etc/exports`. The following configuration is needed to disable root squashing:

```
/auto/home *(rw,sync,no_root_squash)
```

In this example, `/auto/home` is the NFS mounted directory to be shared.

NOTE: The GID is not set for files copied from the datagrid to an NFS mounted volume, whether root squashing is disabled or not. This is a limitation of NFS.

20.6 Troubleshooting General VM Management Issues

The following sections provide solutions to the problems you might encounter while working with general VM management operations:

- “Volume Tools Hang While Scanning a Suspended Device” on page 304
- “RHEL 5 VMs Running the Kudzu Service Do Not Retain Network Interface Changes” on page 305
- “Policies Applied to VM Resources Are Deleted” on page 305
- “VMs Provisioned from a VM Template Are Not Restarted When a VM Host Crashes” on page 307
- “Admin Password on Windows 2003/2008 Workloads Cannot Be Set by Users” on page 307
- “Unable to Provision a VM to Another Cluster Node Due to Reason “VM Networks Are Not Available”” on page 307
- “If Multiple Workloads are Cloned Simultaneously, They are not Load Balanced Across Repositories” on page 308
- “Block Disks Show up as Regular Vdisks in Orchestrate” on page 308

Volume Tools Hang While Scanning a Suspended Device

Source: Scanned device.

Explanation: When a mapped device is in a suspended state, volume tools such as vgscan, lvscan, and pvscan hang. If the vmprep job is run on such a device, it throws an error such as the following to alert you to the condition:

```
vmquery: /var/adm/mount/vmprep.df8fd49401e44b64867f1d83767f62f5:
Failed to
mount vm image "/mnt/nfs_share/vms/rhel4tmpl2/disk0": Mapped
device
/dev/mapper/loop7p2 appears to be suspended. This might cause
scanning for
volume groups (e.g. vgscan) to hang.
WARNING! You may need to manually resume or remove this mapped
device (e.g.
dmsetup remove /dev/mapper/loop7p2)!
```

Action: Because of this behavior, we recommend against using LVM and similar volume tools on a virtual machine managed by Orchestration Services.

RHEL 5 VMs Running the Kudzu Service Do Not Retain Network Interface Changes

Source: Kudzu service.

Explanation: Anytime you modify the hardware configuration (for example, changing the MAC address or adding a network interface card) of a RHEL 5 VM that is running the Kudzu hardware probing library, the VM does not retain the existing network interface configuration.

Possible Cause: When you start a RHEL 5 VM, the Kudzu service recognizes the hardware changes at boot time and moves the existing configuration for that network interface to a backup file. The service then rewrites the network interface configuration to use DHCP instead.

Action: To work around this problem, disable the Kudzu service within the RHEL VM by using the `chkconfig --del kudzu` command.

Policies Applied to VM Resources Are Deleted

Source: VM clones awaiting provisioning.

Explanation: Provisioning code requires that VMs and VM clones be standalone (that is, they are removed from a template dependency and are no longer considered to be “linked clones”).

Possible Cause: VMs in PlateSpin Orchestrate 2.5 and later must be made standalone to receive and retain associated policies.

Action: Apply a conditional policy to the parent template that can be applied to the clones while they are running. Depending upon the facts set on the clone, the inherited VM host constraint can be conditionally applied to the clone.

The following is an example of a conditional policy that you could apply to the VM template to restrict `vmhost` based on resource attributes (group membership, etc.).

```
<policy>
    <constraint type="vmhost">
        <if>
            <contains fact="resource.groups" value="exclude_me"
                      reason="Only apply this vmhost constraint to
resources NOT in exclude_me resource group" >
            </contains>
            <else>
                <if>
                    <defined fact="resource.some_boolean_fact" />
                    <eq fact="some_boolean_fact" value="true" />
                    <then>
                        <contains fact="vmhost.resource.groups"
value="first_vmhost_group"
                                  reason="When a resource is not in
the exclude_me group, when some_ boolean_fact is true,
                                          provision to a vmhost in
the first_vmhost_group"/>
                    </then>
                    <else>
                        <if>
                            <defined
fact="resource.some_other_boolean_fact" />
                            <eq fact="some_other_boolean_fact"
value="true" />
                            <not>
                                <and>
                                    <eq fact="resource.id"
value="never_use_this_resource"
                                     reason="Specifically exclude
this resource from consideration." />
                                    <or>
                                        <eq fact="vmhost.cluster"

factvalue="resource.provision.vmhost.cluster" />
                                        <eq fact="vmhost.cluster"

factvalue="resource.provision.vmhost" />
                                    </or>
                                </and>
                            </not>
                            <then>
                                <contains
fact="vmhost.resource.groups" value="another_vmhost_group"
                                          reason="When a resource is
not in the exclude_me group, when some_ boolean_fact is false,
                                                  and
some_other_boolean_fact is true, (but also not some other things),
                                                 provision to a vmhost
in another_vmhost_group"/>
                            </then>
                        </if>
                    </else>
                </if>
            </else>
        </if>
    </constraint>
</policy>
```

VMs Provisioned from a VM Template Are Not Restarted When a VM Host Crashes

Source: VM host with VMs provisioned from a template.

Explanation: If a VM host crashes, VMs that were provisioned from a template on that host are not restarted on another active VM host. Instead, the Orchestration Server provisions another VM cloned from the original template, on the next available host. The disk files of the original clone are not destroyed (that is, “cleaned up”) after the crash, but the original VM template files are destroyed.

If a Discover Repository action is issued before the cloned VM is deleted from the crashed host, the Orchestration Server creates a new VM object with the `zombie_` string prepended to the VM object name.

Possible Cause: While hosting a provisioned clone, VM host crashed or the Orchestration Agent on that host went offline.

Action: To work around this issue, you can either remove the VM from the file system before the Orchestration Server rediscovers it, or you can issue a Destroy action on the discovered “zombie” VM.

Admin Password on Windows 2003/2008 Workloads Cannot Be Set by Users

Source: Windows 2003/2008 workloads in the Cloud Manager Web Console accessed by users.

Explanation: In order for a user to set the Administrator pass when configuring a Windows 2003/2008 workload, the VM template (from which the workload is created) must not have an Administrator password set.

Action: To leave the Administrator password unset on a VM template, you must turn off the complex password setting in the password policy.

Unable to Provision a VM to Another Cluster Node Due to Reason “VM Networks Are Not Available”

Source: The Orchestration Console

Explanation: A prerequisite for clustering is that every node contained within a cluster should be symmetric. That is, every node in a cluster should have visibility to all the networks and storage provided by the cluster. In this case, because a VM host cluster must be able to place the VM on any node in the cluster, the networks shown as being available to that cluster are the intersection of all the networks available on the VM host nodes that are members of the cluster (see the `vmhost.networks` fact on the cluster object).

Action: Reconfigure each of the cluster nodes to provide the networks required by the VM host cluster and re-run the **Discover Hosts** action.

Alternatively, you can reconfigure the VM to use another network available to all cluster nodes. After you choose a new network(s) configuration for a VM, ensure that you run the **Save Config** action to commit these changes to the VM configuration.

When you reconfigure the networks on a VM, at least one network option, **all**, is available. This option designates that any network can be suitable for VM placement. Choosing this option allows the network constraint to pass, and the provisioning adapter is then responsible for configuring a new network as it sees fit.

If Multiple Workloads are Cloned Simultaneously, They are not Load Balanced Across Repositories

Source: The Orchestration Server

Explanation: When multiple workloads are being cloned at the same time, the cloning process looks at the current state of the storage repositories to determine which repository should be used.

Possible Cause: With multiple asynchronous cloning processes running concurrently, the utilization of the repository does not reflect the state at the completion of other running clone processes. This leads to one repository being identified as the preferred repository until prior cloning jobs have finished running.

Action: Perform a single cloning operation at a time in order to achieve true load balancing, or be aware that multiple cloning operations can result in workload distributions between repositories that are not truly load balanced.

Block Disks Show up as Regular Vdisks in Orchestrate

Source: The Orchestration Client.

Explanation: The first time discovery is run after adding a new block device to a VM, the block device is marked as a regular vdisk in the repository.

Possible Cause: The VM discovery took place before the pdisk was discovered, and Orchestrator has not matched them up.

Action: Run the VM discovery process a second time.

20.7 Troubleshooting vSphere VM Provisioning Actions

The following sections provide solution to the problems you might encounter while performing provisioning actions on VMs managed by the VMware vCenter hypervisor:

- “Unable to perform any Provisioning Adapter Action after the Save Config Action on the vSphere Managed VM” on page 309
- “(503) Service Unavailable Errors Might Occur while Cloning vSphere VMs” on page 309
- “Invalid Datastore Path Error” on page 309
- “Running Provisioning Operations on a Batch of vSphere VMs Results in JDL Event Handler Errors” on page 310
- “Moving a VM Host in vSphere Results in Duplicate Repositories” on page 311
- “Domain Settings are Not Applied Correctly on Some Windows Workloads” on page 312

Unable to perform any Provisioning Adapter Action after the Save Config Action on the vSphere Managed VM

Source: The Orchestration Console.

Possible Cause: The VM UUID value of the vSphere managed VM is not a 128-bit hexadecimal value. Even though the Save Config action is successful and the VM is provisioned, the hypervisor automatically assigns a different UUID value. Subsequently, any provisioning adapter action performed on the VM fails.

Action: Specify a 128-bit hexadecimal value for the VM UUID.

1. In the Orchestration Console, click **Resources** > the vSphere managed VM.

   The **Info/Groups** tab is displayed by default.
2. In the Virtual Machine Configuration panel, set the value of **VM UUID** to a 128-bit hexadecimal value.
3. Right-click the vSphere managed VM, then click **Save Config**.

(503) Service Unavailable Errors Might Occur while Cloning vSphere VMs

Source: The Orchestration Console

Explanation: Running the **Clone** action repeatedly on vSphere VM templates might result in the following error:

```
Clone : (503)Service Unavailable
```

Possible Cause: This error indicates that the server is currently unable to handle the request due to a temporary overloading or maintenance of the server. Testing has shown that this error is most likely to occur when vSphere and the Orchestration Agent are both installed on the same Windows Server 2003 computer.

Action: If you encounter this error, we recommend that you download and apply the appropriate Microsoft hotfix (http://support.microsoft.com/kb/979230) to the vCenter server.

Invalid Datastore Path Error

Source: The Orchestration Server.

Explanation: When attempting to **Save Config** a vSphere VM with an ISO-backed vDisk (for example, a vDisk that specifies a location in the `/vmmimages` folder and does not have its repository fact set), the job fails with a message similar to the following:

```
VMSaveConfig : Invalid datastore path '/vmimages/tools-isoimages/
linux.iso'
```

Action: To work around this issue, associate a policy with the ISO-backed vdisk object that prepends an empty datastore string ([]) to the beginning of the `vdisk.location` fact. For example:

```
<policy>
  <vdisk>
    <fact name="location"
          type="String"
          value="[] /vmimages/tools-isoimages/linux.iso" />
  </vdisk>
</policy>
```

Running Provisioning Operations on a Batch of vSphere VMs Results in JDL Event Handler Errors

Source: The Orchestration Console

Explanation: If you write JDL scripts to automate provisioning actions for a large number of vSphere VMs, you might receive a failure notice similar to the following:

```
Job 'testadmin.r_testvm_resync_batch.15684' terminated because of
failure.
Reason: job exceeded max limit of jdl event handler
Job 'testadmin.r_testvm_resync_batch.15684' terminated because of
failure.
Reason: job exceeded max limit of jdl event handler
```

You also see the following error in `server.log`:

```
08.24 17:32:59: JobManager,NOTICE: job instance
'testadmin.r_testvm_resync_batch.15082' failed
08.24 17:46:25: JobManager,NOTICE: job instance
'testadmin.r_testvm_resync_batch.15684' failed
08.24 17:46:25: Broker,ERROR: Exception in thread "JDL Event
(job_failed_event) jobId (testadmin.r_testvm_resync_batch.15684)"
08.24 17:46:25: Broker,ERROR: ValueError: I/O operation on closed
file
```

Possible Cause: This error indicates that maximum number of JDL threads allowed by the server have been exceeded. Testing has shown that numerous instances of the `provisioner_completed_event` are blocked and waiting for the provisioner job to finish its `job_started_event`.

Action: Rewrite the original script. The original script might look like this:

```
import time
class testvm_resync(Job):
    def job_started_event(self):
        vms_group = getMatrix().getGroup(TYPE_RESOURCE, 'VMs') #
gets the matrix object id for 'VMs' group
        vms = vms_group.getMembers() # gets the group members of
'VMs' group
        for vm in vms:
            id = vm.getFact("resource.id") #gets the resource.id
fact of a vm
            thevmtype = vm.getFact("resource.type") # find the vm
type
          if id.startswith("c-") and thevmtype == 'VM':  # search
criteria
              vmstate = vm.getFact("resource.provision.state") #
find the vm state
               thevm = getMatrix().getGridObject(TYPE_RESOURCE,
id); #gets the vm's id
               thevm.check() # vm life cycle operations
               time.sleep(2*60) #pause time - 2 min
```

The rewritten script might look like this:

```
import time
class testvm_resync(Job):
    def job_started_event(self):
        timer = Timer(self.prov,0)
    def prov(self):
        vms_group = getMatrix().getGroup(TYPE_RESOURCE, 'VMs') #
gets the matrix object id for 'VMs' group
        vms = vms_group.getMembers() # gets the group members of
'VMs' group
        for vm in vms:
            id = vm.getFact("resource.id") #gets the resource.id
fact of a vm
            thevmtype = vm.getFact("resource.type") # find the vm
type
          if id.startswith("c-") and thevmtype == 'VM':  # search
criteria
              vmstate = vm.getFact("resource.provision.state") #
find the vm state
               thevm = getMatrix().getGridObject(TYPE_RESOURCE,
id); #gets the vm's id
               thevm.check() # vm life cycle operations
               time.sleep(2*60) #pause time - 2 min
```

This change lets the `job_started_event` end after transferring the process to another JDL event/method to run on a timer basis. In this example, the timer is set for 10 seconds, but you could set it to zero.

The timer is normally used for callback. For example, the vSphere provisioning adapter uses `Timer` to check every 30 seconds whether a vSphere action is still working or dead.

This not isolated to the check() action. It includes other actions such as provision(), shutdown(), suspend(), checkpoint(), saveConfig() and restart().

Moving a VM Host in vSphere Results in Duplicate Repositories

Source: The Orchestration Server.

Explanation: If you move a VM host in your vSphere environment and then you subsequently perform a discovery in the Orchestration Console, the console displays duplicate repositories for the host that was moved.

Action: After you rediscover VM hosts and repositories in the Orchestration Console, you should delete the old repository grid object from the Explorer tree view in the Orchestration Console. Identify the repository to be deleted by checking the name of the datacenter, which is included in the `repository.datacenter` fact. If the value for this fact is the name of the old datacenter, this is the repository you want to delete.

Domain Settings are Not Applied Correctly on Some Windows Workloads

Source: Microsoft Windows Sysprep.

Explanation: Testing has shown that when you configure the Windows settings for a Windows 2008 R2 workload or a Windows 2012 workload in a new business service request, the information you provide when selecting and naming an Active Directory domain for the workload is not applied.

Possible Cause: The issue is a result of a conflict between Windows Sysprep and VMware tools. For Windows 2008 R2 workloads, you need to apply the hotfix (http://support.microsoft.com/kb/2693187) available from Microsoft.

Action: The issue has not been resolved for Windows 2012 workloads. Microsoft has been informed of the problem.

20.8 Multicast Troubleshooting

An exhaustive tutorial of IP multicasting and troubleshooting is beyond the scope of this document. If you are having problems with multicast discovery, ensure that your operating system is configured to provide IP multicasting support. Most modern versions of Linux* and Windows* provide this support, however it might be disabled. Multicast discovery does not work unless IP multicasting is enabled by the operating system. Routing misconfiguration on the system can also lead to problems with multicast discovery.

21 Increasing the Kernel ARP Threshold Value on the Orchestration Server

Testing has shown that Cloud Manager Orchestration grids that have more than 1210 VMs (each with an installed Orchestration Agent) and 124 Resource objects, the Orchestration Server (installed on supported SUSE Linux Enterprise Server computers) in the grid loses connection to the Audit Database Server or other devices installed on other machines. The failure is manifest following the ping command when the following error is displayed:

```
connect : No buffer space is available
```

To correct this problem, it is necessary to increase the kernel ARP threshold on the Orchestration Server (SL. This section contains the following information to help you perform this increase:

- Section 21.1, “Threshold Definitions,” on page 313
- Section 21.2, “Determining the Current Kernel Threshold Value,” on page 313
- Section 21.3, “Changing the Current Kernel Threshold Value,” on page 314

21.1 Threshold Definitions

The following definitions for kernel ARP levels is referenced from the Linux ARP man page (http://linux.die.net/man/7/arp):

- **gc_thresh1:** The minimum number of entries to keep in the ARP cache. The garbage collector will not run if there are fewer than this number of entries in the cache. Defaults to 128.
- **gc_thresh2:** The soft maximum number of entries to keep in the ARP cache. The garbage collector will allow the number of entries to exceed this for 5 seconds before collection will be performed. Defaults to 512.
- **gc_thresh3:** The hard maximum number of entries to keep in the ARP cache. The garbage collector will always run if there are more than this number of entries in the cache. Defaults to 1024.

21.2 Determining the Current Kernel Threshold Value

To determine the current kernel threshold value on the Orchestration Server, use the following commands to determine the values for each threshold:

- `# cat /proc/sys/net/ipv4/neigh/default/gc_thresh1`

  The command might result in a displayed value like this:

  ```
  # cat /proc/sys/net/ipv4/neigh/default/gc_thresh1
  128
  ```

- `# cat /proc/sys/net/ipv4/neigh/default/gc_thresh2 512`

  The command might result in a displayed value like this:

```
# cat /proc/sys/net/ipv4/neigh/default/gc_thresh2
512
```

- `# cat /proc/sys/net/ipv4/neigh/default/gc_thresh3`

  The command might result in a displayed value like this:

  ```
  # cat /proc/sys/net/ipv4/neigh/default/gc_thresh3
  1024
  ```

You can also use this command:

```
#sysctl -A|grep ipv4|grep default |grep gc_thresh
```

The result is a listing similar to the following:

```
net.ipv4.neigh.default.gc_thresh3 = 1024
net.ipv4.neigh.default.gc_thresh2 = 512
net.ipv4.neigh.default.gc_thresh1 = 128
```

21.3 Changing the Current Kernel Threshold Value

When you know the current threshold values, you can change them using one of two methods:

- Section 21.3.1, “Editing the /etc/sysctl.conf File,” on page 314
- Section 21.3.2, “Making Live Changes to the Threshold Values,” on page 314

21.3.1 Editing the /etc/sysctl.conf File

1. Open `/etc/sysctl.conf` in a text editor.
2. Add the following lines to the `.conf` file:

   ```
   net.ipv4.neigh.default.gc_thresh1 = 256

   net.ipv4.neigh.default.gc_thresh2 = 1024

   net.ipv4.neigh.default.gc_thresh3 = 2048
   ```

3. Reboot the server.

21.3.2 Making Live Changes to the Threshold Values

To make changes to a given threshold on the Orchestration Server you can run a command for each threshold that you want to change, for example:

```
# echo '256' > /proc/sys/net/ipv4/neigh/default/gc_thresh1
```

After you run the command, perform a `/etc/init.d/network restart` command to restart the Orchestration Server and put the changes in place.

NOTE: This method of changing the threshold values is volatile: if you reboot the SLES server, the changes are lost.

---

22 Orchestration Server Logs

As you discover and manage resources with the Cloud Manager Orchestration layer, that activity is logged in a number of locations. You can access that information for troubleshooting your environment yourself or for sending the information to NetIQ Technical Support for analysis and consultation.

Activity Logs

The following list includes the available activity logs and their respective locations on the Cloud Manager Orchestration Server:

- **Installation activity log:** `/var/opt/novell/novell_zenworks_orch_install.log`
- **Orchestration Server activity log:** `/var/opt/novell/zenworks/server/logs/server.log`
- **Orchestration Agent activity log:** `/var/opt/novell/zenworks/zos/agent/node.default/agent.log`
- **Orchestration Job activity log:** `/var/opt/novell/zenworks/server/dataGrid/files/users/<user_id>`
- **Orchestration Web Service activity log:** `/var/opt/novell/pso-ws/jetty/logs/`

Enabling a Log Trace of Calls to the Orchestration Server REST Interface

Cloud Manager includes a Karaf command that surfaces the timings of calls to the REST interface of the Cloud Manager Orchestration Server.

The following command turns on a log trace for the elapsed time of all REST calls to the server:

```
karaf> log:set TRACE com.novell.cm.psoservice.impl
```

The following command resets the log level to default so these timing messages no longer appear in the log:

```
karaf> log:set INFO com.novell.cm.psoservice.impl
```

Enabling NIXU API Logging

Cloud Manager searches the NIXU IPAM host only to check if a specified IP is registered in the network. Excessive logging by the NIXU API client is disabled by default.

To enable detailed NIXU API logging, you can enable the `apilogging` option in the `/opt/netiq/cloudmanager/etc/cloudmanager-ipam-nixu.cfg` file:

```
apilogging=true
```

The supportconfig Tool

NetIQ Technical Services maintains a tool called `supportconfig` that uses product plug-ins to gather logs and other information useful in troubleshooting. When directed by Support, you can run this tool on the Orchestration Server and send the needed information to NetIQ through a secure Internet connection.

For more information, contact NetIQ Technical Support:

Worldwide:	www.netiq.com/Support/contactinfo.asp
North and South America:	1-713-418-5555
Europe, Middle East, and Africa:	+353 (0) 91-782 677
Email:	support@netiq.com
Website:	www.netiq.com/support

V Application Server and Console Features and Functions

This section provides reference information about some features and functions of the NetIQ Application Server and Application Server Console you might need to know in order to adequately understand or efficiently troubleshoot and use these features and functions.

- Chapter 23, “The File Upload Feature,” on page 319

23 The File Upload Feature

Using the file upload procedure, you can upload any number of files to a Cloud Manager workload managed by the vSphere provisioning adapter. You can perform an upload in one session, as long as total size of the files does not exceed a maximum of 4.37 GB (the standard size of a DVD).

The file upload process can take a very long time, particularly when vSphere is configured to accept requests on the HTTP protocol only (we recommend that you allow both HTTP and HTTPS connections when configuring vSphere).

- Section 23.1, “Prerequisites to Using the Feature,” on page 319
- Section 23.2, “Feature Behaviors,” on page 319
- Section 23.3, “The “Preflight” Check,” on page 320
- Section 23.4, “Stages of the File Upload,” on page 320
- Section 23.5, “Canceling a File Upload,” on page 322
- Section 23.6, “Configuring vSphere for HTTP Connections,” on page 322

For information about how to execute a file upload, see “Adding Files to a Workload” in the *Cloud Manager Procedures Guide*.

23.1 Prerequisites to Using the Feature

The file upload feature does not function unless the following conditions are in place:

- The feature must be enabled in the System Configuration dialog.
- The workload where you intend to copy the files must have VNC console functionality.
- The workload where you intend to copy the files must include a virtual CD-ROM as an integrated drive. the selected CD-ROM needs to be empty/ejected. If not empty, the user will get a message indicating they need to eject the existing media first.
- The user executing the feature must be a valid, logged-in user.
- The Application Server, the Orchestration Server, and the vSphere Server all need `Write` rights to the various locations where the files are copied (see Section 23.4, “Stages of the File Upload,” on page 320), and each of these locations must have sufficient disk space to receive the files.

23.2 Feature Behaviors

The file upload feature has some behaviors you might notice as you upload a file (or files) to the VNC console of the workload:

- Uploading a directory is not supported:
  - Any directories in the set of files you drag to the VNC console are skipped in the upload.
  - If you attempt to upload directories only, you will receive an email message notifying you that no files were uploaded.
- Some Cloud Manager functions are not available during the file upload process:
  - You cannot submit a Change Request for a business service.

- You cannot use the lifecycle buttons (Start, Stop, Pause, Resume, etc.) on the workload dialog.
- If a file upload is already in progress, another upload cannot proceed until the first is complete.

- If multiple CD-ROM devices are defined, the feature uses them in this order:
  1. **A CD-ROM without media** (a blank CD-ROM).
  2. **A CD-ROM containing a previously uploaded set of files**. The VNC console displays an Eject icon. You must select this icon (that is, the CD must be “ejected” from the drive) before the upload can begin.
  3. **Any other CD-ROM device** (that is, a physical device, or one that points to some other ISO image). The VNC console displays an Eject icon. You must select this icon (that is, the CD must be “ejected” from the drive) before the upload can begin.
- Duplicate filenames being uploaded are renamed:
  - The first file with the same name keeps its name (`file_name.ext`).
  - The second file with the same name is renamed with a number (`file_name(1).ext`).

23.3 The “Preflight” Check

The Cloud Manager Web Console sends a “preflight” check request prior to beginning the upload of each file. This allows the Application Server to examine all the required conditions before attempting to upload the file. Should anything fail, you will see a failure message before the file is sent. This check cannot guarantee all of the desired conditions for the upload:

- Servers might run out of space during the file upload, even though conditions were favorable when it began.
- The Orchestration Server cannot examine many of the file upload error conditions during the preflight check. It reports such errors only during the final mounting stage.

23.4 Stages of the File Upload

Because the size of files you can upload can be quite large (several GB), unexpected interruption of the upload process (such as when a power failure occurs) might cause these large files to remain in one of the transfer locations, which could cause a significant and unanticipated load on server resources. Although Cloud Manager makes every effort to clean up these files, you need to know where to locate them so that you can manually delete them in the event of unusual circumstances.

This section includes the following information:

- Section 23.4.1, “The File Uploads to the Application Server,” on page 320
- Section 23.4.2, “The File is Published from the Application Server to the Orchestration Grid Repository,” on page 321
- Section 23.4.3, “The CD-ROM Mounts on the VM,” on page 321

23.4.1 The File Uploads to the Application Server

When you drag a file from the local file system and drop it on the VNC window, the first move for that file is to the Cloud Manager Web Service. Cloud Manager displays messages on the VNC console, such as this:

`Uploading (18%): filename.xxx`

The files are physically uploaded to <*ncm_writable_location*>`/.cloudmanager/plugins/cloudManager/workloads/vmUploadedFiles`.

If you are uploading multiple files, some browsers are able to copy them in parallel. You will see these messages change, referencing one file and then another, and back again. The upload does not progress to the next stage (publishing) until all of the files are copied to the Web Server.

The Web Server performs an auto-delete of files copied in the upload session if:

- The Web Server is rebooted or restarted.
- The user performing the file upload logs out, ending the session.
- Errors occur in the client or server during the file upload.
- When files are moved to the publishing stage of the upload process.

23.4.2 The File is Published from the Application Server to the Orchestration Grid Repository

The next stage of the file upload publishes the files to the Orchestration datagrid repository located on the Orchestration Server at `/var/opt/novell/zenworks/zos/server/dataGrid/files/upload/virtualmachines/`<*vmid*>`/`. At this transfer stage, Cloud Manager displays messages on the VNC console such as:

`Publishing (18%): filename.xxx`

NOTE: If you are publishing multiple files, Cloud Manager moves them in sequence, so messages on the VNC toolbar do not show operations alternating between files. The upload does not progress to the next stage (mounting) until all of the files are copied to the grid repository.

23.4.3 The CD-ROM Mounts on the VM

After all of the files are uploaded from the Application Server location to the datagrid repository, a `Mounting CD-ROM` message is displayed on the VNC Console while the Cloud Manager “back end” handles the files in the following sequence:

1. The Orchestration Server starts an Apply Config job to compress and consolidate the uploaded files into a single `.iso` file following the pattern `DataImage_`<*vdiskId*>`.iso`.
2. An Orchestration joblet runs and copies the `.iso` file to a temporary folder location used by the Orchestration Agent that is responsible for running vsphere jobs: <*Agent_Install_dir?*>`/node.default\<HostName>\dataGrid\upload\virtualmachines\`<*vmid*>`\`<*vdiskId*>`.iso.`
3. The vCenter Server uploads the `.iso` file from the Agent location to the vSphere ESXi host (that is, the location where the VM is running), and then into the `/file_transfer` subfolder of the VM’s repository.

4. An Orchestration Server job removes the file from the temporary, agent location.
    - The CD mounting stage can take the most time in the file upload process because of the hops between locations that are managed by the Orchestration Server. Cloud Manager retains the successfully uploaded files in its internal system for two days. At the end of this period, the Orchestration Server uses a scheduled job to auto-eject any mounted CD-ROM and to delete the associated uploaded files.
    - This `cleanupFileTransfers` job runs every day to detect old files (more than two days old) and remove the files from the datastore. It can be configured to increase or decrease its effective date range, depending on your need to maintain file upload storage and usage.

23.5 Canceling a File Upload

You can cancel the file upload process after it has started, no matter which stage the process is in. Cancelling the process auto-deletes any temporary files that might have been created.

IMPORTANT: If another user with access to the virtual environment ejects the CD on the VM, you can recover Cloud Manager using one of two methods:

- Run the **Synchronize** action, available as an option on the Workloads list.
- Click the Eject button on the VNC console.

23.6 Configuring vSphere for HTTP Connections

If upload speed to the vCenter server is problematic, you can change the server's security settings to allow it to receive both HTTP and HTTPS connections, then configure the vsphere job in the Orchestration Server to disable SSL. To change the settings:

1. At the vCenter server, change to the `C:\Users\All Users\VMware\VMware VirtualCenter` directory.
2. Use a text editor to open the `proxy.xml` file.
3. Navigate to this section of the file:

```
<e id="0">
<_type>vim.ProxyService.NamedPipeServiceSpec</_type>
<accessMode>httpsWithRedirect</accessMode>
<pipeName>\\.\pipe\vmware-vpxd-webserver-pipe</pipeName>
<serverNamespace>/</serverNamespace>
</e>
```

4. Change the value of the `accessMode` parameter from `httpsWithRedirect` to `httpANDhttps`.
5. Save the file.
6. Restart both the VMware VirtualCenter Server service and VMware VirtualCenter Webservices service.
7. Configure the Orchestration vsphere job to disable SSL:
    - **7a** In the Explorer tree of the Orchestration Server Console, expand the Jobs object to display the vsphere provisioning adapter job.
    - **7b** In the admin view, select the Job Configuration tab, then in the **Joblet Settings** panel of the page, deselect the **Enable SSL for Datastore File Upload** check box.
    - **7c** Save the settings and exit the console.

IMPORTANT: It is not possible to disable SSL for uploading files from the vCenter Server to the ESXi 5.1 host where the files are to reside. However, disabling SSL is possible to an ESXi 5.5 host, which might improve upload speed. For more information about the SSL disablement in vSphere 5.5, see VMware KB Article 2056830.

VI Rebranding the Cloud Manager Interface

If you have purchased a Cloud Manager license, you might be in a position to use the product as a service provider, as a reseller, or as an organization manager. In any of these roles, you can give Cloud Manager the look and feel that is proprietary to your own organization. This section includes information that you can use to make these branding customizations.

- Chapter 24, “The Cloud Manager Branding Infrastructure,” on page 327
- Chapter 25, “Rebranding the Application Server Console Interface,” on page 331
- Chapter 26, “Rebranding Reports,” on page 335
- Chapter 27, “Rebranding Email Messages,” on page 337
- Chapter 28, “Rebranding the Cloud Marketplace Mobile App,” on page 343

24 The Cloud Manager Branding Infrastructure

When you install Cloud Manager, it establishes an infrastructure that lets you customize the look and feel of various parts of the Cloud Manager interface, including email messages sent from the Cloud Manager system, costing reports, and the appearance of the Cloud Manager Application Console itself. This section includes conceptual information that explains the infrastructure installed by Cloud Manager to facilitate the branding of all these components.

- Section 24.1, “The Theming Directory Structure,” on page 327
- Section 24.2, “Creating a New Theme,” on page 327
- Section 24.3, “The Theming Configuration File,” on page 328
- Section 24.4, “Theming Search Order,” on page 329

24.1 The Theming Directory Structure

Cloud Manager uses a standard search method to find and implement “themes,” which consist of strings, colors, graphics, and layout in the visible parts of the product. At installation time, a default `/themes` folder is laid down at the root of the Cloud Manager Application Server:

```
/themes/CloudManager
```

Inside the `/CloudManager` folder is a localization layer, the default for which is `/en`:

```
/themes/CloudManager/en
```

Inside the `/en` folder is a “components” layer. By default, several components are laid down in this layer:

- `/themes/CloudManager/en/server`
- `/themes/CloudManager/en/reports`
- `/themes/CloudManager/en/webapp`

Inside the respective component folders, there might be an `/images` subfolder, populated with proprietary branding images that you can replace with your own, along with a number of other files (including XML or HTML files, style sheets, and so on) that you can modify to suit your branding needs.

24.2 Creating a New Theme

The `themes/CloudManager` folder (and its files) serves as a template for your branding effort. When you want to create a new theme, copy the default and customize it according to your needs. The procedures for customizing your own theme for these components, along with some examples, are included in other sections of this guide. See Chapter 25, “Rebranding the Application Server Console Interface,” on page 331, Chapter 26, “Rebranding Reports,” on page 335, Chapter 27, “Rebranding Email Messages,” on page 337. We recommend that you do not change anything in the default `/`

`CloudManager` theme folder, keeping it in pristine condition to copy and customize for new themes. This also allows for upgrading the `/CloudManager` default theme without affecting any other theme you are using.

Currently, the method used for applying branding to the Cloud Manager Mobile Clients does not use this theming method. For information on rebranding these clients, see Chapter 28, “Rebranding the Cloud Marketplace Mobile App,” on page 343.

24.3 The Theming Configuration File

When Cloud Manager components are called and their respective graphical interfaces rendered, Cloud Manager first looks at `/etc/cloudmanager-themes.cfg`, a configuration file that the theming system applies to a defined object (such as an organization). The default contents of the file are shown below:

```
server.theme = CloudManager
server.productname = NetIQ Cloud Manager
server.shortproductname = Cloud Manager
server.producturl = http://localhost:8182/CloudManager
```

These lines define the name you want to apply to the product throughout your theme, along with a URL link to the login of the Cloud Manager Application Server (a convenient way to embed server access into an email message).

This default theme is applied universally throughout the product, unless you create a new theme. You can do this by changing the values inside `cloudmanager-themes.cfg`, for example:

```
server.theme = DigitalAirCloudServices
server.productname = Digital Airlines Cloud Services
server.shortproductname = DAL Cloud
server.producturl = http://cloud.dal.com
```

NOTE: You should change only the contents of `cloudmanager-themes.cfg` when you want to use a new global theme. Do not change the name of the file.

You can also change the configuration file to discretely associate a new theme to the object where you want it applied. For example, if you create a new theme called “AcmeCloudServices” (a customized copy of the CloudManager theme), you might modify `/etc/cloudmanager-themes.cfg` to look like this:

```
server.theme = CloudManager
server.productname = NetIQ Cloud Manager
server.shortproductname = Cloud Manager
server.producturl = http://localhost:8182/CloudManager

164.theme = AcmeCloudServices
164.productname = Acme Cloud Services
164.server.shortproductname = Acme Cloud
164.server.producturl = http://localhost:8182/AcmeCloud
```

In the example above, the second section provides customization information for an organization with an ID of `164`. This informs Cloud Manager that when a user from the “Acme” organization (164) logs in, he or she will see the branding defined in the `AcmeCloudServices` theme.

You can add as many new theming sections to the configuration file as you want.

TIP: To obtain a list of all the organizations in your Cloud Manager system, open a browser to this URL:

`<your_cloudmanager_server:port>/cloudmanager-api/organizations`

This URL displays a report of all the organizations in your Cloud Manager system, along with their respective Cloud Manager IDs.

You can also find the organization ID for an individual organization on the Edit Organization dialog in the Application Server Console.

24.4 Theming Search Order

After Cloud Manager scans its theming configuration file to identify the themes it is to apply, it begins at the top of the theme hierarchy and searches for the presence of a localization layer or “locale” that matches system settings. Cloud Manager uses the following order as it searches for a locale and theme to apply to its components:

1. The locale of the User object (defined in LDAP).
2. The theme (defined in `cloudmanager-themes.cfg`, which points to `/themes/<theme>`) and the locale for the Organization object where the User object resides (see Section 24.3, “The Theming Configuration File,” on page 328).
3. The theme and locale defined for the system being used for viewing.
4. The default theme (`CloudManager`) and the default locale (`en`).

If Cloud Manager does not find a locale or theme in its first search that it can apply to its components, it proceeds to the next search level, and so on. With this theming hierarchy in mind, you can localize the interface, the system email messages, the reports for users in one of your organizations, or the look and feel of the Cloud Manager Application Server Console.

24.4.1 More About the Locale Folder

The Cloud Manager theming algorithm searches for the most specific variant of a language it can find. For example, if you define an `/en_gb`, locale folder, Cloud Manager looks in that folder for any files it might need before defaulting to the sibling, “less-specific” locale, `/en`. This makes it possible to customize your message text or colors to avoid any graphics, colors, or phrases that might be offensive in one locale but acceptable in another.

Because of the fallback nature of the search order, you might only customize a few files in a specific variant of a language. Cloud Manager only uses what it can find, so if a specific file is not present, it defaults to the sibling, “less specific” locale. If the parent locale is not found, Cloud Manager falls back to the installed locale, `/themes/CloudManager/en`.

25 Rebranding the Application Server Console Interface

This section includes information that can help you to apply a new look and feel (also known as a "theme") to the Cloud Manager Application Server web interface.

- Section 25.1, "Additional Theming Infrastructure," on page 331
- Section 25.2, "Theming Resources for the Cloud Manager Service Reseller," on page 331
- Section 25.3, "Theming the Remote Console Logo," on page 333

25.1 Additional Theming Infrastructure

In addition to the standard Cloud Manager theming infrastructure used to search for common themes and locales, the Cloud Manager Application Server Console adds more search structure to determine the search order for themes and resources. The list below details the order of this structure:

1. The theme and locale you define by specifying an additional parameter (`?theme=&locale=`) in the Cloud Manager URL.

   For example, if your Cloud Manager URL is `http://ncmdev.netiq.com/cloudmanager.html`, you could change its theming search to launch the console with a new theme and locale by adding the parameter: `http://ncmdev.netiq.com/cloudmanager.html ?theme=DigitalAirlines&locale=zh_CN`

   The parameter exhibits the following behaviors:

   - Persists in the console session until you reset the URL with a new parameter.
   - Invokes the settings when other parts of the console are accessed away from the login and home pages, even if the parameter is not included in the target URLs.
2. The theme and locale you define in an HTTP request header. You can set up this request to create a virtual host showing a specific theme.
3. The locales specified by the browser's `Accept-Language` request header, which lists locales in preferred order.
4. The common search methods outlined in Section 24.4, "Theming Search Order," on page 329.

25.2 Theming Resources for the Cloud Manager Service Reseller

One major theming use case is that of a retailer who wants to brand Cloud Manager branded with the identity of his service. In this scenario the retailer usually wants:

- A unique URL that identifies his service. For example, instead of `www.acmecloud.demo` he wants his version of Cloud Manager to be accessed with a URL such as `cloud.digitalairlines.demo`.
- A visually unique identity, to be shown throughout the Cloud Manager session, including before the user logs in.

To create your own retail theme in a similar scenario, use the following steps (shown with sample names):

1. Copy the `/opt/netiq/cloudmanager/themes/CloudManager` directory to `/opt/netiq/cloudmanager/themes/DigitalAirlines`.

   Remember, you create a new name for your own theme.

2. Customize the images, style sheets, HTML help files and other theme resources you want to brand in the `/opt/netiq/cloudmanager/themes/DigitalAirlines` folder, as you desire.

3. On the Cloud Manager Application Server, set up an Apache 2.*x* web server instance, configuring and enabling the `mod_proxy` and `mod_header` modules.

4. Create a virtual host for the new identity:

   4a Add the new virtual host by including a directive similar to the following in the Apache `httpd.conf` file:

   ```
   <VirtualHost *:80>
      ServerAdmin webmaster@digitalairlines.demo
      DocumentRoot "/usr/htdocs/cloud.digitalairlines.demo"
      ServerName cloud.digitalairlines.demo
      ErrorLog "/var/log/apache2/cloud.digitalairlines.demo-error_log"
      CustomLog "/var/log/apache2/cloud.digitalairlines.demo-access_log"
   common
      ProxyRequests off
      ProxyPass / http://www.acmecloud.demo:8182/
      ProxyPassReverse / http://www.acmecloud.demo:8182/
      RequestHeader set theme "DigitalAirlines"
   </VirtualHost>
   ```

   or

   4b Configure the Apache host to use SSL:

   4b1 In the `/etc/sysconfig/apache2` file, locate the line reading `APACHE_SERVER_FLAGS` and enable the SSL flag. The line should read as follows:

   ```
   APACHE_SERVER_FLAGS="SSL"
   ```

   4b2 Generate a SSL certificate for the new theme (`cloud.digitalairlines.demo`). For information about creating a self-signed SSL certificate, perform a search for information similar to what is found at Setting Up a Secure Server with SSL (https://www.suse.com/documentation/sles11/book_sle_admin/data/sec_apache2_ssl.html) in the *SUSE Linux Enterprise Server Administration Guide*.

   4b3 Copy `/etc/apache2/vhosts.d/vhost-ssl.template` to `/etc/apache2/vhosts.d/digitalairlines-ssl.conf` file, then modify the new `.conf` file as follows:

   4b3a In the file, change `ServerName` to `cloud.digitalairlines.demo:443`.

   4b3b Add the following prior to the `SSLEngine on` line:

   ```
   SSLProxyEngine on
   ProxyRequests off
   ProxyPreserveHost on
   ProxyPass / https://www.acmecloud.demo:8183/
   ProxyPassReverse / https://www.acmecloud.demo:8183/
   RequestHeader set theme "DigitalAirlines"
   ```

   4c (Conditional) If you want to create multiple virtual hosts with different IP addresses, change `<VirtualHost*:443>` to `<VirtualHost 1.1.1.1>:443`, replacing *1.1.1.1* with the virtual host's IP address. You also need to change the certificate lines to point to that virtual host's specific certificate.

TIP: To configure the server for multiple SSL retailers, you can either use a wildcard certificate with the same TLD (for example, the certificate is for `*.cloud.demo` and the virtual hosts are `digitalairlines.cloud.demo` and `acme.cloud.demo`), or you need to provide different IP address for each name, with a different certificate configured for each IP address (for example, `cloud.acme.demo` with an IP address of 1.1.1.1, and `cloud.digitalairlines.demo` with and IP address of 1.1.1.2).

25.3 Theming the Remote Console Logo

The VNC console logo image is controlled by the `/console/logo.js` file in the `theme/locale/web` folder. The file contains a JSON-encoded representation of an image file. To change the image, open the file and replace the base64 encoded data (located between '`base64`,' and '`}";)'` with data resulting from encoding a different `.jpg` file in base64.

TIP: You can easily find a base64 encoder online that you can use to convert the new branded image you want into base64 code.

26 Rebranding Reports

When you create a new theme for a Cloud Manager organization, you can customize the look and feel of the costing reports available to the users of that organization. For more information about creating a new theme and how Cloud Manager recognizes themes, see Chapter 24, “The Cloud Manager Branding Infrastructure,” on page 327.

26.1 File Locations and Contents for Report Theming

To customize the reports within a given theme, navigate to the `/themes/<new_theme_name>/<language>/reports` folder. There are two files in this folder that you can customize with your own look and feel:

- `header_logo.png`
- `styles.jrtx`

The corporate logo of the organization is embedded in the costing report `.pdf`. Copy the `header_logo.png` file to a graphic editor to determine its dimensions, then scale the corporate logo you want to use to the same dimensions, rename it to `header_logo` and save it in `.png` format in the `/reports` folder.

The following example shows the contents of the `styles.jrtx` file, which is a style sheet in `.xml` format:

```
<?xml version="1.0?>
<!DOCTYPE jasperTemplate
  PUBLIC "-//JasperReports//DTD Template//EN"
  "http://jasperreports.sourceforge.net/dtds/jaspertemplate.dte">

<jasperTemplate>

        <style name="title" backcolor=#660066" forecolor="#FF0000"/>
        <style name="pageHeader" backcolor=#FF00FF" forecolor="#0000FF"/>
        <style name="header1" backcolor=#9999FF" forecolor="#FF00FF"/>
        <style name="header2" backcolor=#9CB8FF" forecolor="#00FFFF"/>
        <style name="header3" backcolor=#E6E6FF" forecolor="#FF00000"/>
        <style name="header4" backcolor=#F5F5FF" forecolor="#FF00FF"/>
        <style name="tableHeader" backcolor=#F0FF0F" forecolor="#00FFFF"/>
        <style name="separator" backcolor=#00FFFF" forecolor="#FF0000"/>
        <style name="body" backcolor=#FFFFFF" forecolor="#F000000"/>
```

You need to experiment with this style sheet to determine the suitability of the organization’s colors or complementary colors when embedded in the report layout.

27 Rebranding Email Messages

When you create a new theme for a Cloud Manager organization, you can customize the look and feel of the email message formatted as reminders or pending tasks sent to the users of that organization. For more information about creating a new theme and how Cloud Manager recognizes themes, see Chapter 24, “The Cloud Manager Branding Infrastructure,” on page 327.

- Section 27.1, “File Locations and Contents for Mail Message Theming,” on page 337
- Section 27.2, “Tokens,” on page 338
- Section 27.3, “Email Testing Tools,” on page 341

27.1 File Locations and Contents for Mail Message Theming

To customize the email messages within a given theme, navigate to the `/themes/<new_theme_name>/<language>/server` folder. This folder holds a variety of theming resources.

- Section 27.1.1, “The main_skin.msg file,” on page 337
- Section 27.1.2, “The /images Folder,” on page 337
- Section 27.1.3, “The /styles Folder,” on page 337
- Section 27.1.4, “The /tasks Folder,” on page 338
- Section 27.1.5, “The .msg and .subject Files,” on page 338

27.1.1 The main_skin.msg file

The `main_skin.msg` file is the main “template” for the “wrapper” of the email message. Whenever a message is generated, Cloud Manager drops its contents into this template, which generates the general “look-and-feel” of the email message as specified, including the logo.

The file includes comments to help you understand how to insert your own image. You can reference the image by URL or by file location (see Section 27.1.2, “The /images Folder,” on page 337).

27.1.2 The /images Folder

The `/images` folder contains the images that are embedded in the email messages. You can use your organization logo by substituting it for the default. The sizing of the image is up to you.

27.1.3 The /styles Folder

Most email servers do not support cascading styles, so to provide a separation of the style from the message, each message can contain references to “style” files. These references let you define a theme once (for example, “error messages”) in a style file (located in the `/styles` folder) and reference it throughout several documents. To change that style later, you would change the style file.

27.1.4 The /tasks Folder

Cloud Manager has a tasks subsystem that places messages in the `/tasks` folder these tasks are dropped into the `main.skin` template and are sent as alerts to the Cloud Manager user who is assigned to complete a task in the Cloud Manager workflow.

27.1.5 The .msg and .subject Files

These are the actual messages (and subjects) that are injected into the compiled message wrapper. They can reside in the `/server` folder or in the `/tasks` subfolder. These are the actual files that you can customize to meet the branding needs of your organization. To help you identify each unique task or message file, the email message template displays its filename in light gray text at the bottom of the message.

27.2 Tokens

Some of the contents of the message and task files are tokenized, allowing the system to run macros that fetch elements to include in the body or subject of the message.

$$INCLUDE:<*filename*.msg>$$:

This is the name of a file you want to invoke and display in a message. It could include such things as predefined tables. The macro infers that the message file is found in the current folder. This allow a part of a message (such as a details table that might be included in multiple messages) to be defined only once and then referenced where needed. This minimizes challenges when you change the included message.

You can change the name or the look and feel (including the contents) of the default tables or you can customize your own table, based on one of these:

- `businessServiceDetailsTable.msg`
- `businessServiceWithWorkloadsDetailsTable.msg`
- `changeRequestDetailsTable.msg`
- `changeRequestWithWorkloadsChangesDetailsTable.msg`
- `changeRequestWithWorkloadsDetailsTable.msg`

$$STYLE:<*filename*.style>$$:

This is the name of a style file you want to invoke and display in a message. It could include such things as headers. The macro infers that the style file is found in the /styles folder.

- **$$PRODUCT_NAME$$:** This is the formal name of the Cloud Manager system, renamed as you want it to appear in your rebranded version of the product. The value for this macro is found in the theming configuration file (`/etc/cloudmanager-themes.cfg`). For more information, see Section 24.3, “The Theming Configuration File,” on page 328.
- **$$PRODUCT_SHORT_NAME$$:** This is the informal or “short” name of the Cloud Manager system, renamed as you want it to appear in your rebranded version of the product. The value for this macro is found in the theming configuration file (`/etc/cloudmanager-themes.cfg`). For more information, see Section 24.3, “The Theming Configuration File,” on page 328.
- **$$PRODUCT_URL$$:** This is the URL link to Cloud Manager system you have installed and configured for use in your organization. Because it is used in all messages, this macro is included in the `main_skin.msg` file.

- **$$CONTENT$:** This token is found only in the `main_skin.msg` template. It is the insertion point for the editable `.msg` file content you control (that is, the various messages, subjects, and task messages found in the theming infrastructure).

$$EMBED_IMAGE:<*image_name*>$$:

This token references a file in the `/themes` directory. It directs Cloud Manager to calculate a Base-64 encoded string, and then embeds that string in the message.

- **$$ATTACH_IMAGE:<image_name>$$:** This token references a file in the `/themes` directory. It directs Cloud Manager to attach an image to the email. It also adds a content inside the email message that references the attachment.

The following list of macros is available for use only in core messages. The purpose of each is self-explanatory. Use the form ${*MACRO_NAME*} when you tokenize a message.

- ATTACH_IMAGE
- BS_ID
- BS_NAME
- BS_ID_NAME
- BS_DESCRIPTION
- BS_STATUS
- BS_SUBMISSION_DATE
- BS_START_DATE
- BS_EXPIRE_DATE
- BS_REQUESTER
- BS_REQUESTER_ID
- BS_DECOMMISSIONED_DAT
- BS_BUSINESS_GROUP_NAME
- BS_ORG_NAME
- BS_STATUS_ADDL_INFO
- CURRENT_TIME_STAMP
- ZONE_ID
- ZONE_NAME
- ZONE_DESCRIPTION
- ZONE_CONNECTION_INFO
- WL_NAME
- WL_OSTYPE
- WL_OSFAMILY
- WL_BASECOST
- WL_NUMBER_OF_CPUS
- WL_AMOUNT_OF_MEMORY
- WL_VM_NAME
- WL_REPOSITORY_ID
- WL_TEMPLATE_NAME
- WL_SERVICE_LEVEL_NAME

- WL_ZONE_NAME
- WL_RESOURCE_GROUP_NAME
- WL_IMPORTED
- WL_DISKS_SUMMARY
- WL_NIC_CONFIGS_SUMMARY
- WL_AOS_SUMMARY

The following list of macros is available for use only for messages in the /tasks subfolder. The purpose of each is self-explanatory. Use the form ${*MACRO_NAME*} when you tokenize a message.

- ATTACH_IMAGE
- BS_ID
- BS_NAME
- BS_ID_NAME
- BS_DESCRIPTION
- BS_STATUS
- BS_SUBMISSION_DATE
- BS_START_DATE
- BS_EXPIRE_DATE
- BS_REQUESTER
- BS_REQUESTER_ID
- BS_DECOMMISSIONED_DATE
- BS_BUSINESS_GROUP_NAME
- BS_ORG_NAME
- CR_ORIGBS_ID
- CURRENT_TIME_STAMP
- TASK_ID
- TASK_SUBJECT
- TASK_ID_SUBJECT
- TASK_MESSAGE
- TASK_STATUS
- TASK_CALLOUT_GUID
- TASK_DATE_CREATED
- TASK_DATE_COMPLETED
- TASK_COMPLETE_BY_DATE
- TASK_DELEGATED_BY
- TASK_DELEGATED_BY_ID
- TASK_ASSIGNED_USER
- TASK_ASSIGNED_USER_ID
- TASK_ASSIGNED_PERMISSION

27.3 Email Testing Tools

As you customize the contents of the mail messages according to your needs, you might want to use the following tools to help you in the customization process:

- Karaf shell commands.

  Cloud Manager uses a caching system in its search tools to track the locales it has searched and what is available for use in localizing messages. Some Karaf commands help to manage that cache.

  - Karaf shell command: `cm:flush-branding-resources-cache`

    The configuration file and some of the resources (ie styles) are read only once. This command flushes the resources from the cache so that new modified values can take effect.

  - Karaf shell command: `cm:dump-branding-resources-cache`

    This command displays the current values of certain stored settings options.

  - Karaf shell command: `cm:send-branded-test-message`

    This command mimics sending a branded test message from the Cloud Manager file system as if you were a user. The command lets you use the default email messaging system to test the look and feel of the message.

    The following table shows the options that you can use with the command to pass in information that might be required to send a tokenized message.

Short Option Name	Long Option Name	Description
`-b`	`-bsID`	The business service ID. This option is useful when you want to test a message with a tokenized business service ID and workloads that are members of the business service.
`-t`	`-to`	The message recipient ID. Cloud Manager uses this option to determine the email address to send to, along with the correct locale of the recipient. **NOTE:** If you always test sending to an administrative user, you will always get the system theme. To test custom organization-based themes, you need to test sending to a user in that organization.
`-z`	`-zoneID`	The zone ID. This option is useful when you want to test a message with a tokenized zone ID.
`-m`	`-msgID`	The message ID (that is, the filename of the message, displayed at the bottom of the content body).

    Syntax: `cm:send-branded-test-msg [`*`option`*`] <`*`parameter`*`>`

    Example: `cm:send-branded-test-msg -m ncm.businessservice.deployed -b 221 -t 141`

- `https://putsmail.com`

  This is a web-based tool provided by Litmus Putsmail that you can use to mock up the messages and send test mails to validate the look and feel.

28 Rebranding the Cloud Marketplace Mobile App

If you want users of the Cloud Marketplace mobile app for iPad to see proprietary branding of Cloud Manager customized for their organization, you need to copy the re-themed resources (that is, proprietary text and graphics) to the Cloud Manager Application Server.

Cloud Marketplace uses images installed on the Cloud Manager Application Server to rebrand its user interface. The default images you can use to re-theme Marketplace are located in a downloadable `.zip` file (https://www.netiq.com/documentation/cloud-manager-24/resources/marketplace_rebrand.zip) at the Cloud Manager documentation website.

To re-theme these images:

1. Extract the contents of the `.zip` file.
2. Evaluate the contents. You can change the colors, the logo, and the product name (embedded in the `.json` file) in your re-theming. Some things to remember:
   - When a graphic designer creates the new brand images, the image size, format, and filename must remain the same. The image resources can help the designer maintain these criteria.
3. After you customize the images to match your own branding, save them and compress them in a file named `marketplace_rebrand.zip`.

   NOTE: Remember to save the rebranding files in a flat hierarchy; there should be no folder inside the `.zip` file.
4. Upload your `marketplace_rebrand.zip` file to the Cloud Manager Application Server at

   `/opt/netiq/cloudmanager/themes/CloudManager/en/mobile/marketplace_rebrand.zip`

   You must create this file path if it does not already exist.
5. Verify that the images display properly in the Cloud Marketplace app.

   A Cloud Manager Application servlet serves up `marketplace_rebrand.zip`. When the Cloud Marketplace iPad app launches, it retrieves this `.zip` file, extracts it and uses the images and information within to change its appearance.

   You can check the contents of the newly themed `.zip` file from a browser at

   `http://<cloud_manager_application_server>/theme/mobile/marketplace_rebrand.zip`
6. If the servlet attempts to find the file and fails, or if the image you replaced displays incorrectly in Marketplace, you need to use the following Karaf command to flush the cache and refresh it with a new image from the `marketplace_rebrand.zip` file:

   ```
   cm:flush-branding-resources-cache
   ```

NOTE: The language aspect of Marketplace rebranding is determined by headers that the Marketplace client provides to the server. The theme name is based on the hostname that the iPad application uses to access the server. The server can be configured to respond to several virtual host names, and each can have its own theme.

For more information, see Section 25.1, “Additional Theming Infrastructure,” on page 331.

VII Application Server Troubleshooting

Use the information in this section to troubleshoot your Cloud Manager Application Server.

- Chapter 29, “Troubleshooting Cloud Manager Application Issues,” on page 347
- Chapter 30, “Cloud Manager Application Server Logs,” on page 351

29 Troubleshooting Cloud Manager Application Issues

This section provides information that identifies some ongoing known issues in Cloud Manager Application components and the methods that you can use to address those issues.

29.1 Business Services Issues

The following sections provide solution to the problems you might encounter while using thCloud Manager:

- “Business Service Workloads Remain in the Building or Provision State” on page 347
- “Remote Console for a Workload Does Not Launch from Internet Explorer with Chrome Frame Plug-in” on page 347
- “Multi-Select Behavior Is Incorrect for Multiple Virtual Machines in a List” on page 348
- “Enabling Email Notifications for Users with Specific Permission” on page 348

Business Service Workloads Remain in the Building or Provision State

Source: Cloud Manager Application Server

Explanation: During the building phase of a new business service’s workload and the startup phase of a deployed workload, it is possible for the workload to be unable to be assigned to a host. This can occur when no hosts in the host group have the available resources to meet the workload resource requirements.

In the build phase, the business service remains in the **Building** state until a host becomes available for the workload. In the startup phase, the workload remains in the **Provision** state until a host becomes available.

Action: You have several options to resolve this issue:

- Shut down a workload to free up the required resources on a host. If possible, select a workload that can be restarted on another host.
- Add another host to the host group.

Remote Console for a Workload Does Not Launch from Internet Explorer with Chrome Frame Plug-in

Source: Cloud Manager Application Server

Explanation: If you are using Microsoft Internet Explorer 9 with the Google Chrome Frame plug-in and a self-signed security certificate (in a test environment) to view the Cloud Manager Application Server console, and if you try to launch a remote session of a workload in a business service, the remote session fails.

Testing has shown that the failure occurs because Chrome Frame does not allow a self-signed security certificate.

NOTE: This procedure is effective only if the VNC Proxy is configured to use the same certificate as the Cloud Manager Application Server console. For information about configuring the VNC proxy, see "Configuring Remote Console Access to Workloads" in the *Cloud Manager Procedures Guide*.

Action: You can force Chrome Frame to accept the self-signed certificate as a trusted root in Windows by using the following steps:

1. Open Internet Explorer as an Administrator, then, using Internet Explorer to open the Cloud Manager console, note the warning from the browser about an unrecognized certificate.
2. At the warning dialog, specify that you want to continue to the website.
3. In the browser address bar, click the certificate warning, then click **View Certificates**.
4. In the certificate pop-up, click **Install Certificate** then click **Next**.

   TIP: The **Install Certificate** button does not display unless you are running the browser as Administrator.
5. Select **Place all certificates in the following store**, then click **Browse**.
6. From the list, select **Trusted Root Certification Authorities**, then click **OK**, click **Next**, then click **Finish**.
7. At the confirmation dialog, click **Yes**, then, at the next confirmation dialog, click **Close**.
8. At the certificate popup, click **Close**.

Multi-Select Behavior Is Incorrect for Multiple Virtual Machines in a List

Source: Cloud Manager Application Server

Explanation: Testing has shown that the normal procedure for selecting a group of virtual machines (for instance, select **Business Services** > *Virtual Machines*, select a VM, scroll down the list, then Shift+click to multi-select) listed across several pages does not multi-select the desired VMs in the list. Further, in such a long list, your first selection on the list might be "lost" (deselected) when you scroll down the list to multi-select.

Action: Use the Ctrl+A command to select all VMs in the list. This issue will be fixed in the next product release.

Enabling Email Notifications for Users with Specific Permission

Source: Cloud Manager Application Server

Explanation: You can control the individuals who can receive email notification when a business service is deployed or a change request has completed if you add the `ncm.bs.deploy.perms` property to the `/opt/netiq/cloudmanager/etc/system.properties` file.

Action: As you add this new property, you also need to add the permissions to be honored on the business service. Any user with those permissions will receive the email notifications. For example, if you wanted a user with `SYSTEM_SUPPORT` permission or `MODIFY_BS` or `VIEW_BS` permission to receive emails when the business service being deployed or changed, you would modify the `/opt/netiq/cloudmanager/etc/system.properties` file like this:

```
...
...
ncm.bs.deploy.perms=SYSTEM_SUPPORT,MODIFY_BS,VIEW_BS
...
...
```

Additional Info: You might occasionally add these permissions:

Name	Description
`MODIFY_BS`	Modify a business service request
`VIEW_BS`	View a business service request
`SYSTEM_SUPPORT`	Designated system support personnel

Additional Info: You might add these permissions less frequently:

Name	Description
`MODIFY_SYSTEM`	Modify the system object
`MODIFY_ORG`	Modify an organization
`VIEW_ORG`	View an organization
`MODIFY_BG`	Modify a business group
`VIEW_BG`	View a business group

30 Cloud Manager Application Server Logs

As you manage resources and workflow with the Cloud Manager application layer, that activity is logged in a number of locations. You can access that information for troubleshooting your environment yourself or for sending the information to NetIQ Technical Support for analysis and consultation.

The following list includes the available activity logs and their respective locations on the Cloud Manager Application Server:

- **Application Server Installation activity log:** `/opt/netiq/cloudmanager/logs/netiq_cloudmanager_config.log`
- **Application Server activity log:** `/opt/netiq/cloudmanager/logs/cloudmanager_server.log`

The supportconfig tool

NetIQ Technical Services maintains a tool called supportconfig that uses product plug-ins to gather logs and other information useful in troubleshooting. When directed by Support, you can run this tool on the Application Server and send the needed information to NetIQ through a secure Internet connection.

For more information, contact NetIQ Technical Support:

Worldwide:	+1-416-203-4799
North America:	+1-800 -858-4000
Email:	support@netiq.com
Website:	www.netiq.com/support