

Oracle® Enterprise Manager

Cloud Control Extensibility Programmer's Reference

12c Release 3 (12.1.0.3)

E25161-09

July 2013

Copyright © 2013, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle USA, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Contents

Preface	xxi
Audience	xxi
Documentation Accessibility	xxi
Related Documents	xxi
Conventions	xxi
 1 Getting Started with Plug-in Development	
1.1 About the Plug-in Creation Process	1-1
1.2 About the Extensibility Development Kit (EDK)	1-2
1.2.1 Contents of the EDK	1-2
1.3 Installing the Extensibility Development Kit (EDK)	1-3
1.4 Designing the Plug-in	1-4
1.5 Creating a Basic Plug-in	1-4
1.6 Creating an Intermediate Plug-in	1-5
1.7 Creating an Advanced Plug-in	1-5
 2 Defining the Plug-in	
2.1 Introduction to Defining the Plug-in	2-1
2.2 Basic Plug-in Metadata	2-2
2.2.1 Defining the Plug-in ID	2-2
2.2.2 Defining the Plug-in Version	2-2
2.3 Creating Plug-in Definition Files	2-3
2.4 Creating the plugin.xml File	2-3
2.4.1 Overview of plugin.xml Elements	2-4
2.4.2 Certifying Plug-ins	2-5
2.5 Creating the plugin_registry.xml File	2-7
2.5.1 Overview of plugin_registry.xml Elements	2-7
2.6 Validating the Plug-in Definition Files	2-8
2.7 Adding Log Viewer Support to Your Plug-in	2-9
 3 Creating Target Metadata Files	
3.1 Introduction to Creating Target Metadata Files	3-1
3.2 Overview of Target Definition Files	3-2
3.3 Creating the Target Type Metadata File	3-3
3.3.1 Creating a Basic Target Type Metadata File	3-3

3.3.2	Naming the Target Type Metadata File	3-4
3.3.3	Defining the Target Type Metadata	3-4
3.3.4	Defining Target Credentials	3-5
3.3.5	Defining Type Properties.....	3-5
3.3.6	Defining Instance Properties	3-7
3.4	Defining Metrics to Collect from the Target	3-8
3.4.1	Metric Definition Files.....	3-9
3.4.2	Defining the Basic Response Metric Group	3-9
3.4.3	Defining Advanced Metrics	3-11
3.4.4	Overview of Key Metric Metadata Elements.....	3-12
3.5	Creating the Default Collection File	3-15
3.5.1	Grouping Similar Metrics For Collection	3-16
3.5.2	Defining Basic Metric Collection	3-17
3.5.3	Defining Advanced Metric Collection	3-17
3.5.4	Defining Target Configuration Data Collections	3-18
3.5.5	Overview of Key Default Collection Metadata Elements	3-18
3.6	Guidelines for Creating Target Metadata.....	3-21
3.6.1	Defining Target Metadata	3-21
3.6.2	Defining Collections	3-23
3.7	Testing Your Target Type Definitions	3-24
3.7.1	Activate the Metric Browser.....	3-24
3.7.2	View Your Metrics	3-25
3.8	Validating Your Metadata XML	3-25
3.9	Troubleshooting the Target Creation Process.....	3-25

4 Adding Information Publisher Reports

4.1	Introduction to Adding Information Publisher Reports	4-1
4.1.1	Assumptions and Prerequisites	4-2
4.2	Overview of SYSTEM Reports	4-2
4.2.1	About the Report Definitions Page	4-2
4.3	Understanding the Report Definition File.....	4-2
4.4	Creating a Report Definition File.....	4-3
4.4.1	About the Report Definition File Development Process.....	4-3
4.4.2	About the Report Lifecycle: Updating Report Definitions	4-5
4.5	Understanding the XML Report Definition Interface.....	4-5
4.5.1	About Report Definition Tags.....	4-5
4.5.2	Using Element Parameters	4-7
4.5.3	Understanding the Metric Details Element	4-26
4.5.4	Using Text Element Parameters	4-28
4.5.5	About Report-Wide Parameters	4-29
4.6	Using the ImportExport.xsd File	4-29
4.7	About Enterprise Manager Command Line Interface (EM CLI) Verbs	4-32
4.8	About Development Guidelines.....	4-33

5 Developing BI Publisher Reports

5.1	Introduction to Oracle BI Publisher	5-1
5.1.1	Assumptions and Prerequisites	5-1

5.2	Training and Resources.....	5-2
5.3	About the Report Data Source	5-2
5.4	Developing a Report.....	5-2
5.5	Using the Enterprise Manager EDK for Staging and Deploying BI Publisher Reports ...	5-3

6 Collecting Target Configuration Data

6.1	Introduction to Collecting Target Configuration Data.....	6-1
6.1.1	Assumptions and Prerequisites	6-2
6.2	About the Configuration Definition Files.....	6-2
6.3	Modeling Enterprise Configuration Management Tables	6-2
6.3.1	Defining Configuration Collection Tables	6-4
6.3.2	Overview of Configuration Management Snapshot Metadata Elements	6-9
6.3.3	Packaging Configuration Metadata	6-13
6.3.4	Registering Metadata With the Configuration Management Framework	6-13
6.3.5	Supporting Translation	6-14
6.3.6	Upgrading Configuration Data	6-15
6.3.7	Modifications to Standard Collection Metrics and RAW Metrics	6-17
6.3.8	Testing the Configuration Collection Data	6-19
6.3.9	Troubleshooting.....	6-19

7 Adding Job Types

7.1	Introduction to Adding Job Types	7-1
7.2	About Job Types.....	7-2
7.3	Introducing New Job Types	7-3
7.4	Specifying a New Job Type in XML	7-3
7.4.1	Understanding Job Type Categories.....	7-4
7.4.2	Using Agent-Bound Job Types	7-4
7.4.3	About Job Steps.....	7-5
7.5	Using Commands	7-8
7.5.1	About the remoteOp Command.....	7-9
7.5.2	Using the fileTransfer Command.....	7-10
7.5.3	About the putFile Command	7-11
7.5.4	Using the getFile Command	7-12
7.5.5	Using the execAndSuspend Command	7-13
7.6	About Command Error Codes	7-13
7.7	Executing Long-Running Commands at the Oracle Management Service.....	7-14
7.7.1	Configuring the Job Dispatcher to Handle Long-Running Commands.....	7-14
7.8	Specifying Parameter Sources	7-14
7.8.1	Understanding SQLParameter Source	7-15
7.8.2	About the User Parameter Source	7-17
7.8.3	About the Inline Parameter Source	7-18
7.8.4	Using the checkValue Parameter Source.....	7-19
7.8.5	About the properties Parameter Source	7-19
7.8.6	Understanding Parameter Sources and Parameter Substitution.....	7-20
7.8.7	About Parameter Encryption	7-20
7.9	Specifying Credential Information.....	7-20

7.9.1	About Credential Usage	7-21
7.9.2	Overview of Credential Binding	7-21
7.9.3	XSD Elements – Credential Usage and Credential Binding	7-22
7.10	Specifying Security Information	7-23
7.11	Specifying Lock Information	7-25
7.12	Suspending a Job or Step	7-28
7.13	Restarting a Job.....	7-28
7.13.1	Restarting Versus Resubmitting.....	7-28
7.13.2	Default Restart Behavior.....	7-28
7.13.3	Using the restartMode Directive	7-29
7.14	Adding Job Types to the Job Activity and Job Library Pages	7-32
7.14.1	Adding a Job Type to the Job Activity Page	7-32
7.14.2	Adding a Job Type to the Job Library Page	7-34
7.15	Examples: Specifying Job Types in XML.....	7-34
7.16	About Performance Issues.....	7-42
7.16.1	Using Parameter Sources.....	7-43
7.17	Adding a Job Type to Enterprise Manager	7-43

8 Defining a Management User Interface

8.1	Introduction to Defining a Management User Interface	8-2
8.1.1	Flex Implementation	8-3
8.1.2	Metadata-only Implementation.....	8-4
8.1.3	Assumptions and Prerequisites	8-5
8.2	MPCUI Concepts.....	8-5
8.2.1	Integration Class	8-5
8.2.2	Activity	8-5
8.2.3	Page.....	8-5
8.2.4	Services.....	8-6
8.2.5	URL	8-6
8.3	UI Options for a Plug-in.....	8-6
8.3.1	Metadata-only Implementation.....	8-7
8.3.2	Flex Implementation	8-7
8.4	Creating the MPCUI Metadata File.....	8-8
8.4.1	Overview of MPCUI Metadata Elements.....	8-10
8.5	Defining Metadata	8-12
8.5.1	Limitations of the Metadata Implementation.....	8-12
8.5.2	Defining Integration Metadata	8-12
8.5.3	Defining Navigation.....	8-18
8.6	Defining the MPCUI Application.....	8-20
8.6.1	Defining the Application Activities (Integration Class).....	8-21
8.6.2	Defining Pages	8-22
8.6.3	Defining Dialogs	8-25
8.6.4	Defining Trains and Train Pages	8-26
8.6.5	Defining URLs.....	8-27
8.7	Packaging the MPCUI Implementation With the Plug-in	8-27
8.8	Converting a Metadata-based UI to a Flex-based UI.....	8-28
8.9	Defining System Home Pages	8-28

8.9.1	Defining systemUiIntegration Metadata.....	8-31
8.9.2	Defining System Regions.....	8-33
8.10	Defining Navigation.....	8-34
8.10.1	Navigation to Activities.....	8-35
8.10.2	URL and Links	8-35
8.10.3	Adding Links to External Applications.....	8-36
8.11	Accessing Enterprise Manager Data	8-37
8.11.1	Metric Services	8-37
8.11.2	Custom Data Source	8-41
8.11.3	Computed Data Source	8-43
8.11.4	Packaged SQL and the Query Service	8-45
8.11.5	Working With Target Services.....	8-48
8.11.6	Monitoring Service Request Performance.....	8-50
8.11.7	Automated Polling of Service Requests	8-51
8.11.8	Batching of Service Requests.....	8-52
8.11.9	Software Library Search Service	8-53
8.12	Performing Task Automation	8-54
8.12.1	Automation Services	8-54
8.12.2	Working With Credentials	8-60
8.13	Storing Session State	8-66
8.14	Defining Page Layout Components	8-67
8.14.1	Defining Regions.....	8-68
8.15	Including Packaged Regions	8-69
8.15.1	Availability Region.....	8-69
8.15.2	Incidents and Problems Region.....	8-69
8.15.3	Job Summary Region.....	8-70
8.15.4	Credentials Region	8-70
8.16	Defining Charts	8-70
8.16.1	Line Chart	8-70
8.16.2	Area Chart.....	8-72
8.16.3	Bar (Horizontal) Chart	8-73
8.16.4	Column (Vertical Bar) Chart	8-74
8.16.5	Pie Chart.....	8-75
8.17	Defining Tables	8-75
8.17.1	Data Service	8-75
8.17.2	Custom Data Provider	8-76
8.17.3	Getting Selected Rows.....	8-77
8.18	Defining Dialogs	8-77
8.18.1	Dialog Registration.....	8-77
8.18.2	Displaying a Dialog and Waiting for Close Events	8-78
8.19	Defining Trains.....	8-79
8.19.1	Train Definition Example	8-79
8.19.2	Train Controller	8-80
8.19.3	Train State	8-80
8.19.4	Train Events.....	8-80
8.20	Defining Information Item and Information Displays (Label-Value Pairs).....	8-81
8.21	Using Built-in Renderers	8-82

8.22	Defining Links	8-83
8.23	Including Enterprise Manager Images	8-83
8.24	Displaying a Processing Cursor.....	8-83
8.25	Defining a Processing Window	8-84
8.26	Defining Icons for Target Types	8-85
8.27	Displaying the Target Navigator	8-86
8.28	Defining a UI for Guided Discovery	8-87
8.28.1	About Guided Discovery.....	8-87
8.28.2	Supporting Guided Discovery	8-87
8.28.3	Constructing the Guided Discovery User Interface	8-88
8.28.4	Structure of the Discovery Application.....	8-90
8.28.5	Using Discovery Service	8-90
8.28.6	Using Target Information Services.....	8-92
8.28.7	Using Target Management Services.....	8-92
8.29	About Logging	8-93
8.29.1	Add Logging to your Code	8-93
8.29.2	Options for Capturing Log Output.....	8-94
8.30	Development Environment Options.....	8-95
8.30.1	Developing MPCUI with Flex SDK and Apache Ant	8-96
8.30.2	Developing MPCUI in Adobe Flash or Flex Builder	8-97
8.30.3	Setting Up an Adobe Flash Builder Project for MPCUI.....	8-101
8.31	Migrating Home Page Customizations	8-107
8.32	Accessibility Guidelines	8-107
8.32.1	Accessibility Options in Enterprise Manager	8-108
8.32.2	Summary of Critical Issues.....	8-108
8.32.3	Using MPCUI Components.....	8-109
8.32.4	Controlling Reading and Tabbing Order	8-109
8.33	Localization Support	8-110
8.34	Providing Online Help	8-110

9 Customizing Incident Manager

9.1	Introduction to Customizing Incident Manager	9-1
9.2	Understanding Supported Customizations	9-2
9.3	Creating Event-Specific Customization XML	9-3
9.3.1	Overview of Event-Specific Customization Metadata Elements.....	9-4
9.3.2	About Events	9-5
9.4	Adding Customized Details About the Event.....	9-12
9.5	Providing Content in the Guided Resolution Region	9-14
9.5.1	Adding Recommendations using XML	9-15
9.5.2	Customizing Sections	9-16
9.6	Defining a Search String for My Oracle Support Knowledge	9-16
9.7	Defining Conditions for Customization	9-17
9.8	Registering Customizations	9-18
9.9	Testing Incident Manager After Customization.....	9-18

10 Using Derived Associations

10.1	Introduction to Derived Associations.....	10-1
------	---	------

10.1.1	Assumptions and Prerequisites	10-2
10.2	Understanding Enterprise Manager Association Concepts	10-3
10.2.1	About Out-of-Box Association Types	10-3
10.2.2	Using Association Derivation	10-4
10.2.3	About Automated Discovery and Promotion of Associations	10-4
10.2.4	Understanding Association Creation During Guided Discovery	10-4
10.2.5	Using Associations Derived from a System Stencil	10-5
10.2.6	Associations Derived from Rule.....	10-5
10.3	About Association Derivation Rules Management	10-5
10.3.1	Using Association Derivation Rules Syntax and Semantics.....	10-6
10.3.2	Understanding XML Metadata File Syntax and Semantics.....	10-9
10.3.3	Using Rule Semantics.....	10-14
10.3.4	Maintaining Performance	10-14
10.3.5	About Regular Query and Trigger Patterns	10-15
10.3.6	Diagnosing Issues	10-16
10.3.7	Useful Examples	10-16
10.3.8	Applying the Mechanical Steps of Integration.....	10-19
10.3.9	Special Triggers: Host Name Change Triggers	10-20
10.3.10	Understanding Activation Expressions.....	10-20
10.3.11	About Debugging	10-23
10.4	Ensuring Performance.....	10-25
10.4.1	Using Custom Configuration Specifications for Data Collection.....	10-26
10.5	Using Overlapping Associations	10-26
10.5.1	Overlap Between Associations Derived by Rules	10-26
10.6	Creating Associations for Composite and System Target Types	10-27
10.6.1	Composite Membership and the Containment Association	10-27
10.6.2	Other Non-Composite Associations (Composite Topology)	10-27
10.6.3	System Membership Associations.....	10-27
10.6.4	Associations to External Targets	10-27
10.6.5	Regarding the Timing of Association Creation.....	10-28
10.7	Frequently Asked Questions	10-28
10.7.1	Which tables can I reference in a rule query?	10-28
10.7.2	Are there guidelines for when to use target properties?	10-29
10.7.3	What is the relationship between discovered and derived associations?	10-29

11 Defining Target Discovery

11.1	Introduction to Defining Target Discovery	11-1
11.2	Creating Discovery XML	11-2
11.2.1	Generic Discovery Integration Example	11-2
11.2.2	Discovery Script Example.....	11-3
11.2.3	Overview of the Discovery Metadata Elements.....	11-4
11.3	Creating the Discovery Script	11-5
11.3.1	Discovered Targets DTD	11-5
11.4	Packaging Discovery XML and Discovery Content	11-6
11.4.1	Location of the Discovery Metadata File.....	11-6
11.4.2	Package Discovery Content.....	11-6
11.5	Setting Up and Testing Discovery	11-8

11.6	Manually Adding Targets	11-8
11.6.1	Manually Adding Host Targets.....	11-9
11.6.2	Manually Adding Non-Host Targets.....	11-9
11.7	Configuring and Promoting Targets for Monitoring by Enterprise Manager	11-10
11.8	Examples for Using Generic Discovery Framework	11-10
11.8.1	Discovery Integration Example Requiring User Input	11-10
11.9	Configuring Automatic Discovery For Plug-ins	11-12

12 Adding Compliance Standards

12.1	Introduction to Adding Compliance Standards.....	12-1
12.1.1	Assumptions and Prerequisites	12-2
12.2	About the Compliance Standard Rules	12-2
12.2.1	Defining Repository Check-based Rules.....	12-2
12.2.2	Defining Real-time Monitoring Rules.....	12-8
12.3	Defining Compliance Standards.....	12-21
12.4	Defining a Compliance Framework	12-24
12.5	Defining Compliance Content	12-26
12.6	Removing Compliance Content.....	12-29
12.7	Supporting Translation	12-29
12.8	Packaging Compliance XML	12-32
12.9	Setting Up and Testing Compliance Standards and Rules	12-32
12.9.1	Install Compliance Content.....	12-32
12.9.2	Test Compliance Standard	12-32
12.9.3	Constraints for Testing.....	12-34
12.10	More Compliance Examples	12-34
12.11	Publishing Compliance Content Using Self Update.....	12-49

13 Validating, Packaging, and Deploying the Plug-in

13.1	Introduction to Validating, Packaging, and Deploying the Plug-in.....	13-1
13.2	Staging the Plug-in.....	13-2
13.3	Validating the Plug-in	13-5
13.4	Creating the Plug-in Archive	13-6
13.5	Importing and Deploying the Plug-in Archive into Enterprise Manager	13-8
13.5.1	Prerequisites for Importing the Plug-in.....	13-8
13.5.2	Importing the Plug-in Archive	13-8
13.5.3	Deploying the Plug-in on Oracle Management Service (OMS)	13-9
13.5.4	Important Details Regarding Plug-in Deployment	13-10
13.6	Adding a Target Instance.....	13-11
13.7	Updating Deployed Metadata Files Using the Metadata Registration Service (MRS)	13-12
13.7.1	Target Types and Default Collections	13-13

14 Defining Software Library Metadata

14.1	Introduction to Software Library Framework	14-1
14.2	Defining Software Library Metadata	14-2
14.2.1	Defining Folders.....	14-2
14.2.2	Defining Types	14-3

14.2.3	Defining Subtypes	14-3
14.2.4	Defining Entities	14-6
14.3	Organizing Software Library Metadata Files	14-7
14.4	Adding the Software Library Metadata to Enterprise Manager	14-7
14.4.1	Step 1: Validating Metadata XML	14-7
14.4.2	Step 2: Adding Metadata XML to OPAR	14-8
14.5	Using Software Library Entities.....	14-8
14.5.1	Using Job Types	14-9
14.5.2	Using EMCLI Verbs.....	14-11
15	Defining Credentials	
15.1	Introduction to Security Concepts.....	15-1
15.1.1	Understanding Credential Types	15-2
15.1.2	About Named Credentials	15-2
15.1.3	Authenticating Target Types	15-3
15.1.4	Overview of Credential Sets.....	15-3
15.1.5	Using the Credential Store	15-3
15.1.6	About the Credential Reference	15-3
15.2	Defining Credential Metadata.....	15-4
15.2.1	Overview of Credential Elements	15-5
16	Converting an Existing Metadata Plug-in	
16.1	Introduction to Converting an Existing Metadata Plug-in	16-1
16.2	Impact of the New Plug-in Framework on Existing Plug-ins	16-2
16.2.1	Plug-in Metadata Must Be Converted to the New Format.....	16-2
16.2.2	Plug-ins Must Be Packaged as Oracle Plug-in Archive (OPAR) Files	16-2
16.2.3	In-place Upgrade of Existing Plug-ins Not Supported	16-2
16.2.4	Plug-ins Should Be Registered Before Upgrading	16-2
16.3	Using the convert_mp Utility to Convert Plug-in Metadata	16-2
16.3.1	Converting Plug-In Metadata Stored In A Management Repository	16-3
16.3.2	Converting a Metadata Plug-in Archive (MPA)	16-5
16.4	Post-Conversion Steps.....	16-6
16.4.1	Modifying the plugin.xml and plugin_registry.xml Files	16-6
16.4.2	Converting Report Definitions Created With PL/SQL.....	16-7
16.4.3	Converting Job Type Definitions.....	16-8
16.4.4	Packaging the Plug-in	16-9
17	Monitoring Using Web Services and JMX	
17.1	Overview	17-1
17.2	Monitoring Using Web Services in Enterprise Manager.....	17-2
17.2.1	Creating Metadata and Default Collection Files	17-2
17.3	Monitoring Using WS-Management in Enterprise Manager	17-11
17.3.1	Creating Metadata and Default Collection Files	17-11
17.4	Monitoring JMX Applications Deployed on Oracle Application Servers (OC4J)	17-18
17.4.1	Creating Metadata and Default Collection Files	17-18
17.4.2	Displaying Target Status Information	17-28

17.5	Monitoring a Standalone JMX-instrumented Java Application or JVM Target.....	17-31
17.5.1	Generating Metadata and Default Collection Files.....	17-32
17.5.2	Using the Metadata and Default Collection Files	17-37
17.6	Monitoring JMX Applications Deployed on Oracle WebLogic Application Servers ..	17-37
17.6.1	Creating Metadata and Default Collection Files using jmxcli.....	17-37
17.6.2	Using the Metadata and Default Collection Files	17-45
17.7	Adding a Target to a Management Agent.....	17-46
17.7.1	Adding a Web Services Target - CalculatorService	17-46
17.7.2	Adding a WS-Management Target - TrafficLight.....	17-47
17.7.3	Configuring a Standalone Java Application or JVM Target.....	17-47
17.7.4	Adding a Target Instance for a Custom J2EE Application on WebLogic.....	17-51
17.8	Monitoring Credential Setup	17-53
17.9	Viewing Monitored Metrics	17-54
17.10	Creating JMX Metric Extensions.....	17-55
17.10.1	Using the Enterprise Manager Console.....	17-55
17.10.2	Using the JMXCLI to create a Metric Extension Archive.....	17-68
17.11	Surfacing Metrics from a Standalone JVM or Oracle Coherence.....	17-71
17.11.1	Using the Enterprise Manager Console.....	17-71
17.11.2	Using JMXCLI	17-71

18 Using the Web Services Framework

18.1	Using APIs to Write Java Clients.....	18-1
18.1.1	About SOAPMessageBuilder.....	18-1
18.1.2	About WebServiceInvoker	18-2
18.1.3	Using the InvokeAPITester Sample Program.....	18-3

19 Using Management Repository Views

19.1	Overview	19-2
19.1.1	Using Repository Views	19-2
19.2	Application Deployment Views	19-3
19.2.1	MGMT\$J2EE_APPLICATION	19-3
19.2.2	MGMT\$J2EEAPP_EJBCOMPONENT	19-4
19.2.3	MGMT\$J2EEAPP_JRFWS.....	19-4
19.2.4	MGMT\$J2EEAPP_JRFWSOPER.....	19-5
19.2.5	MGMT\$J2EEAPP_JRFWSPOLICY	19-5
19.2.6	MGMT\$J2EEAPP_JRFWSUPPORT	19-6
19.2.7	MGMT\$J2EEAPP_WEBAPPCOMPONENT	19-7
19.2.8	MGMT\$J2EEAPP_WSCONFIG	19-8
19.2.9	MGMT\$J2EEAPP_WSPORTCONFIG	19-8
19.3	Blackout Views.....	19-9
19.3.1	MGMT\$BLACKOUT_HISTORY	19-9
19.3.2	MGMT\$BLACKOUTS.....	19-10
19.4	Chargeback Views	19-11
19.4.1	MGMT\$EMCT_CBA_CHARGE_HOURLY	19-11
19.4.2	MGMT\$EMCT_CBA_CHARGE_DAILY	19-13
19.5	Compliance Views	19-15
19.5.1	MGMT\$COMPLIANCE_STANDARD_RULE.....	19-15

19.5.2	MGMT\$COMPLIANCE_STANDARD.....	19-17
19.5.3	MGMT\$COMPLIANCE_STANDARD_GROUP	19-18
19.5.4	MGMT\$CS_EVAL_SUMMARY	19-19
19.5.5	MGMT\$COMPOSITE_CS_EVAL_SUMMARY	19-21
19.5.6	MGMT\$CS_RULE_EVAL_SUMMARY	19-23
19.5.7	MGMT\$CS_GROUP_EVAL_SUMMARY.....	19-24
19.5.8	MGMT\$CS_TARGET_ASSOC.....	19-24
19.5.9	MGMT\$CSR_CURRENT_VIOLATION	19-25
19.5.10	MGMT\$CSR_VIOLATION_CONTEXT	19-26
19.5.11	MGMT\$EM_RULE_VIOL_CTXT_DEF	19-27
19.6	Compliance Real-time Monitoring Views.....	19-27
19.6.1	MGMT\$CCC_ALL_OBS_BUNDLES	19-27
19.6.2	MGMT\$CCC_ALL_OBSERVATIONS.....	19-28
19.6.3	MGMT\$CCC_ALL_VIOLATIONS	19-29
19.6.4	MGMT\$COMPLIANT_TARGETS	19-30
19.6.5	MGMT\$COMPLIANCE_SUMMARY	19-31
19.6.6	MGMT\$COMPLIANCE_TREND.....	19-31
19.7	Configuration Management Views	19-32
19.7.1	MGMT\$CSA_COLLECTIONS.....	19-32
19.7.2	MGMT\$CSA_FAILED	19-35
19.7.3	MGMT\$CSA_HOST_OS_COMPONENTS.....	19-36
19.7.4	MGMT\$CSA_HOST_SW	19-37
19.7.5	MGMT\$CSA_HOST_COOKIES	19-37
19.7.6	MGMT\$CSA_HOST_CUSTOM.....	19-37
19.7.7	MGMT\$CSA_HOST_RULES	19-38
19.7.8	MGMT\$CSA_HOST_CPUS.....	19-38
19.7.9	MGMT\$CSA_HOST_IOCARDS.....	19-39
19.7.10	MGMT\$CSA_HOST_NICS.....	19-39
19.7.11	MGMT\$CSA_HOST_OS_PROPERTIES.....	19-40
19.7.12	MGMT\$CSA_HOST_OS_FILESYSEMS	19-40
19.8	Custom Configuration Specification Views.....	19-40
19.8.1	MGMT\$CCS_DATA.....	19-41
19.8.2	MGMT\$CCS_DATA_SOURCE	19-41
19.8.3	MGMT\$CCS_DATA_VISIBLE.....	19-42
19.8.4	MGMT\$CCS_DATA.....	19-43
19.9	Database Configuration Views	19-44
19.9.1	MGMT\$DB_TABLESPACES.....	19-44
19.9.2	MGMT\$DB_DATAFILES	19-45
19.9.3	MGMT\$DB_CONTROLFILES	19-46
19.9.4	MGMT\$DB_DBNINSTANCEINFO	19-46
19.9.5	MGMT\$DB_FEATUREUSAGE	19-47
19.9.6	MGMT\$DB_INIT_PARAMS.....	19-48
19.9.7	MGMT\$DB_LICENSE.....	19-49
19.9.8	MGMT\$DB_REDOLOGS.....	19-49
19.9.9	MGMT\$DB_ROLLBACK_SEGS.....	19-50
19.9.10	MGMT\$DB_SGA	19-51
19.9.11	MGMT\$DB_TABLESPACES_ALL.....	19-52

19.9.12	MGMT\$DB_OPTIONS.....	19-52
19.10	Events Views.....	19-53
19.10.1	MGMT\$INCIDENTS.....	19-53
19.10.2	MGMT\$INCIDENT_CATEGORY.....	19-54
19.10.3	MGMT\$INCIDENT_TARGET.....	19-55
19.10.4	MGMT\$INCIDENT_ANNOTATION.....	19-55
19.10.5	MGMT\$EVENTS_LATEST.....	19-55
19.10.6	MGMT\$EVENTS.....	19-56
19.10.7	MGMT\$EVENT_ANNOTATION.....	19-57
19.10.8	MGMT\$PROBLEMS.....	19-58
19.10.9	MGMT\$PROBLEM_ANNOTATION.....	19-59
19.11	Glassfish Views.....	19-60
19.11.1	MGMT\$EMAS_GLASSFISH_DOMAIN.....	19-60
19.11.2	MGMT\$EMAS_GLASSFISH_NODES.....	19-60
19.11.3	MGMT\$EMAS_GLASSFISH_SERVER.....	19-61
19.11.4	MGMT\$EMAS_GLASSFISH_SVR_PROP.....	19-61
19.11.5	MMGMT\$EMAS_GLASSFISH_NW_LSTNR.....	19-61
19.11.6	MGMT\$EMAS_GLASSFISH_DATASOURCE.....	19-62
19.11.7	MGMT\$EMAS_GLASSFISH_DS_PROP.....	19-63
19.12	Hardware Views.....	19-63
19.12.1	MGMT\$HW_CPU_DETAILS.....	19-63
19.12.2	MGMT\$HW_NIC.....	19-64
19.12.3	MGMT\$HW_NIC_BONDS.....	19-64
19.12.4	MGMT\$HW_IO_DEVICES.....	19-65
19.13	Inventory Views.....	19-65
19.13.1	MGMT\$TARGET.....	19-65
19.13.2	MGMT\$TARGET_TYPE.....	19-66
19.13.3	MGMT\$TARGET_TYPE_DEF.....	19-68
19.13.4	MGMT\$TARGET_ASSOCIATIONS.....	19-68
19.13.5	MGMT\$TARGET_MEMBERS.....	19-69
19.13.6	MGMT\$TARGET_FLAT_MEMBERS.....	19-69
19.13.7	MGMT\$TARGET_TYPE_PROPERTIES.....	19-70
19.13.8	MGMT\$TARGET_PROPERTIES.....	19-70
19.14	Job Views.....	19-71
19.14.1	MGMT\$CA_TARGETS.....	19-71
19.14.2	MGMT\$CA_EXECUTIONS.....	19-71
19.14.3	MGMT\$JOBS.....	19-74
19.14.4	MGMT\$JOB_TARGETS.....	19-75
19.14.5	MGMT\$JOB_EXECUTION_HISTORY.....	19-76
19.14.6	MGMT\$JOB_STEP_HISTORY.....	19-78
19.14.7	MGMT\$JOB_ANNOTATIONS.....	19-79
19.14.8	MGMT\$JOB_NOTIFICATION_LOG.....	19-79
19.15	Linux Patching Views.....	19-80
19.15.1	MGMT\$HOSTPATCH_HOSTS.....	19-80
19.15.2	MGMT\$HOSTPATCH_GROUPS.....	19-80
19.15.3	MGMT\$HOSTPATCH_GRP_COMPL_HIST.....	19-81
19.15.4	MGMT\$HOSTPATCH_HOST_COMPL.....	19-81

19.16	Management Template Views	19-81
19.16.1	MGMT\$TEMPLATES.....	19-81
19.16.2	MGMT\$TEMPLATE_POLICY_SETTINGS.....	19-82
19.16.3	MGMT\$TEMPLATE_METRICCOLLECTION.....	19-83
19.16.4	MGMT\$TEMPLATE_METRIC_SETTINGS.....	19-84
19.17	Metric Views	19-86
19.17.1	MGMT\$METRIC_CATEGORIES	19-86
19.17.2	MGMT\$METRIC_COLLECTION	19-87
19.17.3	MGMT\$METRIC_ERROR_CURRENT.....	19-88
19.17.4	MGMT\$METRIC_ERROR_HISTORY	19-88
19.18	Monitoring Views	19-89
19.18.1	MGMT\$ALERT_CURRENT.....	19-89
19.18.2	MGMT\$TARGET_METRIC_COLLECTIONS.....	19-91
19.18.3	MGMT\$TARGET_METRIC_SETTINGS	19-92
19.18.4	MGMT\$AVAILABILITY_CURRENT	19-94
19.18.5	MGMT\$AVAILABILITY_HISTORY.....	19-95
19.18.6	MGMT\$ALERT_HISTORY	19-95
19.18.7	MGMT\$METRIC_DETAILS.....	19-97
19.18.8	MGMT\$METRIC_CURRENT	19-98
19.18.9	MGMT\$METRIC_HOURLY	19-99
19.18.10	MGMT\$METRIC_DAILY.....	19-101
19.19	Operating System Views.....	19-102
19.19.1	MGMT\$OS_SUMMARY.....	19-102
19.19.2	MGMT\$OS_COMPONENTS.....	19-103
19.19.3	MGMT\$OS_HW_SUMMARY	19-103
19.19.4	MGMT\$OS_PATCH_SUMMARY.....	19-104
19.19.5	MGMT\$OS_FS_MOUNT.....	19-105
19.19.6	MGMT\$OS_KERNEL_PARAMS.....	19-105
19.19.7	MGMT\$OS_PATCHES	19-106
19.19.8	MGMT\$OS_PROPERTIES.....	19-106
19.19.9	MGMT\$OS_MODULES.....	19-106
19.19.10	MGMT\$OS_LIMITS	19-107
19.19.11	MGMT\$OS_INIT_SERVICES.....	19-107
19.20	Oracle Home Directory Patching Views.....	19-107
19.20.1	MGMT\$EM_HOMES_PLATFORM.....	19-107
19.20.2	MGMT\$HOMES_AFFECTED.....	19-108
19.20.3	MGMT\$APPL_PATCH_AND_PATCHSET	19-108
19.20.4	MGMT\$APPLIED_PATCHES	19-108
19.20.5	MGMT\$APPLIED_PATCHSETS.....	19-109
19.21	Oracle Home Directory Views	19-109
19.21.1	MGMT\$OH_HOME_INFO	19-109
19.21.2	MGMT\$OH_DEP_HOMES	19-110
19.21.3	MGMT\$OH_CRS_NODES.....	19-110
19.21.4	MGMT\$OH_CLONE_PROPERTIES	19-111
19.21.5	MGMT\$OH_COMPONENT.....	19-111
19.21.6	MGMT\$OH_COMP_INST_TYPE	19-112
19.21.7	MGMT\$OH_COMP_DEP_RULE.....	19-112

19.21.8	MGMT\$OH_PATCHSET.....	19-113
19.21.9	MGMT\$OH_VERSIONED_PATCH	19-113
19.21.10	MGMT\$OH_PATCH.....	19-114
19.21.11	MGMT\$OH_PATCHED_COMPONENT	19-115
19.21.12	MGMT\$OH_PATCH_FIXED_BUG	19-115
19.21.13	MGMT\$OH_PATCHED_FILE.....	19-116
19.21.14	MGMT\$OH_FILE	19-116
19.21.15	MGMT\$PA_RECOM_METRIC_SOURCE.....	19-117
19.21.16	MGMT\$OH_INV_SUMMARY	19-117
19.21.17	MGMT\$OH_INSTALLED_TARGETS.....	19-118
19.22	Oracle WebLogic Server Views.....	19-118
19.22.1	MGMT\$WEBLOGIC_APPLICATIONS	19-118
19.22.2	MGMT\$WEBLOGIC_EJBCOMPONENT	19-119
19.22.3	MGMT\$WEBLOGIC_FILESTORE	19-119
19.22.4	MGMT\$WEBLOGIC_JDBCdatasource.....	19-119
19.22.5	MGMT\$WEBLOGIC_JDBCMULTIDS	19-121
19.22.6	MGMT\$WEBLOGIC_JMSConnectionFactory.....	19-121
19.22.7	MGMT\$WEBLOGIC_JMSQUEUE	19-122
19.22.8	MGMT\$WEBLOGIC_JMSSERVER.....	19-122
19.22.9	MGMT\$WEBLOGIC_JMSTOPIC	19-123
19.22.10	MGMT\$WEBLOGIC_JOLTCONNPOOL	19-123
19.22.11	MGMT\$WEBLOGIC_JVMSYSPROPS.....	19-124
19.22.12	MGMT\$WEBLOGIC_MACHINE	19-124
19.22.13	MGMT\$WEBLOGIC_NETWORK_CHANNELS.....	19-125
19.22.14	MGMT\$WEBLOGIC_NODEMANAGER.....	19-125
19.22.15	MGMT\$WEBLOGIC_RACONFIG.....	19-126
19.22.16	MGMT\$WEBLOGIC_RAOUTBOUNDCONFIG.....	19-126
19.22.17	MGMT\$WEBLOGIC_RESOURCECONFIG	19-128
19.22.18	MGMT\$WEBLOGIC_SERVER	19-128
19.22.19	MGMT\$WEBLOGIC_STARTSHUTCLASSES.....	19-130
19.22.20	MGMT\$WEBLOGIC_VIRTUALHOST	19-130
19.22.21	MGMT\$WEBLOGIC_WEBAppComponent.....	19-131
19.22.22	MGMT\$WEBLOGIC_WORKMANAGER.....	19-131
19.22.23	MGMT\$WEBLOGIC_WSCONFIG	19-132
19.22.24	MGMT\$WEBLOGIC_WSPORTCONFIG.....	19-132
19.23	Oracle WebLogic Domain Views.....	19-133
19.23.1	MGMT\$WEBLOGIC_DOMAIN.....	19-133
19.23.2	MGMT\$WEBLOGIC_OPSSSYSPROP	19-133
19.23.3	MGMT\$WEBLOGIC_OAMCONFIG.....	19-134
19.24	Oracle WebLogic Cluster Views.....	19-134
19.24.1	MGMT\$WEBLOGIC_CLUSTER	19-134
19.25	Security Views.....	19-135
19.25.1	MGMT\$ESA_ALL_PRIVS_REPORT	19-135
19.25.2	MGMT\$ESA_ANY_DICT_REPORT	19-135
19.25.3	MGMT\$ESA_ANY_PRIV_REPORT	19-135
19.25.4	MGMT\$ESA_AUDIT_SYSTEM_REPORT	19-136
19.25.5	MGMT\$ESA_BECOME_USER_REPORT	19-136

19.25.6	MGMT\$ESA_CATALOG_REPORT.....	19-136
19.25.7	MGMT\$ESA_CONN_PRIV_REPORT.....	19-137
19.25.8	MGMT\$ESA_CREATE_PRIV_REPORT	19-137
19.25.9	MGMT\$ESA_DBA_GROUP_REPORT.....	19-137
19.25.10	MGMT\$ESA_DBA_ROLE_REPORT	19-138
19.25.11	MGMT\$ESA_DIRECT_PRIV_REPORT.....	19-138
19.25.12	MGMT\$ESA_EXMPT_ACCESS_REPORT.....	19-138
19.25.13	MGMT\$ESA_KEY_OBJECTS_REPORT	19-138
19.25.14	MGMT\$ESA_OH_OWNERSHIP_REPORT	19-139
19.25.15	MGMT\$ESA_OH_PERMISSION_REPORT.....	19-139
19.25.16	MGMT\$ESA_POWER_PRIV_REPORT.....	19-139
19.25.17	MGMT\$ESA_PUB_PRIV_REPORT	19-140
19.25.18	MGMT\$ESA_SYS_PUB_PKG_REPORT.....	19-140
19.25.19	MGMT\$ESA_TABSP_OWNERS_REPORT.....	19-140
19.25.20	MGMT\$ESA_TRC_AUD_PERM_REPORT	19-140
19.25.21	MGMT\$ESA_WITH_ADMIN_REPORT	19-141
19.25.22	MGMT\$ESA_WITH_GRANT_REPORT	19-141
19.25.23	MGMT\$ESM_COLLECTION_LATEST	19-141
19.25.24	MGMT\$ESM_FILE_SYSTEM_LATEST.....	19-142
19.25.25	MGMT\$ESM_PORTS_LATEST	19-142
19.25.26	MGMT\$ESM_SERVICE_LATEST	19-142
19.25.27	MGMT\$ESM_STACK_LATEST	19-142
19.26	Service Tag Views.....	19-142
19.26.1	MGMT\$SERVICETAG_INSTANCES.....	19-142
19.26.2	MGMT\$SERVICETAG_REGISTRY	19-143
19.27	Storage Reporting Views	19-143
19.27.1	MGMT\$STORAGE_REPORT_DATA.....	19-143
19.27.2	MGMT\$STORAGE_REPORT_KEYS	19-144
19.27.3	MGMT\$STORAGE_REPORT_PATHS.....	19-144
19.27.4	MGMT\$STORAGE_REPORT_ISSUES	19-145
19.27.5	MGMT\$STORAGE_REPORT_DISK.....	19-145
19.27.6	MGMT\$STORAGE_REPORT_VOLUME.....	19-145
19.27.7	MGMT\$STORAGE_REPORT_LOCALFS	19-146
19.27.8	MGMT\$STORAGE_REPORT_NFS.....	19-146
19.28	Target Views.....	19-147
19.28.1	MGMT\$AGENTS_MONITORING_TARGETS.....	19-147
19.28.2	MGMT\$EM_ECM_MOS_PROPERTIES.....	19-148
19.28.3	MGMT\$EM_ECM_TARGET_FRESHNESS.....	19-148
19.28.4	MGMT\$MANAGEABLE_ENTITIES.....	19-148
19.29	VT Target Views.....	19-150
19.29.1	MGMT\$VT_VM_CONFIG	19-150
19.29.2	MGMT\$VT_VM_SW_CFG	19-151
19.29.3	MGMT\$VT_VM_VNIC.....	19-151
19.29.4	MGMT\$VT_VM_VNIC_QOS	19-152
19.29.5	MGMT\$VT_VM_EM_CFG	19-152
19.29.6	MGMT\$VT_VM_VDISKS.....	19-152
19.29.7	MGMT\$VT_VM_VDISK_QOS	19-153

19.29.8	MGMT\$VT_EXA_CTRL_VSERVER_TAGS	19-153
19.29.9	MGMT\$VT_VSP_CONFIG.....	19-153
19.29.10	MGMT\$VT_VS_HW_CFG	19-154
19.29.11	MGMT\$VT_VS_HYPERVISOR	19-155
19.29.12	MGMT\$VT_VS_PROCESSORS	19-155
19.29.13	MGMT\$VT_VS_SW_CFG.....	19-156
19.29.14	MGMT\$VT_VS_ATTRIBUTES	19-157
19.29.15	MGMT\$VT_VS_FS_MOUNTS.....	19-157
19.29.16	MGMT\$VT_VS_NET_DEVICE.....	19-157
19.29.17	MGMT\$VT_VS_FILESERVERS	19-158
19.29.18	MGMT\$VT_VS_REPOS.....	19-158
19.29.19	MGMT\$VT_VS_ABILITIES.....	19-158
19.29.20	MGMT\$VT_ZONE_CONFIG	19-158
19.30	Examples	19-159

20 Using Receivelets

20.1	About Receivelets.....	20-1
20.2	SNMP Receivelet.....	20-1

21 Using Fetchlets

21.1	About Fetchlets.....	21-1
21.2	OS Command Fetchlets.....	21-2
21.2.1	OS Fetchlet	21-2
21.2.2	OSLines Fetchlet (split into lines).....	21-4
21.2.3	OSLineToken Fetchlet (tokenized lines).....	21-6
21.2.4	Invoke an OS Fetchlet as a Specific User for Metric Collection	21-8
21.3	SQL Fetchlet	21-9
21.4	SNMP Fetchlet.....	21-15
21.5	HTTP Data Fetchlets.....	21-19
21.5.1	URL Fetchlet (raw).....	21-20
21.5.2	URL Lines Fetchlet (split into lines).....	21-21
21.5.3	URL Line Token Fetchlet (tokenized lines).....	21-22
21.6	URLXML Fetchlet	21-23
21.7	URL Timing Fetchlet	21-25
21.8	Dynamic Monitoring Service (DMS) Fetchlet.....	21-30
21.8.1	Advantages to Using DMS for Oracle Management Agent Integration	21-30
21.8.2	DMS Fetchlet/Oracle Management Agent Integration Instructions	21-33
21.9	JDBC Fetchlet.....	21-35
21.10	WBEM Fetchlet.....	21-37
21.11	JMX Fetchlet.....	21-41
21.12	Web Services Fetchlet.....	21-44
21.12.1	Using Credentials for Authentication.....	21-48
21.13	WS-Management Fetchlet.....	21-50
21.13.1	Using Credentials	21-53
21.14	REST Fetchlet.....	21-54
21.14.1	Response Processing	21-55
21.14.2	Using HTTPS and Self-Signed Certificates	21-58

21.14.3	Using REST CLI to Generate Metadata	21-59
---------	---	-------

A Out-of-Box Associations

Preface

This document covers Enterprise Manager framework extensibility and related reference information.

Note: For the most current version of this document, go to the **Extensibility** page of the Oracle Enterprise Manager Online Documentation set:

<http://www.oracle.com/pls/em121/homepage>

Audience

This document is intended for plug-in developers that want to extend Oracle Enterprise Manager to support the ability to manage custom target types or extend the manageability of out-of-box target types.

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at

<http://www.oracle.com/pls/topic/lookup?ctx=acc&id=docacc>.

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit

<http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info> or visit

<http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs> if you are hearing impaired.

Related Documents

For more information about Oracle Enterprise Manager, see the Oracle Enterprise Manager Online Documentation set:

<http://www.oracle.com/pls/em121/homepage>

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
<code>monospace</code>	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

Getting Started with Plug-in Development

This chapter contains the following sections:

- [About the Plug-in Creation Process](#)
- [About the Extensibility Development Kit \(EDK\)](#)
- [Installing the Extensibility Development Kit \(EDK\)](#)
- [Designing the Plug-in](#)
- [Creating a Basic Plug-in](#)
- [Creating an Intermediate Plug-in](#)
- [Creating an Advanced Plug-in](#)

1.1 About the Plug-in Creation Process

Creating a plug-in involves the following steps, all of which are covered in this book:

1. Designing your plug-in.
2. Developing the plug-in, which includes creating the requisite metadata files that enable the plug-in functionality.
3. Staging the plug-in.
4. Validating the plug-in.
5. Packaging the plug-in as an archive (.opar) file.
6. Importing the plug-in into Enterprise Manager Cloud Control.
7. Deploying the plug-in to Oracle Management Service.
8. Adding a target from your environment to initiate target monitoring. The plug-in files required by the Management Agent to monitor the target will be pushed to the Agent at this time.
9. Testing the plug-in to verify that it is behaving as expected.

As you continue to modify your plug-in metadata, you can upload your new metadata files to Enterprise Manager without re-deploying the plug-in archive using the Metadata Registration Service. See [Section 13.7, "Updating Deployed Metadata Files Using the Metadata Registration Service \(MRS\)"](#) for instructions on using this service.

In addition, to keep track of each updated version of your plug-in, you should incrementally update the plug-in version as follows in the following plug-in metadata files:

- In the `PluginVersion` attribute in the `plugin.xml` file that describes the plug-in to Oracle Management Service the plug-in is deployed to. See [Section 2.4, "Creating the plugin.xml File"](#).
- In the `Version` attribute in the `plugin_registry.xml` file that describes the plug-in to Management Agents the plug-in is deployed to. See [Section 2.5, "Creating the plugin_registry.xml File"](#).

1.2 About the Extensibility Development Kit (EDK)

A key component of the Enterprise Manager Cloud Control architecture is the Extensibility framework. To enable Oracle partners to extend the Enterprise Manager platform, an Extensibility Development Kit (EDK) is provided with the product.

The EDK is a collection of tools, utilities, and documentation, including:

- Enterprise Manager Extensibility documentation: Provide general guidelines for programming Enterprise Manager plug-ins
- Reference Implementation: Provides a reference code implementation, code snippets, and so on for various Enterprise Manager features
- Build time tools to verify EDK conformance: A tool that you can use to validate and report any violations, with respect to Enterprise Manager Extensibility guidelines
- Packaging Tool: A tool to package the plug-in components tool (`empdk`)
- Verification Tool: A tool to validate plug-in code components and to report violations (if any).

Enterprise Manager EDK includes a command line utility called `empdk`. Use this utility to package or validate a plug-in archive. For information associated with the `empdk` commands and their options, see [Section 13.3, "Validating the Plug-in"](#).

After you download the EDK, unzip it on your local system, and change your current directory to the location where you unpacked the EDK. The EDK contains reference documentation and guides to help you with plug-in development as well as the API reference you might need to integrate while developing plug-ins.

For information about downloading the EDK, see [Section 1.3, "Installing the Extensibility Development Kit \(EDK\)"](#).

1.2.1 Contents of the EDK

The EDK archive contains the following directories:

- `/bin`
Contains the `empdk` utility, which you use to:
 - Validate the structure of your plug-in
 - Package your plug-in
 - Convert the metadata for existing (pre-Cloud Control 12) plug-ins to the new metadata formats
- `/doc`
Contains the Oracle Enterprise Manager Extensibility Programmer's Guide and Programmer's Reference, as well as the EDK API Reference documentation,

including documentation on Management Views. Review `overview.html` for links to the documentation provided.

You can also access the EDK API Reference documentation directly through its index page (`sdk_api_ref.html`).

- `/lib`

Contains internal libraries used by the EDK.

- `/oui`

Contains internal libraries used by the EDK.

- `/samples`

Contains a complete reference implementation of a plug-in, packaged as `demo_hostsample.zip`. The sample metadata files included should be used as examples of the files referenced throughout the EDK documentation.

View the README packaged with the archive for instructions on building, deploying, and using the sample plug-in.

Other utilities referenced in this documentation, including EM CLI and EM CTL, are installed with Enterprise Manager and are typically deployed to the Oracle Management Service (OMS) host.

1.3 Installing the Extensibility Development Kit (EDK)

Before you begin developing plug-ins, install the Extensibility Development Kit (EDK).

Note: Before installing the EDK, you must have the following:

- Latest version of the EDK ZIP archive from the Self Update console. (To access the Self Update console, from the Cloud Control console, select **Setup**, then **Extensibility**, and then **Self Update**.)
 - Java version 1.6.0_24 or later
 - Local system running Solaris, Linux, HP-UX, AIX, or Windows with New Technology File System (NTFS)
-

To install the EDK:

1. Download the EDK ZIP archive to your local system using one of the following options:
 - Enterprise Manager Cloud Control
 - a. Log in to Enterprise Manager Cloud Control.
 - b. From the **Setup** menu, select **Extensibility**, then select **Development Kit**. The Extensibility Development Kit (EDK) page appears.
 - c. Under Installing the EDK, select **Download the Extensibility Development Kit to your workstation**.
 - d. Save `12.1.0.2.0_edk_partner.zip` to your local system.
 - Enterprise Manager Command Line Utility (EM CLI)

Note: For information about setting up EM CLI, see [Section 13.5.1.2, "Setting Up the EM CLI Utility"](#).

Open a command prompt and run the following command:

```
emcli get_ext_dev_kit
```

This command downloads the EDK zip archive to the same directory from where you ran the command and does not require any parameters.

2. Set your `JAVA_HOME` environment variable and ensure that it is part of your `PATH`. For example:

```
setenv JAVA_HOME /usr/local/packages/j2sdk1.6.0_24
```

```
setenv PATH $JAVA_HOME/bin:$PATH
```

3. Unpackage the downloaded EDK ZIP archive to a directory on your local system. For example:

```
Unzip 12.1.0.2.0_edk_partner.zip
```

This command creates the following directories under the directory (*release_edk_partner*) where you unpackaged the EDK ZIP archive:

```
release_edk_partner
|
bin
doc
lib
oui
samples
README
```

For more information about the directory contents, see [Section 1.2.1, "Contents of the EDK"](#).

1.4 Designing the Plug-in

Before creating your plug-in, you must first determine what information needs to be collected to monitor and manage the target type. This involves:

- Identifying performance and configuration metrics that should be collected.
- Determining how often each metric should be collected. Oracle recommends that the collection frequency for any metric should not be less than once every five minutes.
- Based on customer-specific operational practices, specifying default warnings and critical thresholds on these metrics. Whenever a threshold is crossed, Enterprise Manager generates an alert, informing administrators of potential problems.

1.5 Creating a Basic Plug-in

Ideally, begin by creating a basic monitoring plug-in that includes the basic required metadata:

- The target type definition file, which defines:

- A required "Response" metric group, which includes a status metric and a performance metric.
- Credentials needed to authenticate with the target.

For more information, see [Chapter 3, "Creating Target Metadata Files"](#).

- A default collection file defining the frequency at which metrics and configuration data will be collected. For more information, see [Section 3.5, "Creating the Default Collection File"](#).
- Developing Oracle Business Intelligence Publisher (BI Publisher) reports to display collected target data.

For more information, see [Chapter 5, "Developing BI Publisher Reports"](#).

Once created, you will validate and package your plug-in. See [Chapter 13, "Validating, Packaging, and Deploying the Plug-in"](#) for instructions.

As the final step, deploy your plug-in to Enterprise Manager Cloud Control and begin testing to ensure that data for the Response metric is being returned.

See Also: The EDK provides a sample host plug-in. This contains examples of all features mentioned previously. For more information, see the README bundled with the sample plug-in.

1.6 Creating an Intermediate Plug-in

When you have created a basic plug-in, you might want to enhance the plug-in's capabilities by adding more complex functionality.

- Add more complex metrics.

For more information, see [Section 3.4.3, "Defining Advanced Metrics"](#).

- Provide the ability to collect configuration data for the target, which is used to create a "snapshot" of the target's configuration at a specific point in time.

For more information, see [Chapter 6, "Collecting Target Configuration Data"](#).

- Create a metadata-based metadata custom user interface, which will essentially add a custom target home page for displaying target metrics to Enterprise Manager.

For more information, see [Chapter 8, "Defining a Management User Interface"](#).

- Defining target associations, which can be used to create topology models of the targets managed by the plug-in.

For more information, see [Chapter 10, "Using Derived Associations"](#).

- Define a job type that executes specific tasks specific to the target type.

For more information, see [Chapter 7, "Adding Job Types"](#).

1.7 Creating an Advanced Plug-in

When you have successfully created and validated your basic or intermediate plug-in, you might want to enhance it with advanced features. Among the enhancements to consider:

- Adding the ability to automatically discover newly-added instances of the target type managed by the plug-in.

For more information, see [Chapter 11, "Defining Target Discovery"](#).

- Building a Flex-based custom management user interface. This page will be accessed with other target home pages through the Cloud Control console.

For more information, see [Chapter 8, "Defining a Management User Interface"](#).

- Define compliance standards and monitoring rules specific to the target type.

For more information, see [Chapter 12, "Adding Compliance Standards"](#).

Defining the Plug-in

This chapter contains the following sections:

- [Introduction to Defining the Plug-in](#)
- [Basic Plug-in Metadata](#)
- [Creating Plug-in Definition Files](#)
- [Creating the plugin.xml File](#)
- [Creating the plugin_registry.xml File](#)
- [Validating the Plug-in Definition Files](#)
- [Adding Log Viewer Support to Your Plug-in](#)

2.1 Introduction to Defining the Plug-in

As a plug-in developer, you are responsible for the following steps within the plug-in definition process:

1. Define the plug-in identifier (ID).
For more information, see [Section 2.2.1, "Defining the Plug-in ID"](#).
2. Define the plug-in version.
For more information, see [Section 2.2.2, "Defining the Plug-in Version"](#).
3. Create the plug-in definition files:
 - a. Create the plugin.xml file.
The plugin.xml file provides the metadata describing the plug-in.
For more information, see [Section 2.4, "Creating the plugin.xml File"](#).
 - b. Create the plugin_registry.xml file.
The plugin_registry.xml file provides the metadata required by the Management Agent to which the plug-in will be deployed.
For more information, see [Section 2.5, "Creating the plugin_registry.xml File"](#).
4. Package the plug-in definition files in the plug-in staging directory (*plugin_stage*):
 - `plugin_stage/plugin.xml`
 - `plugin_stage/agent/plugin_registry.xml`For more information, see [Chapter 13, "Validating, Packaging, and Deploying the Plug-in"](#).

5. Validate the plug-in definition files.

For more information, see [Section 2.6, "Validating the Plug-in Definition Files"](#)

2.2 Basic Plug-in Metadata

The most basic plug-in requires metadata for the plug-in itself, including basic information such as name and version that is used by Oracle Management Service and Management Agents; definition of a basic metric indicating whether the target being monitored is up or down; and specifying the frequency at which metric data should be collected.

2.2.1 Defining the Plug-in ID

Plug-ins are identified by a unique plug-in identifier (ID). The plug-in ID has three parts:

- Vendor ID (8 chars). For example: `acme`
- Product ID (8 chars). For example: `switch`
- Plug-in Tag (4 chars). For example: `xkey`.

Note: The Vendor ID, Product ID, and Plug-in Tag *must not* begin with a number or include any special characters. All these parts must contain alphanumeric characters only.

The Plug-in ID created from the previous example would be:

`acme.switch.xkey`

Note: ■The plug-in ID must be unique across Enterprise Manager.

- If you are developing a plug-in within Oracle, contact Enterprise Manager Release Management to get the plug-in ID.
-
-

2.2.2 Defining the Plug-in Version

Each plug-in must be assigned a version. The plug-in versioning syntax is as follows:

`a.b.c.d.e`

- `a.b` = The version of the Enterprise Manager Extensibility Development Kit (EDK) used for development (12.1, 12.2, and so on).
- `c` = This value is always 0
- `d` = The developer-assigned plug-in version. This value must be incremented with each plug-in release on the same Enterprise Manager Cloud Control release.
- `e` = for future use. The default value is 0.

Putting it all together, the first version of a plug-in created for Enterprise Manager Cloud Control is:

`12.1.0.1.0`

Note: As a best practice, you should update the plug-in version incrementally as you create and deploy each iteration of your plug-in. For example: 12.1.0.1.0, 12.1.0.2.0, and so on.

2.3 Creating Plug-in Definition Files

Two new metadata files are required for all plug-ins deployed to Enterprise Manager Cloud Control 12c.

- **plugin.xml**
This file is used during plug-in deployment. It contains properties that identify the plug-in, such as name and version, and declares the set of target types that will be added to Enterprise Manager Cloud Control.
- **plugin-registry.xml**
This file declares those components included in the plug-in that are to be deployed to the Management Agent.

2.4 Creating the plugin.xml File

The plugin.xml file provides the metadata describing the plug-in.

The following sections describe the required and some of the optional tags that you can include in the plugin.xml file. The `TargetTypeList` element contains a `TargetType` element, which in turn contains a reference to the target type metadata file location in the plug-in archive.

[Example 2-1](#) provides a sample plugin.xml for a plug-in with basic features:

Example 2-1 Sample plugin.xml

```
<?xml version = '1.0' encoding = 'UTF-8'?>
<Plugin xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
        xsi:schemaLocation="http://www.oracle.com/EnterpriseGridControl/plugin_
        metadata plugin_metadata.xsd"
        xmlns="http://www.oracle.com/EnterpriseGridControl/plugin_metadata">

  <PluginId vendorId="acme" productId="demo" pluginTag="hostsample"/>

  <PluginVersion value="12.1.0.1.0"/>

  <PluginOMSOSAruId value="2000">
  </PluginOMSOSAruId>

  <ShortDescription>Test plugin for the Acme Demo Plug-in.</ShortDescription>
  <Readme><![CDATA[Acme Demo plug-in]]></Readme>
  <AgentSideCompatibility>
    <Version>12.1.0.1.0</Version>
  </AgentSideCompatibility>

  <TargetTypeList>
    <TargetType name="acme_demo_hostsample.xml" isIncluded="TRUE">
      <VersionSupport>
        <SupportedVersion supportLevel="Basic" minVersion="9.2.0.1"
                           maxVersion="9.8.0.0.0"/>
        <NonSupportedVersion minVersion="9.1.0.0"/>
      </VersionSupport>
    </TargetType>
  </TargetTypeList>
```

```

    </TargetType>
</TargetTypeList>

<PluginDependencies>
  <DependentPlugin pluginDependencyType="RunTime">
    <DepPluginId vendorId="acme" productId="switch" pluginTag="type3" />
    <BaseVersion version="11.2.0.1.0" />
  </DependentPlugin>
  <DependentPlugin pluginDependencyType="RunTime">
    <DepPluginId vendorId="acme" productId="switch" pluginTag="type4" />
    <BaseVersion version="11.2.0.1.0"></BaseVersion>
  </DependentPlugin>
</PluginDependencies>
<PluginAttributes Type="MP" Category="Others" />
</Plugin>

```

2.4.1 Overview of plugin.xml Elements

Table 2–1 describes the key elements included within the plugin.xml files.

Table 2–1 Key Elements Within the plugin.xml File

Element	Required Y/N	Description
Plugin	Y	The root element for the file.
PluginID	Y	The unique identifier assigned to the plug-in. For more information, see Section 2.2.1, "Defining the Plug-in ID" .
PluginVersion	Y	The plug-in version. For more information, see Section 2.2.2, "Defining the Plug-in Version" .
PluginOMSOSAruId	Y	The operating system (OS) ID for the Oracle Management Service to which the plug-in will be deployed. Usually, this element is set to 2000 (generic). For more information, see Section 2.4.2, "Certifying Plug-ins" .
Readme	Y	Provides the name of the plug-in that is displayed on the Plug-ins page of the Cloud Control console. To access the Plug-ins page, from the Setup menu, select Extensibility , then Self Update , and then Plug-ins .

Table 2–1 (Cont.) Key Elements Within the plugin.xml File

Element	Required Y/N	Description
PluginAttributes	N	<p>Defines plug-in attributes such as plug-in type, display name, category, and so on.</p> <p>The default Plug-in Type for metadata plug-ins is “MP”. The default Category is “Others”.</p> <p>Valid Type values:</p> <ul style="list-style-type: none"> ■ MP Metadata plug-in with default UI ■ MPP Metadata plug-in with custom UI <p>Valid Category values:</p> <ul style="list-style-type: none"> ■ Applications ■ Databases ■ Middleware ■ Cloud ■ Engineered Systems ■ Servers, Storage and Network ■ Others
TargetTypeList	N	<p>Contains one or more TargetType elements, each specifying the path and file name for a target type metadata file packaged with the plug-in.</p> <p>For information about the target type metadata file, see Section 3.3, “Creating the Target Type Metadata File”.</p> <p>Each TargetType element can also include a VersionSupport element identifying supported or non-supported versions of the target type.</p>
PluginDependencies	N	<p>Describes any dependencies that exist for the plug-in. Dependencies can be described as the following:</p> <ul style="list-style-type: none"> ■ CompileTime: This dependency indicates that the dependent plug-in should exist prior to deployment of current plug-in. ■ RunTime: This dependency indicates the feature dependencies and deployment of current plug-in can go ahead without the dependent plug-in. Along with dependencies one can also describe the prerequisites. Currently supported prerequisite is of type bug.
AgentSideCompatibility	N	<p>Identifies the Management Agent version with which the plug-in can communicate.</p>

2.4.2 Certifying Plug-ins

Note: All metadata plug-ins must be generic on the OMS side and are implicitly certified on all platforms. However, the plug-in can specify the OS certification for the Management Agent.

Because Enterprise Manager is released on a number of OS platforms, you must consider how your plug-in will behave on different OS platforms. The plugin.xml file contains elements and attributes that support a certification mechanism.

In cases, where the plug-in is applicable to only a subset of OS platforms, you can use the tags defined in [Table 2–2](#). If you do not specify any information in the <Certification> section, the plug-in is assumed certified on all platforms.

Table 2–2 Certification Tags

Tag	Description
Component type	Specifies the plug-in component. Valid values: <ul style="list-style-type: none"> ■ Agent: Management Agent component ■ Discovery: Discovery component
PortARUID value	Specifies the ARU ID for the OS or platform. Valid values: <ul style="list-style-type: none"> ■ 46: Linux x86 (32-bit) ■ 212: AIX 5L and 6.1 (64-bit) ■ 226: Linux x86-64 (64-bit) ■ 23: Solaris Sparc (64-bit) ■ 267: Solaris x86-64 (64-bit) ■ 233: Microsoft Windows x86-64 (64-bit)

[Example 2–2](#) indicates that the plug-in is designed to work on Linux 32 and Linux 64 platforms only. If you do not specify a certified port, then by default your plug-in is certified on all operating systems and platforms. But if you specify at least one PortARUID element, then that component is certified on those specified platforms *only*.

Note: The Management Agent and Discovery components must have the same value

Example 2–2 Certifying Generic Plug-ins

```
<PluginOMSOSAruId value="2000">
</PluginOMSOSAruId>

<Certification>
  <Component type="Agent">
    <CertifiedPorts>
      <PortARUID value="46" />
      <PortARUID value="226" />
    </CertifiedPorts>
  </Component>
  <Component type="Discovery">
    <CertifiedPorts>
      <PortARUID value="46" />
      <PortARUID value="226" />
    </CertifiedPorts>
  </Component>
</Certification>
```

2.5 Creating the plugin_registry.xml File

The plugin_registry.xml file provides the metadata required by the Management Agent that the plug-in will be deployed to. It is packaged in the /agent directory within the plug-in archive and is deployed to the Management Agent that will monitor a target.

[Example 2-3](#) provides a sample plugin_registry.xml file. The TargetTypes element contains a reference to the target type metadata file location in the plug-in archive. The location is relative to the *plugin_stage* directory root, that is, starting from the Management Agent subdirectory or the same location where the plugin_registry.xml file is located.

Similarly, the TargetCollections element contains a reference to the plug-in's default collection metadata file, which is also packaged with the plug-in.

Example 2-3 Sample plugin_registry.xml File

```
<?xml version="1.0"?>
<PlugIn ID="acme.demo.hostsample" Description="Demo Sample Host Plugin"
Version="12.1.0.1.0" HotPluggable="false"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.oracle.com/EnterpriseGridControl/plugin
plugin.xsd">
  <TargetTypes>
    <FileLocation>metadata/acme_switch_key.xml</FileLocation>
  </TargetTypes>
  <TargetCollections>
    <FileLocation>default_collection/acme_switch_key_
collection.xml</FileLocation>
  </TargetCollections>
  <PlugInLibrary>
    <FileLocation>archives/em-as-fmw-fetchlet.jar</FileLocation>
    <FetchletRegistration>
      <Fetchlet ID="DMS" ExecutionClass="oracle.sysman.as.fetchlets.DMSFetchlet"
Version="" Description="" Adapter=""/>
    </FetchletRegistration>
    <AdditionalClassPath>
      <FileLocation>archives/dms.jar</FileLocation>
    </AdditionalClassPath>
  </PlugInLibrary>
</PlugIn>
```

2.5.1 Overview of plugin_registry.xml Elements

[Table 2-3](#) describes the key elements included within the file.

Table 2–3 Key Elements Within the *plugin_registry.xml* File

Element	Required (Y/N)	Description
Plugin	Y	<p>The root element for the file. It includes the following attributes:</p> <ul style="list-style-type: none"> ■ ID: Required. The unique identifier assigned to the plug-in. For more information, see Section 2.2.1, "Defining the Plug-in ID". ■ Description: Optional. A title describing the plug-in. ■ Version: Required. The plug-in version. For more information, see Section 2.2.2, "Defining the Plug-in Version".
TargetTypes	N	<p>Contains one or more <code>FileLocation</code> elements, each specifying the path and file name for a target type metadata file packaged with the plug-in.</p> <p>For information about target type files, see Section 3.3, "Creating the Target Type Metadata File".</p>
TargetCollections	N	<p>Contains one or more <code>FileLocation</code> elements, each specifying the default collection for a target type.</p> <p>For information about this file, see Section 3.5, "Creating the Default Collection File".</p>
PlugInLibrary	N	<p>Lists the different types of artifacts (fetchlets, receivelets, and so on) packaged in the plug-in.</p> <p>The <code>PlugInLibrary</code> element includes the following subelements:</p> <ul style="list-style-type: none"> ■ <code>FileLocation</code>: Mandatory. Defines the location of the JAR containing the implementations of the following listed fetchlets. ■ <code>FetchletRegistration</code>: Optional. Creates an entry that maps a fetchlet id (DMS in Example 2–3) to the class that contains the implementation of the fetchlet interface. ■ <code>ReceiveletRegistration</code>: Optional. Creates an entry that maps a receivelet id to the class that contains the implementation of the receivelet interface. ■ <code>AdditionalClassPath</code>: Optional. Specifies additional JAR files to be loaded by the plug-in for a specific library.
AdditionalClassPath	N	Specifies additional JAR files to be loaded by the plug-in that are shared by all the plug-in libraries

2.6 Validating the Plug-in Definition Files

To verify that your `plugin.xml` and `plugin_registry.xml` files are defined correctly, enter the following command from the `/bin` directory of the EDK:

```
empdk validate_plugin -stage_dir plugin_stage
```

In the preceding command, `plugin_stage` represents the plug-in staging directory.

For information about the EDK, see [Section 1.2, "About the Extensibility Development Kit \(EDK\)"](#) and for more information about the `empdk` command and its usage, see [Section 13.3, "Validating the Plug-in"](#).

2.7 Adding Log Viewer Support to Your Plug-in

Beginning with Enterprise Manager Cloud Control Release 2 (12.1.0.3), it is possible to enable log files for your deployed plug-in to be viewable with Cloud Control's Log Viewer. This component is accessed by selecting **Monitoring**, then **Logs**, from the **Enterprise** menu in Cloud Control.

Follow these steps to enable this feature:

1. Create the log viewer registration XML file for your plug-in. The DTD for this XML file is:
 `oracle/sysman/emSDK/logmgmt/registration/LogMgmtTargetTypeRegistration.xsd`
2. Package this file in `oms/metadata/logmgmt/` within the plug-in directory structure.

An example of a log viewer registration file is shown below.

Example 2–4 Sample Log Viewer Registration File

```
<LogMgmtUITargetConfig TARGET_TYPE="%your target type%">
  <LogViewerImpl CLASS_NAME="oracle.sysman.emas.model.logmgmt.MASLogViewer"/>
  <VersionProperties VALID_VERSIONS="11" MIN_META_VER="11.00000"VERSION_
    CATEGORY_PROP_WILDCARD_CHAR="*" />
</LogMgmtUITargetConfig>
```

Creating Target Metadata Files

This chapter contains the following sections:

- [Introduction to Creating Target Metadata Files](#)
- [Overview of Target Definition Files](#)
- [Creating the Target Type Metadata File](#)
- [Defining Metrics to Collect from the Target](#)
- [Creating the Default Collection File](#)
- [Guidelines for Creating Target Metadata](#)
- [Testing Your Target Type Definitions](#)
- [Validating Your Metadata XML](#)
- [Troubleshooting the Target Creation Process](#)

3.1 Introduction to Creating Target Metadata Files

As a plug-in developer, you are responsible for the following steps within the target metadata files creation process:

1. Create the target definition file.

The target type metadata file tells the Management Agent what data to retrieve and how to obtain that data for this particular target type.

For more information, see [Section 3.3, "Creating the Target Type Metadata File"](#).

2. Define metrics to collect from the target.

A metric refers to a specific piece of data collected from the target. A set of related metrics collectively comprise a metric group.

For more information, see [Section 3.4, "Defining Metrics to Collect from the Target"](#).

3. Define target configuration data to collect.

You can collect configuration data for a target and save it in the Management Repository as a *snapshot* representing the target's configuration at a specific point in time. Each configuration snapshot is associated with a specific target instance.

For more information, see [Chapter 6, "Collecting Target Configuration Data"](#).

4. Create the default collection file.

The default collection file defines the metric data to be collected from targets and written to the Management Repository along with information such as the collection frequency.

For more information, see [Section 3.5, "Creating the Default Collection File"](#).

5. Package the various definition files in the plug-in staging directory (*plugin_stage*):

- Target type metadata file

- *plugin_stage/oms/metadata/targetType/target_type.xml*
- *plugin_stage/agent/metadata/target_type.xml*

Note: An identical copy of this file must be placed in both the /oms and /agent directories.

- Default collection file

- *plugin_stage/oms/metadata/default_collection/target_type.xml*
- *plugin_stage/agent/default_collection/target_type.xml*

Note: An identical copy of this file must be placed in both the /oms and /agent directories.

- Configuration metadata file

plugin_stage/oms/metadata/snapshotlive/target-type_ecmdef.xml

For more information, see [Chapter 13, "Validating, Packaging, and Deploying the Plug-in"](#).

3.2 Overview of Target Definition Files

Two XML metadata files are required to define the target type that your plug-in will enable Enterprise Manager Cloud Control to monitor and manage:

- Target type metadata file

The definition of a target type primarily consists of the metrics you want the Management Agent to collect for the target. The file contains a list of all metrics that will be collected for the target type, along with specifics on how to compute each metric.

For more information, see [Section 3.3, "Creating the Target Type Metadata File"](#).

- Default collection file

This file defines the interval at which metric data will be collected or received from the target. It also specifies alert thresholds and optional corresponding warning messages for each metric. Cloud Control users can override the default collection intervals, but the default values must be provided in this file.

For more information, see [Section 3.5, "Creating the Default Collection File"](#).

This chapter also describes the metadata definitions required to collect configuration data for plug-in targets. This is an advanced feature but can be useful for many

plug-ins. For more information, see [Section 6.2, "About the Configuration Definition Files"](#).

The following sections provide the summary of creating the target type and default collection metadata files, as well as an overview of target configuration data collection.

3.3 Creating the Target Type Metadata File

The target type metadata file tells the Oracle Management Agent what data to retrieve and how to obtain that data for this particular target type.

At the highest definition level, the target type metadata file is composed of four key XML elements as described in [Table 3–1](#).

Table 3–1 Key Elements of the Target Type Metadata File

Element	Description
TargetMetadata	Specifies information about the plug-in, such as name and version.
Metric	Defines a metric group, which in turn contains one or more metrics that each define a specific piece of data collected from the target.
InstanceProperties	Defines properties that are populated when a target instance is created.
CredentialTypes/CredentialSets	Specifies credentials required to by the plug-in authenticate with a target instance.

Enterprise Manager ships with predefined target type metadata files that cover the most common target types. In situations where the predefined target metadata files do not fit the types of targets you want to monitor, you can either:

- Define a new target type by creating a target type metadata file
- Use one of the predefined metadata files as a template for defining a new target type, and then repack the files as a new plug-in

This section briefly introduces the structure of the target type metadata file. A complete example of a target type metadata file is provided with the EDK:

`edk/samples/plugins/SampleHost/oms/metadata/targetType/demo_hostsample.xml`

In the preceding directory path, *edk* represents the directory where you expanded the EDK archive. For information about the EDK archive, see [Section 1.2, "About the Extensibility Development Kit \(EDK\)"](#).

For additional information about creating target type metadata files, [Section 3.6, "Guidelines for Creating Target Metadata"](#).

3.3.1 Creating a Basic Target Type Metadata File

[Example 3–1](#) shows the minimum required information that a target type file must contain.

Example 3–1 Target Type File

```
<TargetMetadata META_VER="2.0" TYPE="demo_hostsample">
  <Display>
    <Label NLSID="hs_displayname">Demo Plugin Sample Host</Label>
  </Display>
```

```

<Metric NAME="Response" TYPE="TABLE">
  <Display>
    <Label NLSID="hs_response_displayname">Response</Label>
  </Display>
  <TableDescriptor>
    <ColumnDescriptor NAME="Status" TYPE="NUMBER" IS_KEY="FALSE">
      <Display>
        <Label NLSID="hs_response_status">Status (up/down)</Label>
      </Display>
    </ColumnDescriptor>
  </TableDescriptor>
  <QueryDescriptor FETCHLET_ID="OSLineToken">
    <Property NAME="scriptsDir" SCOPE="SYSTEMGLOBAL">scriptsDir</Property>
    <Property NAME="fake" SCOPE="INSTANCE"
      OPTIONAL="TRUE">USE_FAKE_DATA</Property>
    <Property NAME="perlBin" SCOPE="SYSTEMGLOBAL">perlBin</Property>
    <Property NAME="script" SCOPE="GLOBAL">%scriptsDir%/emx/demo_
hostsamle/datacollector.pl --collect Response --fake %fake%</Property>
    <Property NAME="startsWith" SCOPE="GLOBAL">em_result=</Property>
    <Property NAME="delimiter" SCOPE="GLOBAL">|</Property>
  </QueryDescriptor>
</Metric>
<InstanceProperties>
  <InstanceProperty NAME="SAMPLE_DATA" CREDENTIAL="FALSE" OPTIONAL="TRUE">
    <Display>
      <Label NLSID="EMPLOYEE_ID_iprop">Employee ID</Label>
    </Display>
  </InstanceProperty>
</InstanceProperties>
</TargetMetadata>

```

The following sections provide information about the XML definitions shown in [Example 3-1](#).

3.3.2 Naming the Target Type Metadata File

Oracle recommends that users adding new target types adhere to Enterprise Manager naming conventions. This naming convention allows for file naming consistency in environments where similar products from multiple vendors are used. The target naming convention follows the form *vendorID_productID_PluginTag*.

For example:

```
acme_demo_targetType.xml
```

3.3.3 Defining the Target Type Metadata

The first lines after the header of the target definition file identify the target type. The following excerpt defines the metadata version (META_VER="2.0") and target type (TYPE="acme_demo_targetType").

```
<TargetMetadata META_VER="2.0" TYPE="acme_demo_targetType">
```

Metadata versioning allows different versions of the same target type metadata to exist concurrently within the managed environment, although only one metadata version is allowed per Management Agent. You should update the metadata version each time you update the target metadata file.

The syntax for the meta value is MajorNumber.MinorNumber. The value for MinorNumber can be either one or two digits. It is important that you choose to use either one or two digits for MinorNumber, and continue to use that syntax throughout the plug-in's lifecycle.

A one-digit MinorNumber can be updated up to nine times, after which the MajorNumber value must be increased. For example:

```
1.0 1.1 1.2 1.9 2.0 2.1
```

A two-digit MinorNumber can be revised up to 99 times. For example:

```
1.01 1.02 1.09 1.10 1.11. 1.99 2.01
```

Note that you cannot combine single-digit and two-digit formats. For example, increasing the version from 1.9 to 1.10 is not valid.

If you modify the metadata for one or more metrics changed during plug-in development, be sure to increment the value of the META_VER attribute of the TargetMetadata element to the next digit value. For example, if the current version is "1.0", set the value to "1.1" if you modify the metric metadata:

```
<TargetMetadata META_VER="1.1" TYPE="demo_hostsample">
```

Continue to increment the version each time you modify your metrics. Otherwise, NullPointerException errors may occur if the plug-in is deployed to the same development/test installation that the previous version was deployed to.

Once plug-in development is complete, the META_VER value can safely be set to the actual production version.

3.3.4 Defining Target Credentials

In most cases, the plug-in will be required to authenticate with each target instance that it will collect data for, or run jobs against. Credential types and credential sets needed to enable authentication are defined in the CredentialInfo element within the target type metadata file.

For more information, see [Chapter 15, "Defining Credentials"](#).

Credentials information for the target includes and defines the credentials fields (referred to as columns) and the credentials sets specific to the target type. Enterprise Manager's security framework provides facilities for managing these credentials and using them when performing various management functions.

3.3.5 Defining Type Properties

The extensibility framework uses type properties to internally categorize the target type for framework processing. They are not visible to the end user. Corresponding subsystems use the type properties to enable features for the target type or to perform appropriate validation checks.

The value set at the type property level applies to all targets of that type and across all metaversions, unlike instance properties which only apply to a specific target.

The following example specifies that the target type is a system class of target. The extensibility framework uses this setting to display the target on all system pages.

```
<TypeProperties>
  <TypeProperty PROPERTY_NAME="is_system" PROPERTY_VALUE="1"/>
</TypeProperties>
```

Table 3–2 provides a description of the available type properties.

Table 3–2 Type Properties

Property Name	Description
is_system	<p>Specifies the type as modelling a system type. You must set this value for all system types.</p> <ul style="list-style-type: none"> is_cluster: Specifies the type as modelling a cluster. Clusters are subsets of systems. is_end_user: Set for systems constructed from user-chosen entities; these are subsets of systems. <p>Note: The property value is always set to 1 for all the is_name properties.</p>
is_service	<p>Specifies the type as modelling a service.</p> <p>Note: The property value is always set to 1 for all the is_name properties.</p>
is_aggregate	Enterprise Manager sets this value automatically. Do not modify.
is_group	Do not use
is_composite	Do not use
is_install	Set this value for an install home manageable entity (for example, Oracle home)
is_existence	<p>Set this value for a discovered entity with an existence-only state, that is, an entity that is discovered but cannot be managed by Oracle yet.</p> <p>Possible value:</p> <ul style="list-style-type: none"> 1: Indicates a discovered entity with an existence-only state <p>Note: When the entity becomes a managed entity by Oracle, you must remove this entry from the target type metadata file and register the target type again.</p>
priv_propagation_mode	<p>This property is used for privilege propagation and specifies the privilege propagation mode.</p> <p>Possible values:</p> <ul style="list-style-type: none"> 0: No privilege propagation 1: Privilege propagation at instance level 2: All targets are privilege propagating
disallow_redundancy_group	<p>Used by the redundancy group feature to disable redundancy groups for certain target types (which have disallow redundancy group set). Specifies whether redundancy group can contain this type as a member</p> <p>Possible value:</p> <ul style="list-style-type: none"> 1: Do not allow redundancy
member_target_type	Used by the redundancy group feature to lock the target type to the specified member_target_type.
TargetVersion	Specifies the name of the instance or dynamic property that represents the target version for the target type (for all target pages and plug-in certification)

Example 3–2 Defining Type Properties

```
<TargetMetadata META_VER="1.1" TYPE="oracle_dbsys" CATEGORY_PROPERTIES=""
RESOURCE_BUNDLE_PACKAGE="oracle.sysman.db.rsc">
```



```

<Display>
  <Label NLSID="oracle_dbsys_nlsid">Database System</Label>
</Display>

<TypeProperties>
  <TypeProperty PROPERTY_NAME="is_system" PROPERTY_VALUE="1"/>
  <TypeProperty PROPERTY_NAME="priv_propagation_mode" PROPERTY_VALUE="2"/>
</TypeProperties>

  <MonitoringMode MEDIATOR="Repository"/>
</TargetMetadata>

```

3.3.6 Defining Instance Properties

Instance properties are populated when a target instance is created. The `InstanceProperties` descriptor within the target type metadata file defines what properties an administrator must specify in the Enterprise Manager Cloud Control console when adding a new target instance of this particular target type.

Although the `InstanceProperties` section can be defined at various locations within the target type metadata file, Oracle recommends defining this section at the very end of the file for consistency. Instance properties defined in the target type metadata file are resolved to values specified for these instance properties in the target type metadata file.

Target instance properties are named values that can be used for computing the metrics of the target, or for display in the home page of the target. The list of target instance properties is specified in the metadata to allow data driven user interfaces to register targets, and for the Oracle Management Agent to validate that a target instance is complete.

3.3.6.1 Static Instance Properties

Instance properties are populated when a target instance is created. In this example, the property is required (`OPTIONAL="FALSE"`) and it is a credential property.

```

<InstanceProperties>
  <InstanceProperty NAME="password" OPTIONAL="FALSE" CREDENTIAL="TRUE">
    <Display>
      <Label NLSID="USER_PASSWORD">User Password</Label>
    </Display>
  </InstanceProperty>
</InstanceProperties>

```

3.3.6.2 Dynamic Instance Properties

The values for dynamic instance properties are passed back by the Management Agent collecting data from the target instance. They are typically used within a `QueryDescriptor` to define properties passed to the fetchlet responsible for metric collection. For more information about the `QueryDescriptor` element, see [Table 3-3](#) on page 3-13. For more information about fetchlets, see [Chapter 21, "Using Fetchlets"](#).

The properties in the following example are described in [Section 3.4.2, "Defining the Basic Response Metric Group"](#).

```

<InstanceProperties>
  <DynamicProperties NAME="AruidInfo" FORMAT="ROW" OPT_PROP_LIST="ARUID">
    <QueryDescriptor FETCHLET_ID="OSLineToken">
      <Property NAME="scriptsDir" SCOPE="SYSTEMGLOBAL">scriptsDir</Property>
      <Property NAME="perlBin" SCOPE="SYSTEMGLOBAL">perlBin</Property>
    </QueryDescriptor>
  </DynamicProperties>
</InstanceProperties>

```

```
<Property NAME="command" SCOPE="GLOBAL">%perlBin%/perl</Property>
<Property NAME="script" SCOPE="GLOBAL">%scriptsDir%/hostaruid.pl</Property>
<Property NAME="startsWith" SCOPE="GLOBAL">em_result=</Property>
<Property NAME="delimiter" SCOPE="GLOBAL">|</Property>
</QueryDescriptor>
</DynamicProperties>
</InstanceProperties>
```

Using dynamic properties reduces the work involved in configuring a target by allowing certain properties to be computed rather than requiring the user to correctly specify their values (for example, the "Version" property of a database can be reliably computed given addressing information).

Dynamic properties are computed in the order that they are defined in the XML file so that later dynamic properties can use values from earlier dynamic properties in the XML file if required.

The Management Agent allows for the fact that the target must be up for the successful computation of these dynamic properties by recomputing the properties each time a target restart is detected; that is, each time the target status changes to **Up**.

Note: Some properties can be computed without access to the target; therefore there is some support for computing dynamic properties when the target is down.

To compute a dynamic property when the target is down, include the following attribute:

```
COMPUTE_WHEN_DOWN="TRUE"
```

3.4 Defining Metrics to Collect from the Target

Metrics are at the core of Cloud Control's target monitoring capabilities. When we speak of Cloud Control's ability to monitor and manage various targets, what we are really talking about is its ability to collect, process, and display target metrics.

A *metric* refers to a specific piece of data collected from the target.

Note: Existing metric definitions defined in target-type metadata files are still valid in this release.

Metrics can be viewed as being of two basic types:

- Pull metrics

In this model, the plug-in polls the target for metric data at the frequency specified in the collections file. This is the most common type of metric utilized by plug-ins.

A metric of this type requires the use of a *fetchlet*, a parametrized data access mechanism that takes arguments (for example, a script, a SQL statement, a target instance's properties) as input and returns formatted data.

Predefined fetchlets are provided by Oracle for use with plug-ins. For a list of available fetchlets and information about their usage, see [Chapter 21, "Using Fetchlets"](#).

- Push metrics

In this model, the plug-in receives notifications that are sent asynchronously from the target, without being requested. This type of metric requires a *receivelet*, which enables the plug-in to receive such notifications. As with fetchlets, predefined receivelets are provided by Oracle. For information about using receivelets, see [Chapter 20, "Using Receivelets"](#).

3.4.1 Metric Definition Files

The target metadata must define each type of metric the plug-in will collect, how and when the metric data is to be collected, and what metric thresholds will trigger an incident to be raised within Cloud Control.

The metadata for metric groups and individual metrics is defined in two metadata files packaged with the plug-in:

- The target type metadata file

The content of the target type metadata file actually consists primarily of metric definitions. The fetchlet or receivelet that a metric will use is defined in the `QueryDescriptor` or `PushDescriptor` element within the target type metadata file.

For information about key metric definition elements, see [Section 3.4.4, "Overview of Key Metric Metadata Elements"](#).

- The default collection metadata file

The frequency at which data is collected for each metric is defined in the default collection metadata file. Metric Alert event conditions for each metric and the messages to display for such alerts are also defined in this file.

For information about the default collection metadata file, see [Section 3.5, "Creating the Default Collection File"](#).

3.4.2 Defining the Basic Response Metric Group

As a matter of practice, Oracle recommends that you specify at least a single Response metric group that includes the following metrics for each target type:

- A Status metric that indicates target availability (required for all target types)
- A metric that reports target performance (optional but recommended)

The corresponding default collection file must define a critical condition on the Status metric that represents the target status as up or down. For more information, see [Section 3.5.2, "Defining Basic Metric Collection"](#).

[Example 3–3](#) defines a Status metric. The return value of Status is as follows:

- 0: Target status is down
- 1: Target status is up

The process by which the metric data is collected is defined in the `QueryDescriptor` element. This descriptor specifies that the `OSLineToken` fetchlet invokes a Perl script (`data_collector.pl`) to collect the data. The Perl script returns a standard out (stdout) data stream containing the collected data to the fetchlet.

Each property passed to the `OSLineToken` fetchlet execution is specified in a `Property` tag within the `QueryDescriptor` element.

- The `OSLineToken` fetchlet requires that a `GLOBAL` property called `command` be set to the command that is to be executed. Different tokens typically have specific

required properties. For more information about the OSLineToken fetchlet, see [Section 21.2, "OS Command Fetchlets"](#).

- When a plug-in is deployed to a Management Agent, any scripts or binaries associated that were packaged within the /agent/scripts directory in the plug-in archive are written to the following directory in the Management Agent, where AGENT_HOME is the Management Agent plug-in home directory and *plugin_name* is the name of the plug-in:

```
AGENT_HOME/plugins/plugin_name
```

The `scriptsDir` property is a token that defines this location.

Note: In previous releases, the scripts bundled with the plug-in were copied to the scripts directory under the Management Agent. However, for this release, the scripts included in the plug-in are used directly.

This changes the behavior of the `scriptsDir` property in the `QueryDescriptor` element. Previously, it referred to the directory under the Management Agent, but now it refers to the directory under the plug-in.

If you want to refer to the scripts directory under the Management Agent, use the `sdkScriptsDir` property.

- The `script` property specifies the script (`data_collector.pl`) to be run.

The EDK provides an example of this script:

```
edk/samples/plugins/HostSample/demo_hostsample/stage/agent/scripts/emx/demo_hostsample
```

- The `startsWith` and `delimiter` properties specify the format of the STDOUT of the script executed. In this case, the script will return a single row that looks like this:

```
em_result="value for Load|value for Status"
```

Example 3-3 Response Metric

```
<Metric NAME="Response" TYPE="TABLE">
  <TableDescriptor>
    <ColumnDescriptor NAME="Status" TYPE="NUMBER" IS_KEY="FALSE">
      <Display>
        <Label NLSID="oracle_emrep_resp_status">Status</Label>
      </Display>
    </ColumnDescriptor>
  </TableDescriptor>
  <QueryDescriptor FETCHLET_ID="OSLineToken">
    <Property NAME="perlBin" SCOPE="SYSTEMGLOBAL">perlBin</Property>
    <Property NAME="scriptsDir" SCOPE="SYSTEMGLOBAL">scriptsDir</Property>
    <Property NAME="command" SCOPE="GLOBAL">%perlBin%/perl</Property>
    <Property NAME="script" SCOPE="GLOBAL">%scriptsDir%/emrepresp.pl</Property>
    <Property NAME="startsWith" SCOPE="GLOBAL">em_result</Property>
    <Property NAME="delimiter" SCOPE="GLOBAL">|</Property>
    <Property NAME="ENVEM_TARGET_NAME" SCOPE="INSTANCE">NAME</Property>
    <Property NAME="ENVEM_REPOS_USER" SCOPE="INSTANCE">UserName</Property>
    <Property NAME="ENVEM_REPOS_PWD" SCOPE="INSTANCE">password</Property>
    <Property NAME="ENVHOST" SCOPE="INSTANCE" OPTIONAL="TRUE">MachineName</Property>
```

```

<Property NAME="ENVPORT" SCOPE="INSTANCE" OPTIONAL="TRUE">Port</Property>
<Property NAME="ENVSID" SCOPE="INSTANCE" OPTIONAL="TRUE">SID</Property>
<Property NAME="ENVCONNECTDESCRIPTOR" SCOPE="INSTANCE"
OPTIONAL="TRUE">ConnectDescriptor</Property>
<Property NAME="ENVEM_TARGET_ADDRESS" SCOPE="GLOBAL">
  (DESCRIPTION=(ADDRESS=(PROTOCOL=TCP) (HOST=%ENVHOST%) (Port=%ENVPORT%))
  (CONNECT_DATA=(SID=%ENVSID%)))</Property>
</QueryDescriptor>
</Metric>

```

For a description of these elements, see [Section 3.4.4, "Overview of Key Metric Metadata Elements"](#).

3.4.3 Defining Advanced Metrics

You can define much more complex metrics, such as metrics that collect CPU performance data or metrics for which values are computed from the values of other metrics. Examples of such advanced or complex metric definitions can be seen in the sample target type metadata file provided with the EDK:

```

edk/samples/plugins/HostSample/demo_hostsample/stage/oms/metadata/targetType/demo_
hostsample.xml

```

[Example 3-4](#) shows a metric group containing metrics that collect CPU performance data. As in the previous example, the `QueryDescriptor` element specifies that the `OSLineToken` fetchlet will invoke the `data_collector.pl` script to collect the data.

Example 3-4 Defining Advanced Metrics

```

<Metric NAME="CPUProcessesPerf" TYPE="TABLE">
  <Display>
    <Label NLSID="hs_cpuprocessesperf_displayname">Host Process CPU Performance
      Statistics</Label>
    </Display>
    <TableDescriptor>
      <ColumnDescriptor NAME="ProcPID" TYPE="NUMBER" IS_KEY="TRUE">
        <Display>
          <Label NLSID="hs_cpuprocessesperf_procpid">PID</Label>
        </Display>
      </ColumnDescriptor>
      <ColumnDescriptor NAME="ProcUser" TYPE="STRING" IS_KEY="FALSE">
        <Display>
          <Label NLSID="hs_cpuprocessesperf_procuser">User</Label>
        </Display>
      </ColumnDescriptor>
      <ColumnDescriptor NAME="ProcCPU" TYPE="NUMBER" IS_KEY="FALSE">
        <Display>
          <Label NLSID="hs_cpuprocessesperf_proccpu">CPU Usage (%)</Label>
        </Display>
      </ColumnDescriptor>
      <ColumnDescriptor NAME="ProcCmd" TYPE="STRING" IS_KEY="FALSE">
        <Display>
          <Label NLSID="hs_cpuprocessesperf_proccmd">Command</Label>
        </Display>
      </ColumnDescriptor>
    </TableDescriptor>
    <QueryDescriptor FETCHLET_ID="OSLineToken">
      <Property NAME="scriptsDir" SCOPE="SYSTEMGLOBAL">scriptsDir</Property>
      <Property NAME="perlBin" SCOPE="SYSTEMGLOBAL">perlBin</Property>
      <Property NAME="command" SCOPE="GLOBAL">%perlBin%/perl</Property>
    </QueryDescriptor>
  </Metric>

```

```

    <Property NAME="script" SCOPE="GLOBAL">%scriptsDir%/emx/demo_hostsample/data_
        collector.pl</Property>
    <Property NAME="startsWith" SCOPE="GLOBAL">em_result=</Property>
    <Property NAME="delimiter" SCOPE="GLOBAL">|</Property>
</QueryDescriptor>
</Metric>

```

[Example 3–5](#) shows a test metric. The Management Agent can check some metrics to determine if a target has been specified correctly with valid instance properties. Setting the IS_TEST_METRIC attribute to "TRUE" provides a Test button when adding a target instance.

Example 3–5 Defining a Test Metric

```

<Metric NAME="Ping" TYPE="TABLE" IS_TEST_METRIC="TRUE" USAGE_TYPE="HIDDEN">
  <Display>
    <Label NLSID="label_metrics_ping">Ping Test</Label>
  </Display>
  <TableDescriptor>
    <ColumnDescriptor NAME="tcpIpPing" TYPE="NUMBER">
      <Display>
        <Label NLSID="test_ping">TCP Ping, Milliseconds</Label>
      </Display>
    </ColumnDescriptor>
  </TableDescriptor>
  <QueryDescriptor FETCHLET_ID="OSLineToken">
    <Property NAME="perlBin" SCOPE="SYSTEMGLOBAL">perlBin</Property>
    <Property NAME="scriptsDir" SCOPE="SYSTEMGLOBAL">scriptsDir</Property>
    <Property NAME="command" SCOPE="GLOBAL">%perlBin%/perl
      %sdkScriptsDir%/osresp.pl</Property>
    <Property NAME="ENVEM_TARGET_NAME" SCOPE="INSTANCE">hostName</Property>
    <Property NAME="startsWith" SCOPE="GLOBAL">em_result=</Property>
    <Property NAME="delimiter" SCOPE="GLOBAL">|</Property>
  </QueryDescriptor>
</Metric>

```

3.4.4 Overview of Key Metric Metadata Elements

[Table 3–3](#) describes the key elements that define metrics. For additional information about defining metrics, see [Section 3.6, "Guidelines for Creating Target Metadata"](#).

Table 3–3 Key Elements Used to Define Metrics

Element	Description
Metric	<p>Required. Defines a metric group containing one or more metrics, each defined in a ColumnDescriptor. The Metric element includes the following attributes:</p> <ul style="list-style-type: none"> NAME: The name of the metric group, up to 64 characters long. TYPE: Valid values are "TABLE" or "RAW". Typically set to "TABLE", indicating tabular data. IS_TEST_METRIC: The Management Agent can check some metrics to determine if a target has been specified correctly with valid instance properties. This attribute marks this metric as one of the test metrics. By default, the value is set to "FALSE". Setting this value to "TRUE" provides a Test button when adding a target instance. USAGE_TYPE: The metric is not viewable and the Management Agent will not upload data collected by the metric to the Management Repository.
TableDescriptor	<p>Required when the Metric TYPE attribute is set to "TABLE". It contains one or more ColumnDescriptor elements, each defining a metric to collect</p>
ColumnDescriptor	<p>Defines a single metric to be collected. It contains the following attributes:</p> <ul style="list-style-type: none"> NAME: The name of the metric, up to 64 characters long. TYPE: Describes how the metric data will be stored in the Management Repository. The value is typically either "NUMBER" or "STRING". IS_KEY: Indicates if the metric is to be treated as a primary key column in the Management Repository. Values are "TRUE" or "FALSE" (default).

Table 3–3 (Cont.) Key Elements Used to Define Metrics

Element	Description
QueryDescriptor	<p>Defines a command to run, which returns the set of data described in the TableDescriptor. The element contains one or more Property elements, each defining a property to pass in with the command invocation.</p> <p>Note: You can refer to earlier defined properties using the <code>%property_name%</code> format. For example:</p> <pre><Property NAME="perlBin" SCOPE="SYSTEMGLOBAL">perlBin</Property> <Property NAME="command" SCOPE="GLOBAL">%perlBin%/perl</Property></pre> <p>The element includes a <code>FETCHLET_ID</code> attribute that identifies the fetchlet mechanism that will be used to process the request. Properties required for the fetchlet invocation are specified in one or more Property elements within the QueryDescriptor.</p> <p>The following are among the fetchlets commonly used by plug-in developers:</p> <ul style="list-style-type: none"> ■ OS: Executes the given operating system command or script and returns the command's output in a single cell table. ■ OSLines: Similar to the fetchlet, but returns the output tokenized by lines. ■ OSLineToken: Similar to the fetchlets, but the output is tokenized first by lines; each line is then tokenized by a given delimiter set. ■ HTTPDataLineToken: Invokes an HTTP request to the specified URL. ■ SQL: Executes the specified SQL script against the specified Oracle Database. ■ Snmp: Invokes the specified SNMP call to the specified SNMP agent. ■ WBEM: Invokes the specified WBEM call to the specified CIMOM object repository. <p>For a complete list of available fetchlets and information about their usage, see Chapter 21, "Using Fetchlets".</p> <p>The SCOPE property defines where the property value is to be obtained. The following scope options are available:</p> <ul style="list-style-type: none"> ■ SCOPE="GLOBAL". Obtain the property value from other variables defined within the current target type metadata file. This includes constants, such as the "I" shown in Example 3–3. ■ SCOPE="INSTANCE". Obtain the property value from instance properties. ■ SCOPE="ENVxxx". Obtain the property value from an environment variable "xxx". ■ SCOPE="SYSTEMGLOBAL". Obtain the property value from the emd.properties file located in the \$AGENT_HOME/sysman/config directory. ■ SCOPE="USER". Obtain the property value from the collector or user.

Table 3–3 (Cont.) Key Elements Used to Define Metrics

Element	Description
PushDescriptor	<p>Defines object identifiers (OIDs). The Management Agent uses an SNMP receivelet to listen for SNMP traps as defined in the target type's Management Information Base (MIB). The PushDescriptor is part of the metric definition of a receivelet.</p> <p>Defines a command to run, which returns the set of data described in the TableDescriptor. The element contains one or more Property elements, each defining a property to pass in with the command invocation.</p> <p>Note: You can refer to earlier defined properties using the <code>%property_name%</code> format. For example:</p> <pre><Property NAME="MatchSpecificTrap" SCOPE="GLOBAL">50</Property> <Property NAME="MatchAgentAddr" SCOPE="INSTANCE">AdminAddress</Property></pre> <p>The element includes a <code>RECVLET_ID</code> attribute that identifies the receivelet mechanism that will be used to process the request. Properties required for the receivelet invocation are specified in one or more Property elements within the PushDescriptor.</p> <p>The SNMP receivelet is the only receivelet commonly used by plug-in developers. For more information about the SNMP receivelet, see Chapter 20, "Using Receivelets".</p> <p>The SCOPE property defines where the property value is to be obtained. The following scope options are available:</p> <ul style="list-style-type: none"> ■ <code>SCOPE="GLOBAL"</code>. Obtain the property value from other variables defined within the current target type metadata file. For example, if a <code>Match*</code> property is GLOBAL-scoped, then it determines with which metric the trap is associated. ■ <code>SCOPE="INSTANCE"</code>. Obtain the property value from instance properties. For example, if the <code>MatchAgentAddr</code> property is an INSTANCE-scoped property, then it determines to which target instance the trap belongs. <p>For examples of SCOPE property definitions, see Example 20–4 through to Example 20–9.</p>
ElementDescriptor	<p>Used to compute aggregation metrics. Specifies the execution plan for evaluating a metric. The Management Agent runs each statement of the plan, in the order it is defined, to produce a Metric Result. The Metric Result generated as result of the evaluation of the last statement of the execution plan will be returned.</p>

3.5 Creating the Default Collection File

The default collection metadata file for a target type defines the following:

- The metric data (including configuration collection metric data) to be collected from targets and written to the Management Repository
- The frequency of at which this metric data is collected
- Thresholds that, when exceeded, will cause a Metric Alert event to be raised
- An optional message to display when a threshold is exceeded

Note: Although you can use any name for a default collection file, Oracle recommends that you use a naming convention that makes it easy to associate the collection file with the corresponding target type metadata file name. For example, using the same file name as the target type metadata file.

For information about naming the target type metadata file, see [Section 3.3.2, "Naming the Target Type Metadata File"](#).

Note that the value of the `TYPE` attribute in the default collection metadata file *must* match the `TYPE` value defined in the target type metadata file to create an association between them.

As noted, you can also specify Metric Alert event conditions on each metric that will be raised as Incidents within Enterprise Manager Cloud Control. Such events are generated when a threshold specified in this file is exceeded. For example, you may want to raise a `WARNING` alert when CPU usage reaches 90% of capacity. You can also specify the message text to be displayed in Enterprise Manager Cloud Control when an alert event is triggered.

The EDK includes an example of a default collection file in the following location:

```
/samples/plugins/HostSample/stage/oms/metadata/default_collection/demo_
hostsample.xml
```

For information about defining the elements in the default collection file, see [Section 3.5.5, "Overview of Key Default Collection Metadata Elements"](#) and [Section 3.6, "Guidelines for Creating Target Metadata"](#).

3.5.1 Grouping Similar Metrics For Collection

For efficiency, metrics are typically grouped together for collection, enabling certain metrics to be collected at the same time or same frequency. This is useful because it guarantees the order of execution of the metrics, which is important if some metrics rely on the results of other metrics.

Each group of metrics to be collected together is defined in a `CollectionItem` within the default collection file. The collection schedule for the group is defined in a `Schedule` element.

Each metric included in the group is in turn defined within a `MetricColl` element within the `CollectionItem`. (Note that if the `CollectionItem` contains just a single metric, like the Response metric example shown in [Section 3.4.2, "Defining the Basic Response Metric Group"](#), it is not necessary to specify the `MetricColl` tag.)

Note that the `UPLOAD` value for the `CollectionItem` is set to 6, meaning that every sixth collection of data will be written to the Management Repository. Because the `IntervalSchedule` specifies that data will be collected every 5 minutes, the data will be written to the Management Repository every 30 minutes (or every sixth data collection).

```
<TargetCollection>
...
<CollectionItem NAME="Perf" UPLOAD="6">
  <Schedule>
    <IntervalSchedule INTERVAL="5" TIME_UNIT="Min"/>
  </Schedule>
  <MetricColl NAME="CPUProcessesPerf">
    ...
  </MetricColl>
  <MetricColl Name="MemoryPerf">
    ...
  </MetricColl>
</CollectionItem>
...
</TargetCollection>
```

You should consider grouping metrics into a `CollectionItem` if:

- The metrics are logically related, such as all metrics related to performance

- The metrics should be collected at the same frequency, such as all metrics that should be collected every 5 minutes
- The metrics should be collected at roughly the same time, such as metrics collected during non-peak times
- You want to collect all of the metrics, or none of the metrics, at the same time

Note that if you have metrics that will be collected on demand, grouping them will improve performance and reduce the communications required by the Management Agent and Oracle Management Service to collect and return metric data from the target.

Ideally, you should avoid having too many independent or singleton collections, as changing the collection schedule for multiple independent metrics is a more cumbersome task. Grouping together unrelated metrics just to avoid having such singletons is not advisable either, as you will not have the ability to turn off collection of just a few metrics in the group without disabling the those metrics that you do need.

3.5.2 Defining Basic Metric Collection

The following represents the `CollectionItem` entry for the basic `Response` metric group, which includes the `Status` metric. It specifies that data for this metric should be collected every 5 minutes, which is the standard collection interval for this type of metric.

A condition has been set on the `Status` metric. For more information about alert conditions, see [Section 3.5, "Creating the Default Collection File"](#) and [Table 3-4](#).

Note that because the `CollectionItem` contains just one metric (`Status`), it is not necessary to include a `MetricColl` tag for the single metric.

```
<TargetCollection META_VER="2.0" TYPE="acme_demo_targetType">
...
<CollectionItem NAME="Response">
  <Schedule>
    <IntervalSchedule INTERVAL="5" TIME_UNIT="Min"/>
  </Schedule>
  <Condition COLUMN_NAME="Status" CRITICAL="0" OPERATOR="EQ" CLEAR_MESSAGE_
NLSID="Response_Status_clearalertmessage"
  MESSAGE="Failed to connect to database instance: %oraerr%."MESSAGE_
NLSID="Response_Status_alertmessage"/>
</CollectionItem>
...
</TargetCollection>
```

3.5.3 Defining Advanced Metric Collection

The next example illustrates collection of a more advanced metric that raises a metric alert when a specified `WARNING` and/or `CRITICAL` thresholds are exceeded. These thresholds, and the message to send to Cloud Control when they are exceeded, are defined in the `Condition` element.

The data for each metric is specified in a `MetricColl` element within a `CollectionItem`, as shown in [Example 3-6](#). For a description of the elements in this example, see [Table 3-4](#).

Example 3-6 Defining Advanced Metric Collection

```
<TargetCollection>
```

```

...
<CollectionItem NAME="Perf" UPLOAD="6">
  <Schedule>
    <IntervalSchedule INTERVAL="5" TIME_UNIT="Min"/>
  </Schedule>
  <MetricColl NAME="CPUProcessesPerf">
    <Condition COLUMN_NAME="ProcCPU" WARNING="75" CRITICAL="90" OPERATOR="GE"
      OCCURRENCES="2"
      MESSAGE="The value for %columnName% is %value%%, which is above the
        critical (%critical_threshold%%) or warning (%warning_threshold%%)
        threshold."
      CLEAR_MESSAGE="The value for %columnName% is %value%%, which is
        below the critical (%critical_threshold%%) or warning (%warning_
        threshold%%) threshold." />
  </MetricColl>
</CollectionItem>
...
</TargetCollection>

```

Note that in addition to a message sent to Enterprise Manager when either the WARNING or CRITICAL thresholds are passed, and “all clear” message to be sent when an alert condition no longer exists has also been defined in the CLEAR_MESSAGE attribute.

3.5.4 Defining Target Configuration Data Collections

As with all other metrics, the frequency at which the configuration metric data is collected is defined default collection file. Given the size of target configuration collections and the infrequent change rate, these metrics should ideally be collected every 24 hours, during off-peak hours.

Note that the value of the TARGET_TYPE attribute of the root METADATA SNAP_TYPE attribute in the configuration metadata file must be identical the TYPE attribute of TargetCollection in the default collection file.

The following example defines the collection frequency for the HostConfig metric

```

<TargetCollection>
...
<CollectionItem NAME="HostSampleSnap" UPLOAD_ON_FETCH="TRUE" CONFIG="TRUE">
  <Schedule OFFSET_TYPE="INCREMENTAL">
    <IntervalSchedule INTERVAL="24" TIME_UNIT="Hr"/>
  </Schedule>
  <MetricColl NAME="HostConfig" />
</CollectionItem>
...
</TargetCollection>

```

3.5.5 Overview of Key Default Collection Metadata Elements

Table 3–4 describes the key elements included in the default collection metadata file.

Table 3–4 Key Elements Within the Default Collection Metadata File

Element	Description
TargetCollection	Required. The root element for the file. It includes a TYPE attribute and the META_VER attribute that must match the TYPE attribute and the META_VER attribute of the TargetMetadata element in the target type metadata file.

Table 3–4 (Cont.) Key Elements Within the Default Collection Metadata File

Element	Description
CollectionItem	<p>Defines a collection frequency and threshold values for a set of metrics. The frequency defined in the included <code>Schedule</code> element will apply to all metrics in the collection group. The element includes the following attributes:</p> <ul style="list-style-type: none"> ■ NAME: Defines the set of metrics to be collected as a group. ■ UPLOAD: Specifies what <i>n</i>th data collection is written to the Management Repository. The default is 1, which means that performance data is uploaded every time it is collected. In Example 3–6, every 6th data collection is uploaded to the Management Repository. ■ UPLOAD_ON_FETCH: When set to TRUE, the first configuration collection is uploaded to the Management Repository immediately. All subsequent collections are bundled. By default, this value is set to FALSE and the UPLOAD_ON_FETCH attribute is ignored.
Schedule	<p>Contains an <code>IntervalSchedule</code> element defining the collection frequency for a <code>CollectionItem</code>. It includes the following attributes:</p> <ul style="list-style-type: none"> ■ INTERVAL: The collection frequency. ■ TIME_UNIT: The unit of time (such as Min for minutes) that the value of INTERVAL corresponds to.
MetricColl	<p>Contains one or more <code>Condition</code> elements corresponding to a single metric group defined in a <code>Metric</code> element in the target type metadata file. The NAME attribute in this element must match the NAME attribute in the corresponding <code>Metric</code> element.</p>

Table 3–4 (Cont.) Key Elements Within the Default Collection Metadata File

Element	Description
Condition	<p>Defines a metric alert condition. It contains the following optional attributes:</p> <ul style="list-style-type: none"> ■ COLUMN_NAME: The name of a metric defined in a ColumnDescriptor element in the target type metadata file. The value of this attribute must match the NAME attribute of the ColumnDescriptor element. ■ WARNING: Defines the threshold at which a "warning" condition exists. A metric alert will be generated when this value is exceeded, which will include the text specified in the MESSAGE attribute. To allow users to set the threshold value, set this attribute to "NotDefined". ■ CRITICAL: Defines the threshold at which a "critical" condition exists. A metric alert will be generated when this value is exceeded, which will include the text specified in the MESSAGE attribute. To allow users to set the threshold value, set this attribute to "NotDefined". ■ OPERATOR: Determines how to apply the threshold values specified in the CRITICAL and WARNING attributes. In Example 3–6, GE specifies that the Warning threshold occurs when ProcCPU is greater than or equal to 75 and the Critical threshold occurs when ProcCPU is greater than or equal to 90. Other values include: LE: Less than or equals EQ: Equals LT: Less than GT: Greater than NE: Not equal CONTAINS: True if the second argument is a substring of the first string. MATCH: True if the first argument (regular expression) matches the second argument. ■ OCCURENCES: Defines the number of successive metric collections that must be returned with the Warning or Critical threshold exceeded before the warning or critical condition is triggered. ■ MESSAGE: Contains a message to display when the thresholds specified in the WARNING and/or CRITICAL attributes have been exceeded. The built-in message attributes such as %columnName% and %critical_threshold% are embedded in the string; the appropriate value will be substituted when the message is generated at runtime. For example: "The value for %columnName% is %value%%%. It has fallen below the critical (%critical_threshold%%%) or warning (%warning_threshold%%%) threshold." ■ CLEAR_MESSAGE: Contains an "all clear" message that will be displayed when the value of the metric returns to a "non-alert" value; that is, when it drops below the thresholds indicated in WARNING and CRITICAL.

3.6 Guidelines for Creating Target Metadata

When developing target type definition files for new plug-ins, special consideration must be paid to the way in which you want a particular target type to be monitored. How a target type is monitored can greatly affect Enterprise Manager performance. Follow these general guidelines for defining target metadata and collections to optimize system performance.

3.6.1 Defining Target Metadata

Metadata is data about data. Generically, the term refers to any data used to help the identification, description, and location of a network entity. Target metadata for an Enterprise Manager target consists of the metrics a user wants to expose and the methods used to compute those metrics.

3.6.1.1 Metadata Version

Whenever the target metadata changes, increment the metadata version (META_VER).

3.6.1.2 Real-time Only Metrics

Performance metrics can be classified into metrics that must be computed to track performance trends and others that are more useful to drill down to get the details at a particular point in time. Real-time only metrics include those that need contextual information to return detailed information about a particular subset of the system, such as the resource utilization for specific processes, to diagnose further.

3.6.1.3 Choice of Key Columns

A key column in a metric is used in the management repository to trend performance data on an axis, such as the tablespace usage per database tablespace. Key-based metrics should be used to model sub-components of the target for which meaningful metric data should be collected, either for target monitoring or target diagnostic purposes. As such, only key columns that are the logical identifiers of the target sub-components should be included in the metric.

Note that including key columns for which the number of distinct values collected across a large number of targets could result in an excessive number of key values being stored in the management repository. For example, using a timestamp (or equivalent, like database SCN or Unix ctime) as a key value will result in a new value for every collection for every target, and is therefore not advisable.

Including a combination of key columns can also be problematic. For example, if you include three key columns in a metric, in which each key can take one of 10 target-specific values (10X10x10) multiplied by the number of targets, you would be collecting data for 1000 keys per target. This could be considered excessive if more than a handful of targets are being managed.

You can have no key columns, but the query descriptor must return a single row.

3.6.1.4 Transient Columns

In some cases, metric columns can be used to compute the values of other more interesting metric columns. When the original columns are not required, then you can mark these columns as transient so that they are not uploaded to the Management Repository, therefore saving space.

3.6.1.5 Metrics and Microsoft Windows

When creating metrics for custom targets, it is important to take into account the cost (CPU usage) of creating additional operating system (OS) processes. This is especially true for systems running Microsoft Windows where process creation is much more CPU intensive compared to UNIX-based systems such as Linux or Oracle Solaris. The percentage CPU utilization increases linearly with creation of child processes. To minimize process creation, avoid executing OS programs or commands from metric collection scripts. For example, when writing Perl scripts, avoid using the system function or backticks (`) to execute an OS command.

3.6.1.6 Target Properties (Static Versus Dynamic)

Target properties are named values that can be used for computing the metrics of the target, or for display in the home page of the target. The list of target properties is specified in the metadata to allow data driven user interfaces to register targets, and for the Management Agent to validate that a target instance is complete.

- **Static Instance Properties:** These are properties whose values need to be specified for a target in the targets.xml entry for the target. An instance property can be marked optional if the target declaration is considered complete even without the specification of the property. The metadata specification of a target property can also provide a default value for use in a configuration user interface.
- **Dynamic Instance Properties:** The Management Agent also allows for target instance properties to be *computed*. Such properties are computed using a QueryDescriptor very similar to the ones used in metrics.

Use of dynamic properties reduces the work involved in configuring a target by allowing certain properties to be computed rather than requiring the user to correctly specify their values (for example, the Version property of a database can be reliably computed given addressing information).

The Management Agent allows for the fact that the target needs to be up for the successful computation of these dynamic properties by recomputing the properties each time a target bounce is detected (each time the target status changes to Up).

3.6.1.7 Metrics

The metric concept, as it pertains to the Management Agent, can be used to denote configuration and performance information.

- **Configuration Metrics:** Configuration metrics collect data similar to target properties that denote the configuration of the target. This information is periodically refreshed and can be used to track changes in the setup of a target. The collection interval on such metrics is typically on the order of about 24 hours.
- **Performance Metrics:** Performance metrics are used to track the responsiveness of a target. These metrics are typically collected more often than configuration metrics though the interval of some performance metrics may vary widely from those of others. Also, performance metrics usually ship with thresholds that are the basis of performance alerts for the target.

A required metric for all targets is the "Response" metric consisting of a "Status" column with a condition on it. This metric is used to track the availability of the target.

3.6.1.7.1 Metric Naming Conventions The conventions used in naming your metrics are extremely important because many areas of the Enterprise Manager user-interface are

data-driven. For example, actual metric column labels and key values can be part of the page title, instruction text, or column headings. Specifically, these elements would appear on the All Metrics page, Metric and Collections Settings page, Event Rules page, Group Charts page, and other pages within the Enterprise Manager user-interface. For this reason, Oracle recommends the following metric naming conventions.

- Metric column names should never be vague or ambiguous, but should be as explicit as possible. For example:
 Vague name: Count
 Better, more descriptive name: Blocking Session Count
- All metric column names (labels) must be unique within a given target type and version, and easily understood by the user (metric units used as required). If the metric refers to a unit of measure, the unit should be included in the metric label inside a parenthesis. For example:
 Network Interface Total I/O Rate (MB/sec)
- All metric column names (labels) should be self-explanatory without dependence on the metric name. For example:
 Tablespace Space Used (%)
- Key column names should be self-explanatory. These names are used when specifying metric thresholds or configuring notifications. For readability, the name of the key column name should fit easily within the phrase "all <key column name> objects". For example:
 all tablespace objects
- Short names (up to 20 characters) associated with the metric column should be both clear and translatable.
- Across target versions, the same columns should use the same labels. This ensures columns, such as metric columns and short names, have the same NLS IDs across different target versions.

3.6.2 Defining Collections

Collections are the mechanism by which the Management Agent periodically computes the metrics of a target and uploads the data to the Management Repository. The most important thing to keep in mind when creating the collections for a target type is to avoid overburdening the Management Repository with excess data. In a large enterprise with hundreds of Management Agents and thousands of targets, the key to scalability is to limit the amount of data collected about a target that is uploaded to the repository. This is especially important since raw data is maintained for 24 hours - rollup benefits only accrue beyond that point.

3.6.2.1 Alert Message Guidelines

Alert messages tell the user when something is wrong. These messages should also assist the user in solving the problem. Oracle recommends following these content guidelines when writing alert messages:

- Alert messages should be meaningful, and should contain all of the information needed to effectively describe the issue to the user. Avoid using terse, ambiguous messages unless the message is only applicable to the specific metric.

The most significant part of the message should be covered within the first 60 characters of the message text. The reason is that the first part of the message is the most visible to users in e-mail notifications, incident tables containing the alert message, etc.

- Target down messages should, in addition to indicating that the target is down, include information indicating possible reasons why the target may be down.
- Error codes/messages should be included whenever possible.

Note that you should not include information on how to resolve or diagnose the problem in the alert message. You should instead provide this content in the Guided Resolution section of Incident Manager. See [Section 9.5, "Providing Content in the Guided Resolution Region"](#) for more information.

3.6.2.2 Metric Evaluation Order

It is important to pay attention to metric evaluation order so as to avoid metric collection failures. For example, the Response metric should be evaluated first in order to prevent a collection failure when a target is down. When a `CollectionItem` tag is used to define a collection, then the Management Agent evaluates all metrics with a collection item in order. However, collection items run independent of each other.

Note: Programmatic logic of the Management Agent distributes the metric evaluations so that each evaluation is separated by approximately 10 seconds.

3.6.2.3 Collection Frequency

In general, there is almost never a good reason to collect information at intervals smaller than 5 minutes. In the rare case where data variations occur at a smaller granularity and administrators need to be notified sooner, the Management Agent provides the capability to use a small collection interval to compute the metrics and threshold information while still only uploading data once in every *nn* computation cycles.

3.6.2.4 Controlling Number of Rows

Some metrics can result in the creation of a large number of rows in a Management Repository table. In some cases, only a subset of these rows may need to be uploaded to the repository. The Management Agent allows the specification of filter conditions that can be used to find rows to skip uploading. Also, a "limit_to" clause can be used on metrics that return sorted metric data to upload only the first *n* rows to the repository.

3.7 Testing Your Target Type Definitions

Test your new target type definitions by using the metric browser. The metric browser is a development utility that is an integral part of the Management Agent. As a subsystem of the Management Agent, it allows you to quickly access the metric values for targets monitored by the Management Agent without the overhead of a Management Repository or other components of the Enterprise Manager framework.

3.7.1 Activate the Metric Browser

To configure the Management Agent's metric browser for debugging metrics without the Enterprise Manager Cloud Control console:

1. Check that the `_enableMetricBrowser` line in the `$AGENT_HOME/sysman/config/emd.properties` file is enabled, where `AGENT_HOME` represents the home directory of the Management Agent:

```
_enableMetricBrowser=True
```

2. Enter the following command to apply the changes that you made to the `emd.properties` file:

```
emctl reload agent
```

3. Open the `emd.properties` file and check the `EMD_URL` line. It has the following format:

```
EMD_URL=http://host:port/emd/browser/main
```

Alternatively, you can use the `emctl` command to activate the metric browser as follows:

```
emctl setproperty agent -name _enableMetricBrowser -value true
```

3.7.2 View Your Metrics

After the target instance has been added to the `targets.xml` file and the new information has been reloaded, you can view available targets and metrics through the metric browser. Access the following URL using any web browser:

```
http://host:port/emd/browser/main
```

Tip: To find the port number used by the Management Agent, open the `$AGENT_HOME/sysman/config/emd.properties` file and search for the `EMD_URL` line.

Note: You must use the Management Agent operating system credentials to log in to the metric browser.

3.8 Validating Your Metadata XML

To verify that your target metadata files are defined correctly, enter the following command from the `bin` directory of the EDK:

```
empdk validate_plugin -stage_dir plugin_stage
```

In the preceding command, `plugin_stage` represents the plug-in staging directory.

For information about the EDK, see [Section 1.2, "About the Extensibility Development Kit \(EDK\)"](#) and for more information about the `empdk` command and its usage, see [Section 13.3, "Validating the Plug-in"](#).

3.9 Troubleshooting the Target Creation Process

This section provides some troubleshooting tips if you encounter any issues with your targets.

My target is not added to Enterprise Manager

If your target is not added, do the following:

- Check the Oracle Management Service trace file (emoms.trc) for exceptions. The OMS trace file is located in the *EM_INSTANCE_BASE/em/OMS_NAME/sysman/log/* directory, where *EM_INSTANCE_BASE* is the OMS Instance Base directory (by default, this directory is under the parent directory of the Oracle Middleware Home).

```
grep -i EntityManager.createEntities *
grep -i EntityUtil *
```

- If your target is added to the Management Repository but not to the Management Agent, go to the *agentStateDir/sysman/log* directory and check the Management Agent log file (gcagent_mdu.log). This log file tracks the metadata updates to the Management Agent.

My target continues to show a pending status

If your target is monitored by the Management Agent and it shows a pending status, then do the following:

- Check if the Management Agent is still monitoring the target.
To listing the name and type of each target being monitored by a Management Agent:
 1. Change directory to the *AGENT_HOME/bin* directory (UNIX) or the *AGENT_HOME\bin* directory (Windows).

2. Enter the following command:

```
emctl config agent listtargets
```

3. Check the output for your target.

- Check that the plug-in is deployed on the Management Agent by reviewing the following log file:

```
agent_inst/sysman/registry.xml
```

- Check that the Management Agent received the request to add the target. Go to the *agentStateDir/sysman/log* directory and review the Management Agent log file (gcagent_mdu.log).
- From a SQL*Plus session, run the *tginfo.sql* script, similar to:

```
@tginfo oracle_database orcl
```

The *tginfo.sql* script includes the following:

```
SELECT type_meta_ver, ':'||category_prop_1||':'||
      category_prop_2||':'||
      category_prop_3||':'||
      category_prop_4||':'||
      category_prop_5||':' category_prop,
       target_guid,
       TO_CHAR(load_timestamp,'DD-MON-YY HH24:MI:SS'),
       timezone_region,owner,host_name,emd_url,broken_reason,broken_str,manage_
status,
       promote_status,
       dynamic_property_status
FROM sysman.em_targets
WHERE target_type='&1'
      AND target_name='&2'
/
```

Note: If you are having issues running the script, edit the script to replace &&1 with the type of the target and replace &&2 with the name of the target.

The output from the script includes the following information:

- TARGET_TYPE
Name of the target, such as oracle_database
- TYPE_META_VER
Metadata version number. Check that the metadata version is correct for the target.
- CATEGORY_PROP_1
Category properties can be used to distinguish different metric definitions based on different configurations. Check that the value is correct for the target.
- BROKEN_REASON
If this value is greater than 0, then target is broken (for example, the target could not be saved or it is missing required properties). The BROKEN_STR value will provide a reason as to why the target is broken.
- MANAGE_STATUS
The manage status of the target. Possible values include:
 - * 0: Ignored
 - * 1: Not yet managed
 - * 2: Managed
 - * 3: Managed target component
- PROMOTE_STATUS
The promotion status of the target. Possible values include:
 - * 0: Cannot promote (an existence-only entity)
 - * 1: Eligible for promotion
 - * 2: Promotion in progress
 - * 3: Promoted to Management Agent
 - * 4: Promotion in progress but target was added to the Management Agent previously
- DYNAMIC_PROPERTY_STATUS
Status of the dynamic properties. Possible values include:
 - * 0: Dynamic properties have not been uploaded by the Management Agent
 - * 1: Dynamic properties are uploaded by the Management Agent

Adding Information Publisher Reports

Defining new target types in Enterprise Manager through metadata plug-ins provides you with the opportunity to add new report definitions. Plug-ins also allow you to add permanent (SYSTEM) target type specific report definitions to Enterprise Manager using the Information Publisher XML file format.

Note: Information Publisher is deprecated as of Enterprise Manager 12c and Oracle recommends using BI Publisher for newly created reports.

This chapter covers the following:

- [Introduction to Adding Information Publisher Reports](#)
- [Overview of SYSTEM Reports](#)
- [Understanding the Report Definition File](#)
- [Creating a Report Definition File](#)
- [Understanding the XML Report Definition Interface](#)
- [About Enterprise Manager Command Line Interface \(EM CLI\) Verbs](#)
- [About Development Guidelines](#)

4.1 Introduction to Adding Information Publisher Reports

As an plug-in developer, you are responsible for the following steps with regard to adding information publisher reports:

1. Design your reports based on the information you wish to show.
2. Create your report definition file
 - a. Define the SQL and PL/SQL queries used to extract information from the management repository.
For more information, see [Defining SQL or PL/SQL Queries](#).
 - b. Create a test report interactively from the Enterprise Manager console
For more information, see [Creating a Test Report Interactively from the Enterprise Manager Console](#).
 - c. Use EM CLI to generate the report definition file
For more information, see [Using EM CLI to Generate the Report Definition File](#).

4.1.1 Assumptions and Prerequisites

This chapter assumes you are familiar with the following:

- EM repository views against which you can write your own queries.
- Familiarity with the XML file format which you will use to create your report definition.

4.2 Overview of SYSTEM Reports

Adding report definitions through metadata plug-ins creates target type specific SYSTEM reports. SYSTEM report definitions are handled differently than definitions created through the Information Publisher user interface. SYSTEM reports are permanent and cannot be deleted or edited by Enterprise Manager administrators. You can add multiple report definitions to a metadata plug-in, thereby allowing you to associate multiple reports with a specific target type.

Adding SYSTEM report definitions using metadata plug-ins and the Information Publisher XML files allows users to access reports from the Enterprise Manager console's Information Publisher Report Definition page.

4.2.1 About the Report Definitions Page

All report definitions added using metadata plug-ins are available from Information Publisher's Report Definitions page. As with out-of-box SYSTEM report definitions, those added using metadata plug-ins are organized according to report category and subcategory. SYSTEM report definitions cannot be deleted from the Enterprise Manager console.

4.3 Understanding the Report Definition File

A report definition file is an XML file that contains code to extract pertinent information from the Management Repository (using repository views) and the report elements used to format and display that data. The Information Publisher API allows you to specify the report elements and parameters that you normally specify when creating a report definition from the Enterprise Manager console. The fully formed report definition file consists of four basic XML tags and takes on a hierarchical tag structure:

- `<ReportDefinition>`
Defines report identification parameters as well as encapsulating all report elements used to build the report
- `<ReportElement>`
Defines the graphical display elements such as tables, charts, or text.
- `<ReportElementParameters>`
Defines specific parameters required by individual report elements.
- `<ReportWideParameters>`
Defines parameters used by all report elements in the report definition file.

4.4 Creating a Report Definition File

As previously mentioned, the content of a report definition file consists of XML tags used to construct a report. You will use both the Enterprise Manager console and EM CLI to develop and generate your report definition file.

Metadata plug-ins allow you to define as many report definition files as required for a particular target type.

4.4.1 About the Report Definition File Development Process

The process of developing a valid report definition file involves three steps:

1. [Defining SQL or PL/SQL Queries](#)
2. [Creating a Test Report Interactively from the Enterprise Manager Console](#)
3. [Using EM CLI to Generate the Report Definition File](#)

Defining SQL or PL/SQL Queries

The first step in creating your report definition is to create the SQL or PL/SQL queries used to extract the requisite report information from the Management Repository. Enterprise Manager provides management views with which you can safely extract data from the Management Repository without reading from the base tables. Using repository views protects your queries from changes to the repository schema that may occur in future releases and ensures your SYSTEM report definitions remain functional.

The following query was used to extract information about blackout history for a target. The query uses the MGMT\$BLACKOUT_HISTORY, MGMT\$BLACKOUTS, and MGMT\$TARGET repository views. The query uses the MGMT\$METRIC_CURRENT repository view.

```
select 'senior mts', count(value) from mgmt$metric_current
where metric_column = 'Title' and LOWER(value) like '%senior member%' and
      target_guid = ??EMIP_BIND_TARGET_GUID??

select bh.created_by "Created by", bh.start_time "Start", bh.end_time "End",
bo.reason "Reason", bo.description "Description"
from
MGMT$BLACKOUT_HISTORY bh, MGMT$TARGET tgt, MGMT$BLACKOUTS bo
where tgt.target_name = bh.target_name and tgt.target_type = bh.target_type
      and tgt.target_guid = ??EMIP_BIND_TARGET_GUID?? and bo.blackout_guid =
      bh.blackout_guid
order by end_time desc

union
select 'consulting mts', count(value) from mgmt$metric_current
where metric_column = 'Title' and LOWER(value) like '%consulting%' and
      target_guid = ??EMIP_BIND_TARGET_GUID?? ;
```

When an administrator views a report from the Enterprise Manager console that contains this SQL query string, Information Publisher automatically binds the unique identifier for the selected target to the ??EMIP_BIND_TARGET_GUID?? placeholder in the SQL query string. The documentation for Chart from SQL and Table from SQL parameters provide information on this bind variable placeholder as well as others you can include in your SQL query string.

Creating a Test Report Interactively from the Enterprise Manager Console

Once you have written and tested the SQL or PL/SQL query, you can use the Enterprise Manager console to generate a version of your report interactively using the Chart from SQL and Table from SQL report elements. By using the Information Publisher user interface, you can easily prototype reports without having to create a report definition file and import Plug-in Archive (OPAR) files.

You can also use this method of interactive prototyping to refine your queries and ensure that the data extracted from the Management Repository and how that information is rendered in your report meets your reporting requirements.

Using EM CLI to Generate the Report Definition File

Once you are satisfied with the way your report is being rendered by Information Publisher, you are ready to create the report definition file. To do this, you use EM CLI to generate the XML based Report Definition file. The EM CLI `export_report` verb is used to export the report definition you developed using the Enterprise Manager console (stored in the Management Repository) and to generate the XML report definition file. For example:

```
>emcli export_report
  -title="resource report"
  -owner="ADMINISTRATOR_JOE"
  -output_file="$HOME/reports/resource_report.xml"
```

Once the report definition file is generated, you will need to edit the XML to insert your own plug-in specific information, such as `product_name`, `component_name`, and `oms_version`.

[Example 4–1](#) shows the content of the report definition file for a report detailing host configuration.

Example 4–1 Host Configuration Report Definition File

```
<?xml version = '1.0' encoding = 'UTF-8'?>
<ReportDefinition title="Host Performance Overview"
description="Overview of host performance" system_report="0"
category="Sample Host Reports" sub_category=
"Performance Reports" show_navigation="1" generate_context="0"
add_toc="0" product_name="EM" component_name="oracle_hostsample" is_jit_multi_
target="0" target_type="oracle_hostsample" is_jit_target="1" style="BLAF" oms_
version="11.1.0.1.0" xmlns="http://www.example.com/DataCenter/ReportDefinition">

<ReportElement element_row="1" suppress_render="0"
  element_name_nlsid="IPMSG_USER_CHART_FROM_SQL" header_nlsid="Average CPU
  Utilization (%)" element_type_nlsid="IPMSG_ANY_TARGET_TYPE" element_order="0">
  <ReportElementParameters
    parameterName="oracle.sysman.eml.ip.render.elem.
    ChartParamController.chartType" parameterValue="pieChart"/>
  <ReportElementParameters parameterName="oracle.sysman.eml.ip.
    render.elem.sqlStatement" parameterValue="select column_label, value
    &quot;CPU Utilization (%)&quot; &#xA;
    from mgmt$metric_current where &#xA;
    target_guid = ??EMIP_BIND_TARGET_GUID??&#xA;
    and metric_name = 'CPUPerf'"/>
  <ReportElementParameters parameterName="oracle.sysman.eml.ip.render.
    elem.ChartParamController.width" parameterValue="200"/>
  </ReportElement>
<ReportElement element_row="3" suppress_render="0"
  element_name_nlsid="IPMSG_USER_CHART_FROM_SQL" header_nlsid="Memory Utilization
```

```
(KB)" element_type_nlsid="IPMSG_ANY_TARGET_TYPE" element_order="1">
  <ReportElementParameters parameterName="oracle.sysman.eml.ip.render.
    elem.ChartParamController.legendPosition" parameterValue="south"/>
  <ReportElementParameters parameterName="oracle.sysman.eml.ip.
    render.elem.ChartParamController.chartType" parameterValue="barChart"/>
  <ReportElementParameters parameterName="oracle.sysman.eml.ip.
    render.elem.sqlStatement" parameterValue="select column_label, value
    &quot;Memory Utilization (KB)&quot;;&#xA;
    from mgmt$metric_current where &#xA;
    target_guid = ??EMIP_BIND_TARGET_GUID??&#xA;
    and metric_name = 'MemoryPerf'"/>
  <ReportElementParameters parameterName="oracle.sysman.eml.ip.render.
    elem.ChartParamController.visualOrientation" parameterValue="horizontal"/>
  <ReportElementParameters parameterName="oracle.sysman.eml.ip.
    render.elem.ChartParamController.width" parameterValue="600"/>
</ReportElement>
</ReportDefinition>
```

4.4.2 About the Report Lifecycle: Updating Report Definitions

With the ability to add report definitions to Enterprise Manager comes the responsibility of maintaining and updating the report definitions. Familiarity with the way in which Enterprise Manager handles report definitions will allow you to anticipate system behavior and plan for backwards compatibility.

When report definitions are deployed using metadata plug-ins, Enterprise Manager allows newer versions of the report definitions to be installed. Report definitions which are not valid with a newer version of Enterprise Manager should be updated and redeployed with the new version of the plug-in. Enterprise Manager will not install older versions of a report definition.

Report definitions, as with metadata plug-ins in general, should be designed with backwards compatibility in mind. Future versions of report definitions should support previous versions of the target type metadata. Report definition-metadata version incompatibility will be most apparent in the following situations:

- Report definitions included with metadata plug-in version 1 and not included with metadata plug-in version 2 will not disappear when version 2 is deployed.
- If version 1 and version 2 of a metadata plug-in are both deployed to the system, Management Agents will collect data based on the metadata of the version installed at that Agent; some will collect for version 1 metadata and some for version 2 metadata. Only the version 2 report definitions will be installed (appear in the Enterprise Manager console). For this reason, version 2 report definitions must support both versions of the metadata.

4.5 Understanding the XML Report Definition Interface

The Information Publisher XML based report definition file provides an easily editable medium for defining and customizing your Information Publisher reports using simple XML tags.

4.5.1 About Report Definition Tags

Use the following XML tags to define and manipulate report information when creating report definition files.

- [<ReportDefinition>](#)

- [<ReportElement>](#)
- [<ReportElementParameters>](#)

4.5.1.1 <ReportDefinition>

The <ReportDefinition> tag is the first XML tag that appears in the report definition file and specifies essential information about your report such as title, description, product name, or Oracle Management Service version. The following example shows the <ReportDefinition> tag as defined for a host configuration report.

Example 4–2 <ReportDefinition> Tag for the Host Configuration Report

```
<ReportDefinition
title="Host Configuration Overview"
description="Overview of host configuration" system_report="0"
category="Sample Host Reports"
sub_category="Configuration Reports"
show_navigation="1"
generate_context="0"
add_toc="0"
product_name="EM"
component_name="oracle_hostsample"
is_jit_multi_target="0"
target_type="oracle_hostsample"
is_jit_target="1"
style="BLAF"
oms_version="11.1.0.1.0"
xmlns="http://www.example.com/DataCenter/ReportDefinition">
```

Tag Attributes

Table 4–1 Tag Attributes for the Host Configuration Report

Attribute	Description
title	report title
description	description
category	category name
sub_category	subcategory name
target_type	target type for late binding, or null if not late binding
add_hoc	1=show 0=hide table of contents
show_navigation	show navigation headers in report (tabs, etc) 1=show, 0=hide
product_name	product name, 'EM' (default)
component_name	product name. This must be set to the metadata plug-in target type.
oms_version	version '11.1' (default)

Report-wide Parameters

```
<ReportWideParameters
  parameterName="oracle.sysman.eml.ip.render.elem.TimePeriodParam"
  parameterValue="0:1"/>
```

4.5.1.2 <ReportElement>

The <ReportElement> tag is used to add a new report element to an existing report definition.

Input

Table 4–2 <ReportElement> Tag

Parameter	Description
element_type_nlsid	The element type name.
header_nlsid	The element header or null.
element_order	The order of this element, 1 based.
element_row	The row for this element, 1 based.

4.5.1.3 <ReportElementParameters>

The <ReportElementParameters> tag is used to declare the parameters used for a report element. Include all of the report element parameters you want to declare within the <ReportElement> tag.

Table 4–3 <ReportElementParameters> Tag

Parameter	Description
parameterName	The parameter name.
parameterValue	The parameter value.

4.5.2 Using Element Parameters

Parameters used by some report elements dictate the operational behavior of those elements. Use the <ReportElementParameters> tag to declare element parameters associated with a <ReportElement>. The parameter names and values are described in this section for each element type.

This section lists the parameters associated with specific report elements.

4.5.2.1 About Table Element Parameters

The Table Element is used to show a tabular view of query results. The queries must be made against management views.

To declare a Table Element, you would use the following in the <ReportElement> tag:

```
<ReportElementParameters
  element_name_nlsid=" IPMSG_USER_TABLE_FROM_SQL"
  element_type_nlsid="IPMSG_ANY_TARGET_TYPE" .....
```

- Element Name nlsid: IPMSG_USER_TABLE_FROM_SQL
- Element Type nlsid: IPMSG_ANY_TARGET_TYPE

Time Period

Table 4–4 Table Element Parameters Time Period

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TimePeriodParam".
Required	No.
Default Value	Null.
Valid Values	"0:0" for last 24 Hours. "0:1" for last 7 Days. "0:2" for last 31 Days.
Summary	Encoded time period.

Sort Column

Table 4–5 Table Render Sort Column

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.initialSortColumn".
Required	No.
Default Value	The first column in result set.
Valid Values	Any valid column name.
Summary	If this parameter is set, the sort column indicator will be shown for the column with this column name. If not set, the sort column indicator is shown on the first column. The SQL query should include an 'order by' clause that sorts by this column.

Sort Order

Table 4–6 Table Render Sort Order

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.initialSortOrder".
Required	No.
Default Value	"ascending".
Valid Values	"ascending" or "descending".
Summary	If this parameter is set, the sort column indicator will be shown either as ascending or descending, according to the value. If not set, the sort column indicator is shown as ascending.

Name Value Pair Display

Table 4–7 Name Value Pair Display

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.nameValueDisplay".
Required	No.
Default Value	<none>.
Valid Values	Positive integer value.

Table 4–7 (Cont.) Name Value Pair Display

Attribute	Description
Summary	If this parameter is set and only one row is returned from the query, the results are displayed in a vertical list of name-value pairs. This value should be set to the number of name/value columns that should be displayed, normally "1".

Number of Rows to Show**Table 4–8 Number of Rows**

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.numRowsToShow".
Required	No.
Default Value	"10".
Valid Values	Positive integer value.
Summary	Number of rows to display at one time in the generated table. The user can scroll through additional rows using the UI controls.

Is PL/SQL Statement**Table 4–9 Is PL/SQL Statement**

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.sqlStatementIsPlSql".
Required	No.
Default Value	"false".
Valid Values	"true" or "false" .
Summary	Whether a SQL statement is PL/SQL.

SQL or PL/SQL Statement**Table 4–10 SQL or PL/SQL Statement**

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.sqlStatementIsPlSql".
Required	No.
Default Value	<None>.
Valid Values	Any valid SQL SELECT statement.

Table 4–10 (Cont.) SQL or PL/SQL Statement

Attribute	Description
Summary	<p>SQL statement can optionally bind values for targets, locale information, and start/end date. The format of the SQL statement should include a bind variable placeholders for the options to be bound.</p> <p>Bind Placeholders:</p> <ul style="list-style-type: none"> ■ ??EMIP_BIND_RESULTS_CURSOR?? For use with PL/SQL statement to bind a return cursor containing results for display. ■ ??EMIP_BIND_TARGET_GUID?? For use with SQL or PL/SQL to bind a target GUID. ■ ??EMIP_BIND_START_DATE?? For use with SQL or PL/SQL to bind a start date. ■ ??EMIP_BIND_END_DATE?? For use with SQL or PL/SQL to bind an end date. ■ ??EMIP_BIND_TIMEZONE_REGION?? For use with SQL or PL/SQL to bind a time zone region. ■ ??EMIP_BIND_LOCALE_COUNTRY?? For use with SQL or PL/SQL to bind a locale country. ■ ??EMIP_BIND_LOCALE_LANGUAGE?? For use with SQL or PL/SQL to bind a locale language. <p>There should be no semi-colon (;) appended to the end of the SQL statement unless it is a PL/SQL statement.</p>

Example 4–3 Specifying an Anonymous PL/SQL Block as a Parameter to an Element Definition

To avoid issues with formatting, generate the report and export it to XML using the EM CLI.

```
<ReportElementParameters
parameterName="oracle.sysman.eml.ip.render.elem.sqlStatement"
parameterValue="BEGIN DECLARE&#xA;&#xA; BEGIN&#xA;&#xA;
open ??EMIP_BIND_RESULTS_CURSOR?? for select &#xA;
bh.created_by &quot;Created by&quot;;, &#xA; bh.start_time &quot;Start&quot;;&#xA;
bh.end_time &quot;End&quot;;&#xA; bo.reason &quot;Reason&quot;;&#xA;
bo.description &quot;Description&quot;;&#xA; from &#xA; MGMT$BLACKOUT_HISTORY bh,
&#xA; MGMT$TARGET tgt,&#xA; MGMT$BLACKOUTS bo&#xA;
where tgt.target_name = bh.target_name&#xA; and tgt.target_type = bh.target_
type&#xA; and tgt.target_guid = ??EMIP_BIND_TARGET_GUID??&#xA; and bo.blackout_
guid = bh.blackout_guid&#xA;
order by end_time desc;&#xA;&#xA;&#xA; END;&#xA;END;" />
```

Named SQL Statement**Table 4–11 Named SQL Statement**

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.NamedSqlStatement".
Required	No.
Default Value	<none>.

Table 4–11 (Cont.) Named SQL Statement

Attribute	Description
Valid Values	Any valid statement name.
Summary	<p>As an alternative to the "oracle.sysman.eml.ip.render.elem.sqlStatement".</p> <p>parameter you may use a Named SQL statement which refers to an actual SQL statement stored in the Enterprise Manager repository.</p> <p>You can register a Named SQL statement by providing an XML file containing the name of the SQL statement as well as the SQL query as part of plug in metadata.</p> <p>For more information, see Section 4.6, "Using the ImportExport.xsd File" for the Named SQL XML file XSD definition.</p>

Maximum Number of Rows

Table 4–12 Number of Rows

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.maxNumberOfRowsAllowed".
Required	No.
Default Value	"2000".
Valid Values	Any scalar numeric value.
Summary	Set the maximum number of rows retrieved for display in the table. For example, show the top 10 xyz's element would set the value to "10".

Null Data String Substitute

Table 4–13 Null Data String Substitute

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.nullDataStringSubstitutue".
Required	No.
Default Value	""
Valid Values	A string.
Summary	A string that will be substituted for null values returned.

Split Table into Multiple Tables by Column

Table 4–14 Split Table

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.nullDataStringSubstitutue".
Parameter String	"oracle.sysman.eml.ip.render.elem.TableRender.tableSplitColumn".
Required	No.
Default Value	Null.
Valid Values	Any valid column name.

Table 4–14 (Cont.) Split Table

Attribute	Description
Summary	If this parameter is set, the table will be split into separate tables with subheaders as the value in this column changes. The data should be ordered by this column.

Column Group Header

Table 4–15 Column Group Header

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.columnGroupHeader"n.
Required	No.
Default Value	Null.
Valid Values	Header string to use for a column group.
Summary	This parameter provides a column header string. This column group header will span columns between the columns specified in "oracle.sysman.eml.ip.render.elem.TableRender.columnGroupStart Col"n and oracle.sysman.eml.ip.render.elem.TableRender.columnGroupEndCol"n. The n suffix is a numeric value starting with 1 for the first column group, sequentially ascending for subsequent column groups.

Column Group Start Column

Table 4–16 Column Group Start Column

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.columnGroupStartCol"n.
Required	No.
Default Value	Null.
Valid Values	Any valid column name.
Summary	Specifies the first column for a given column group. The n suffix is a numeric value starting with 1 for the first column group, sequentially ascending for subsequent column groups.

Column Group End Column

Table 4–17 Column Group End Column

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.columnGroupEndCol"n.
Required	No.
Default Value	Null.
Valid Values	Any valid column name.
Summary	Specifies the first column for a given column group. The n suffix is a numeric value starting with 1 for the first column group, sequentially ascending for subsequent column groups.

Use Separate Rows for Values within a Cell

Table 4–18 *Use Separate Rows for Values Within a Cell*

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.useSeparateRowsColumns".
Required	No.
Default Value	Null.
Valid Values	Comma separated list of valid column names.
Summary	If this parameter is set, the delimited values of the column with the given name specified will be displayed on separate rows within a containing row cell. More than one column can be designated for this treatment by adding comma-separated column names.

Use Separate Rows as Delimiters

Table 4–19 *Use Separate Rows as Delimiters*

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.useSeparateRowsDelimiter".
Required	No.
Default Value	, (comma).
Valid Values	Any string.
Summary	A character used to delimit tokens within a string.

Severity Icon in Column

Table 4–20 *Severity Icon in Column*

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.severityColumn".
Required	No.
Default Value	Null.
Valid Values	Any valid column names.
Summary	A severity icon will be substituted for valid severity values returned. To omit an icon, your result set can contain null values in this column.

Availability Status Icon in Column

Table 4–21 *Availability Status Icon in Column*

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.availabilityStatusColumn".
Required	No.
Default Value	Null.
Valid Values	Any valid column names.

Table 4–21 (Cont.) Availability Status Icon in Column

Attribute	Description
Summary	An availability status icon will be substituted for valid values returned. To omit an icon your result set can contains null values in this column.

Render Image in Column

Table 4–22 Render Image in Column

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.imageFilenameColumns".
Required	No.
Default Value	Null.
Valid Values	Comma separated list of column names.
Summary	Optional parameter to display the given image filename in the indicated columns. Indicate for which columns the given image should be rendered. Specify a comma separated list of column names. The image filename returned should contain a relative path starting with '/images' such as '/images/xyz.gif'. Normally, a SQL decode function would be used to translate a numeric value into the appropriate image filename.

Target Type Column

Table 4–23 Target Type Column

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.targetTypeColumns".
Required	No.
Default Value	Null.
Valid Values	Comma separated list of column names.
Summary	Optional parameter to indicate for which columns the value returned should be used as an internal target type to be translated into a display string for that type. Specify a comma separated list of column names.

4.5.2.1.1 About Filter Elements The following table elements are used to create search filters that allows users to filter on rows for multiple table columns. Three different filter types are permitted:

- text-value
- list of values obtained from a SQL query
- list of values obtained from a comma-separated list in the element definition

Define Filter Name

Table 4–24 Define Filter Name

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.filterNames".
Required	Yes.
Default Value	Null.

Table 4–24 (Cont.) Define Filter Name

Attribute	Description
Valid Values	Comma separated list of filter names.
Summary	Defines filter names in a comma-separated list. This parameter also defines the ordering of filter elements.

Define Filter Prompt

Table 4–25 Define Filter Prompt

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.filterPrompt<name>Filter".
Required	Yes.
Default Value	Null.
Valid Values	CF.
Summary	Defines the prompt used in the Reports page for the filter name. The filter value is accessed from the report element's SQL statement via ??EMIP_BIND_PARAM<name>??. Without any other filter-related parameters, this defines a filter which allows the user to provide a value via a text input field.

SQL Filter

Table 4–26 SQL Filter

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.filterSql<name>".
Required	No.
Default Value	Null.
Valid Values	Any valid SQL SELECT statement.
Summary	Defines the SQL query used to populate a list of values for a filter name that is presented in the UI as a drill-down menu instead of the text input field.

List of Filter Names

Table 4–27 List of Filter Names

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.filterList<name>".
Required	No.
Default Value	Null.
Valid Values	Comma separated list of values.
Summary	Defines a list of values for a filter name which is displayed in the UI as a drill-down menu.

Translate List of Filter Names

Table 4–28 *Translate List of Filter Names*

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.filterTranslateValues<name>".
Required	No.
Default Value	No.
Valid Values	yes or no.
Summary	Defines whether the values provided by filterSql or filterList should be translated to the client locale.

Filter Tip Text

Table 4–29 *Filter Tip Text*

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.filterTip<name>Filter".
Required	No.
Default Value	Null.
Valid Values	Alpha numeric text.
Summary	Defines the text for a tool tip shown if the user moves the mouse over the filter UI elements.

Default Filter Name

Table 4–30 *Default Filter Name*

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.filterDefault<name>".
Required	No.
Default Value	%.
Valid Values	Alpha numeric text string.
Summary	Defines a default value for filter name. If no default value is given, '%' is used instead.

Null Default Filter Name

Table 4–31 *Null Default Filter Name*

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.filterDefaultsToNull<name>".
Required	No.
Default Value	Null.
Valid Values	yes or no.
Summary	When this defines the default value to be NULL instead of '%'.

Global Filter Elements

The following parameters act globally on the filter system.

Display Empty Table

Table 4–32 *Display Empty Table*

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.filterStartEmpty".
Required	No.
Valid Values	yes or no.
Summary	If the value of this parameter is 'yes', then the report initially displays an empty table. The table is populated when the user clicks on the filter button in the UI.

Empty Table Headers

Table 4–33 *Empty Table Headers*

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.filterEmptyTableHeaders".
Required	No.
Default Value	Null.
Valid Values	Comma-separated list of table headers.
Summary	Defines the table headers used when starting with an empty table.

Table Header Type

Table 4–34 *Table Header Type*

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.filterEmptyTableHeaderTypes".
Required	No.
Default Value	VARCHAR.
Valid Values	Comma-separated list of table headers.
Summary	This defines the table header types (column types) used if when starting with an empty table. This is a comma-separated list. If no header types are specified, the table header types default to VARCHAR.

Overwrite Table Header Text

Table 4–35 *Overwrite Table Header Text*

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.filterHeaderText".
Required	No.
Default Value	Search Filter.
Valid Values	Comma separated list of column names.

Table 4–35 (Cont.) Overwrite Table Header Text

Attribute	Description
Summary	Overwrites the default filter section header text.

Overwrite Default Filter Description**Table 4–36 Overwrite Default Filter Description**

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.filterDescriptionText".
Required	No.
Default Value	Enter values to filter what is shown in the table.
Valid Values	Alpha numeric text string.
Summary	Overwrites the default filter section header text.

Overwrite Default Filter Tip Text**Table 4–37 Overwrite Default Filter Tip Text**

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.filterTipText".
Required	No.
Default Value	The search filter is case sensitive. Use '%' as a wildcard.
Valid Values	Alpha numeric text string.
Summary	Overwrites the default filter section tip text.

Overwrite Default Button Text**Table 4–38 Overwrite Default Button Text**

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.filterButtonText".
Required	No.
Default Value	OK.
Valid Values	Alpha numeric text string.
Summary	Overwrites the default filter button text.

Empty Table Text**Table 4–39 Empty Table Text**

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.filterEmptyTableText".
Required	No.
Default Value	(No rows returned).
Valid Values	Alpha numeric text string.
Summary	Specifies the text to be shown in an empty table before the filter is run.

4.5.2.1.2 Using Hyperlinks Within Tables The following parameters are used to implement hyperlinks within tables and incorporate improved link navigation between master/detail views. This method is an alternative to using `oracle.sysman.eml.ip.render.elem.TableRender.columnDestReportTitle<num>`, which first takes the user to the target selector page.

Link to Report

Table 4–40 Link to Report

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.columnDestHomepageReportTitle<num>".
Required	No.
Default Value	Null.
Valid Values	Report definition link.
Summary	Same as <code>oracle.sysman.eml.ip.render.elem.TableRender.columnDestReportTitle<num></code> except that a link to a report definition on the target homepage is created.

Display Number of Columns

Table 4–41 Display Number of Columns

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.numberOfColumnsShown".
Required	No.
Default Value	Number of columns in the SQL.
Valid Values	Number.
Summary	Defines the number of columns from the element SQL to be displayed in the UI. Additional columns from the SQL query are hidden but can be used to create the hyperlinks to expose data in the detail report.

Display Target Name

Table 4–42 Display Target Name

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.columnDestTargetIndex<num>".
Required	No.
Default Value	Null.
Summary	Specifies the column (which may be hidden) that contains the target name. The target name is used in the link to populate the target selection on late binding reports.

Display Target Type

Table 4–43 Display Target Type

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.columnDestTypeIndex<num>".
Required	No.
Default Value	Null.
Summary	Specifies the column (which may be hidden) that contains the target type. The target type is used in the link to populate the target selection on late binding reports.

Display URL

Table 4–44 Display URL

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TableRender.columnDestURLIndex<num>".
Required	No.
Default Value	Null.
Summary	Specifies the column (which may be hidden) that contains an arbitrary URL for a given table element.

Example 4–4 Report definition defining a master report that allows you to drill down using a link to a detail report

```
<?xml version = '1.0' encoding = 'UTF-8'?>
<ReportDefinition title="My Master Report" description=
"A master report to show master/detail" system_report="0"
category="Test Reports" sub_category="Master and Detail"
show_navigation="1" generate_context="0" add_toc="1"
product_name="EM" component_name="SAMPLE" is_jit_multi_target="0" is_jit_
target="0" style="BLAF" oms_version="11.2.0.1.0"
xmlns="http://www.example.com/DataCenter/ReportDefinition">

<ReportElement element_row="1" suppress_render="0"
element_name_nlsid="IPMSG_USER_ TABLE_FROM_SQL" header_nlsid="My Master Report
Table" element_type_nlsid="IPMSG_ANY_TARGET_TYPE" element_order="0">
  <ReportElementParameters parameterName="oracle.sysman.eml.ip.render.
elem.TableRender.filterEmptyTableHeaders" parameterValue=
"Target Name, Target Type"/>
  <ReportElementParameters parameterName="oracle.sysman.eml.ip.render.elem.
TableRender.columnDestParamColumnIndexes1" parameterValue="0,1"/>
  <ReportElementParameters parameterName="oracle.sysman.eml.ip.render.elem.
TableRender.filterHeaderText" parameterValue="My Filter Header"/>
  <ReportElementParameters parameterName="oracle.sysman.eml.ip.render.elem.
TableRender.filterDescriptionText" parameterValue="My Filter description"/>
  <ReportElementParameters parameterName="oracle.sysman.eml.ip.
render.elem.sqlStatement" parameterValue="SELECT TARGET_NAME &quot;Target
Name&quot;, TARGET_TYPE &quot;Target Type&quot;
FROM MGMT$TARGET WHERE TARGET_NAME LIKE
??EMIP_BIND_PARAMNAME?? AND TARGET_TYPE LIKE
??EMIP_BIND_PARAMTYPE??" />
  <ReportElementParameters parameterName="oracle.sysman.eml.ip.render.elem.
```

```

    TableRender.filterPromptNAME" parameterValue="Name"/>
<ReportElementParameters parameterName="oracle.sysman.eml.ip.render.
    elem.TableRender.filterSqlTYPE" parameterValue=
    "select distinct target_type from mgmt$target"/>
<ReportElementParameters parameterName="oracle.sysman.eml.ip.render.elem.
    TableRender.filterStartEmpty" parameterValue="yes"/>
<ReportElementParameters parameterName="oracle.sysman.eml.ip.render.elem.
    TableRender.filterTipTYPE" parameterValue="Filter on the target types"/>
<ReportElementParameters parameterName="oracle.sysman.eml.ip.render.elem.
    TableRender.filterPromptTYPE" parameterValue="Target Type"/>
<ReportElementParameters parameterName="oracle.sysman.eml.ip.render.elem.
    TableRender.filterTipText" parameterValue="My Tip Text"/>
<ReportElementParameters parameterName="oracle.sysman.eml.ip.render.elem.
    TableRender.filterNames" parameterValue="NAME,TYPE"/>
<ReportElementParameters parameterName="oracle.sysman.eml.ip.render.elem.
    TableRender.numberOfColumnsShowed" parameterValue="2"/>
<ReportElementParameters parameterName="oracle.sysman.eml.ip.render.elem.
    TableRender.columnDestReportTitle1" parameterValue="My Detail Report"/>
<ReportElementParameters parameterName="oracle.sysman.eml.ip.render.elem.
    TableRender.filterTipNAME" parameterValue="Filter on the target names"/>
<ReportElementParameters parameterName="oracle.sysman.eml.ip.render.elem.
    TableRender.filterButtonText" parameterValue="My button text"/>
</ReportElement>
</ReportDefinition>

<?xml version = '1.0' encoding = 'UTF-8'?>
<ReportDefinition title="My Detail Report" system_report="0"
    category="Test Reports" sub_category="Master and Detail" show_navigation="1"
    generate_context="0" add_toc="0" product_name="EM"
    is_jit_multi_target="0" is_jit_target="0" style="BLAF"
    oms_version="11.2.0.1.0"
    xmlns="http://www.example.com/DataCenter/ReportDefinition">
<ReportElement element_row="1" suppress_render="0"
    element_name_nlsid="IPMSG_USER_TABLE_FROM_SQL" element_type_nlsid="IPMSG_ANY_
    TARGET_TYPE" element_order="0">
<ReportElementParameters
    parameterName="oracle.sysman.eml.ip.render.elem.headerParam"
    parameterValue="Target Detail"/>
<ReportElementParameters
    parameterName="oracle.sysman.eml.ip.render.elem.sqlStatement"
    parameterValue="SELECT TARGET_NAME &quot;Target Name&quot;;
    TYPE_VERSION &quot;Version&quot;;
    FROM MGMT$TARGET
    WHERE TARGET_TYPE LIKE ??EMIP_BIND_PARAM2??
    AND TARGET_NAME LIKE ??EMIP_BIND_PARAM1??"/>
</ReportElement>
</ReportDefinition>

```

4.5.2.2 About the Chart Element

The Chart Element is used to show a graphical view of query results. The queries must be made against Management Repository views.

- Element Name: IPMSG_USER_CHART_FROM_SQL
- Element Type: IPMSG_ANY_TARGET_TYPE

Chart Type

Table 4–45 Chart Type

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.ChartParamController.chartType".
Required	No.
Default Value	"pie chart".
Valid Values	"barChart", "lineChart", "pieChart", "timeSeriesChart", and "timeSeriesBarChart".
Summary	Chart type to display.

Time Period

Table 4–46 Time Period

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TimePeriodParam".
Required	No.
Default Value	Null.
Valid Values	"0:0" for last 24 Hours. "0:1" for last 7 Days. "0:2" for last 31 Days.
Summary	Encoded time period.

Fill

Table 4–47 Fill

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.ChartParamController.fill".
Required	No.
Default Value	"none".
Valid Values	"none", "absolute", or "cumulative".
Summary	Indicates if a line chart should fill the area under the lines. "none": no fill under lines. "absolute": lines are identical to "none" setting but with the area under the lines filled. "cumulative": causes the values for the lines to be added or stacked, then the areas underneath the lines are filled. Use caution when using the fill attribute to ensure there is no confusion for the report user as to whether the data in the chart is cumulative or absolute.

Height

Table 4–48 *Height*

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.ChartParamController.height".
Required	No.
Default Value	"200".
Valid Values	<i>n</i> , where <i>n</i> is any String that will correctly parse to a positive integer.
Summary	Sets the display height of the chart in pixels.

Horizontal or Vertical

Table 4–49 *Horizontal or Vertical*

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.ChartParamController.visualOrientation".
Required	No.
Default Value	"horizontal".
Valid Values	"horizontal" or "vertical".
Summary	Visual orientation of the chart. This attribute is only valid with the chartType attribute set to barChart or timeSeriesChart. The attribute does not affect the pieChart.

Legend Position

Table 4–50 *Legend Position*

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.ChartParamController.legendPosition".
Required	No.
Default Value	"east".
Valid Values	"default", "east", "south".
Summary	Specifies where the legend should be placed relative to the chart.

Is PL/SQL Statement

Table 4–51 *Is PL/SQL Statement*

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.sqlStatementIsPLSql".
Required	No.
Default Value	"false".
Valid Values	"true" or "false".
Summary	Set to "true" to indicate that the SQL statement is a PL/SQL statement.

SQL or PL/SQL Statement

Table 4–52 SQL or PL/SQL Statement

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.sqlStatement".
Required	No.
Default Value	<none>.
Valid Values	Any valid SQL SELECT statement or PL/SQL block.
Summary	<p>SQL or PL/SQL statement can optionally bind values for targets, locale information, and start/end date. The format of the statement should include a bind variable placeholders for the options to be bound.</p> <p>Bind Placeholders:</p> <ul style="list-style-type: none"> ■ ??EMIP_BIND_RESULTS_CURSOR?? For use with PL/SQL statement to bind a return cursor containing results for display. ■ ??EMIP_BIND_TARGET_GUID?? For use with SQL or PL/SQL to bind a target GUID. ■ ??EMIP_BIND_START_DATE?? For use with SQL or PL/SQL to bind a start date. ■ ??EMIP_BIND_END_DATE?? For use with SQL or PL/SQL to bind an end date. ■ ??EMIP_BIND_LOCALE_COUNTRY?? For use with SQL or PL/SQL to bind a locale country. ■ ??EMIP_BIND_LOCALE_LANGUAGE?? For use with SQL or PL/SQL to bind a locale language. <p>There should be no semi-colon (;) appended to the end of the SQL statement unless it is a PL/SQL statement.</p>

Stacked Bar Chart

Table 4–53 Stacked Bar Chart

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.ChartParamController.stacked".
Required	No.
Default Value	"false".
Valid Values	"true" or "false".
Summary	Indicates if a bar chart should be stacked.

Chart Title

Table 4–54 Chart Title

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.ChartParamController.title".
Required	No.
Default Value	<none>.

Table 4–54 (Cont.) Chart Title

Attribute	Description
Summary	Chart title to identify chart for Americans with Disabilities Act compliance.

Width**Table 4–55 Width**

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.ChartParamController.width".
Required	No.
Default Value	"400".
Valid Values	<i>n</i> , where <i>n</i> is any String that will correctly parse to a positive integer.
Summary	Specifies the display width of the element in pixels.

Y-Axis Label**Table 4–56 Y-Axis Label**

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.yAxisLabel".
Required	No.
Default Value	<none>.
Valid Values	String.
Summary	If this parameter is supplied, it is used as the y-axis label for charts that have an y-axis.

Slices as Percentage**Table 4–57 Slices as Percentage**

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.ChartParamController.pieShowSlicePercentageLabels".
Required	No.
Default Value	<none>.
Valid Values	"true" or "false".
Summary	If this parameter is supplied, it controls whether each slice is labeled with a percentage value. This attribute is ignored for chartTypes other than pieChart.

Show Values in Legend**Table 4–58 Show Values in Legend**

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.ChartParamController.pieValuesInLegend".

Table 4–58 (Cont.) Show Values in Legend

Attribute	Description
Required	No.
Default Value	"value".
Valid Values	"percent", "value" or "none".
Summary	For pie charts, this parameter specifies whether values for pie slices are included in the legend along with the label for the pie slice. The default value for this attributes is "value". If specified as either "percent" or "value" then the numeric value is displayed along with the pie slice label in the form, "pie slice label (numeric value)". If "percent" is specified, then the percentage out of the total of all slice values is calculated and displayed, otherwise, the raw value of the slice is displayed. To omit a value in the legend, specify "none" as a value for this parameter. This attribute is ignored for chartTypes other than pieChart.

4.5.3 Understanding the Metric Details Element

To declare a Metric Details Element, you would use the following in the ReportElement tag:

```
<ReportElementParameters
  element_name_nlsid=" IPMSG_METRIC_DETAILS"
  element_type_nlsid="IPMSG_ANY_TARGET_TYPE" .....
```

Target Type

Table 4–59 Target Type

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.MetDetInternalTargetType".
Required	No.
Default Value	"oracle_database".
Valid Values	Any valid internal target type name.
Summary	The type of target to be shown in the graph.

Metric Name

Table 4–60 Metric Name

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.MetDetSelectedMetric".
Required	Yes.
Valid Values	Valid metric name according to target type selected.
Summary	Metric to be graphed.

Metric Column Name

Table 4–61 Metric Column Name

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.MetDetSelectedMetricColumn".

Table 4–61 (Cont.) Metric Column Name

Attribute	Description
Required	Yes.
Valid Values	Valid column name according to the metric and target type selected.
Summary	Column of metric to be graphed.

Time Period**Table 4–62 Time Period**

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TimePeriodParam".
Required	No.
Default Value	null.
Valid Values	"0:0" for last 24 Hours. "0:1" for last 7 Days. "0:2" for last 31 Days.
Summary	Encoded time period.

Width**Table 4–63 Width**

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.MetDetWidth".
Required	No.
Default Value	300.
Valid Values	n , where n is any String that will correctly parse to a positive integer.
Summary	Width of the image in pixels.

Height**Table 4–64 Height**

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.MetDetHeight".
Required	No.
Default Value	300.
Valid Values	n , where n is any String that will correctly parse to a positive integer.
Summary	Height of the image in pixels.

Legend Position**Table 4–65 Legend Position**

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.MetDetLegendPosition".

Table 4–65 (Cont.) Legend Position

Attribute	Description
Required	No.
Valid Values	"south" (default), "east".
Summary	Position of the legend relative to the chart.

4.5.4 Using Text Element Parameters

The Text Element is used to display any message text you wish to provide for your report. To declare a Text Element, you would use the following in the `ReportElement` tag:

```
<ReportElementParameters
  element_name_nlsid="IPMSG_STYLED_TEXT"
  element_type_nlsid="IPMSG_ANY_TARGET_TYPE" .....
```

Message Text

Table 4–66 Message Text

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TextParamBean.textMessage".
Required	No.
Default Value	"" (empty String).
Valid Values	Any Message.
Summary	Set the message to display in the report.

Message Style

Table 4–67 Message Style

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TextParamBean.textStyleClass".
Required	No.
Default Value	"OraInstructionText".
Valid Values	"OraInstructionText". "OraTipText".
Summary	Specifies the style class for the message text to adopt when displayed.

Link Destination

Table 4–68 Link Destination

Attribute	Description
Parameter Name	"oracle.sysman.eml.ip.render.elem.TextParamBean.textDestination".
Required	No.
Default Value	None.
Valid Values	Any URI.

Table 4–68 (Cont.) Link Destination

Attribute	Description
Summary	Specifies an optional link destination for this text element.

4.5.5 About Report-Wide Parameters

The following parameters apply to all reporting elements within the report definition.

Dynamic Time Selector

You can provide a dynamic time period selector for your report definition that allows the report user to choose a specific time period with which to view the report.

If you are using Table from SQL or Chart from SQL report elements, you can structure your SQL statement such that the start and end dates will be bound automatically for you by Information Publisher. You achieve this by inserting placeholders (for example, `??EMIP_BIND_START_DATE??`) for the start and end date values as shown in [Example 4–5](#).

Example 4–5 Automatic Binding of Start and End Dates

```
'SELECT COLUMN_LABEL, ROLLUP_TIMESTAMP, AVERAGE
  FROM MGMT$METRIC_HOURLY
 WHERE TARGET_GUID = ??EMIP_BIND_TARGET_GUID??
  AND METRIC_LABEL = 'Load'
  AND KEY_VALUE = ' '
  AND ROLLUP_TIMESTAMP > ??EMIP_BIND_START_DATE??
  AND ROLLUP_TIMESTAMP < ??EMIP_BIND_END_DATE??
 ORDER BY ROLLUP_TIMESTAMP'
```

See the online help documentation for Table from SQL or Chart from SQL for detailed information.

4.6 Using the ImportExport.xsd File

The Information Publisher ImportExport.xsd file describes the format of the report definition XML file. A sample NamedSql.xsd file is also shown in [Example 4–7](#).

Example 4–6 Information Publisher ImportExport.xsd

```
<xsd:schema targetNamespace="http://www.example.com/DataCenter/ReportDefinition"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema"
  xmlns:ms="http://www.example.com/DataCenter/ReportDefinition"
  elementFormDefault="qualified" attributeFormDefault="unqualified">

  <xsd:annotation>
    <xsd:documentation>
      <strong>This is the schema definition, used by metadata services and the
report import cli. It is used to fully specify a report definiton</strong>
    </xsd:documentation>
  </xsd:annotation>

  <!-- ***** -->
  <!-- Main Element: ReportDefinition -->
  <!-- ***** -->
  <xsd:element name="ReportDefinition" type="ms:ReportDefinitionT"/>
  <!-- Defining Common Types used in a Report Definition -->
```

```

<!-- ***** -->
<!-- ReportDefinitionT -->
<!-- ***** -->
<!-- Documentation:
      ReportDefinitionT is type for main root element. All the Report
Definitions should validate to this type.
-->
<xsd:complexType name="ReportDefinitionT">
  <xsd:sequence>
    <xsd:element name="ReportWideParameters" type="ms:ReportWideParametersT"
minOccurs="0" maxOccurs="unbounded"/>
    <xsd:element name="ReportElement" type="ms:ReportElementT" minOccurs="0"
maxOccurs="unbounded"/>
  </xsd:sequence>
  <xsd:attribute name="title" type="ms:String100Def" use="required"/>
  <xsd:attribute name="description" type="ms:String500Def"/>
  <xsd:attribute name="system_report" type="ms:BooleanDef" default="0"/>
  <xsd:attribute name="category" type="ms:String100Def" use="required"/>
  <xsd:attribute name="sub_category" type="ms:String100Def" use="required"/>
  <xsd:attribute name="target_type" type="ms:String64Def"/>
  <xsd:attribute name="is_jit_target" type="ms:BooleanDef" default="1"/>
  <xsd:attribute name="is_jit_multi_target" type="ms:BooleanDef" default="0"/>
  <xsd:attribute name="add_toc" type="ms:BooleanDef" default="0"/>
  <xsd:attribute name="pack_name" type="ms:String64Def"/>
  <xsd:attribute name="style" type="ms:String64Def" default="BLAF"/>
  <xsd:attribute name="show_navigation" type="ms:BooleanDef" default="1"/>
  <xsd:attribute name="product_name" type="ms:String100Def" default="EM"/>
  <xsd:attribute name="component_name" type="ms:String100Def"/>
  <xsd:attribute name="generate_context" type="ms:BooleanDef" default="0"/>
  <xsd:attribute name="oms_version" type="ms:NameDef" use="required"/>
</xsd:complexType>

<xsd:complexType name="ReportElementT">
  <xsd:sequence>
    <xsd:element name="ReportElementParameters"
type="ms:ReportElementParametersT" minOccurs="0" maxOccurs="unbounded"/>
  </xsd:sequence>
  <xsd:attribute name="element_name_nlsid" type="ms:String256Def"
use="required"/>
  <xsd:attribute name="element_type_nlsid" type="ms:String100Def"
use="required"/>
  <xsd:attribute name="header_nlsid" type="ms:String100Def"/>
  <xsd:attribute name="element_order" type="xsd:integer"/>
  <xsd:attribute name="element_row" type="xsd:integer"/>
  <xsd:attribute name="suppress_render" type="ms:BooleanDef"/>
</xsd:complexType>

<xsd:complexType name="ReportElementParametersT">
  <xsd:attribute name="parameterName" type="ms:String100Def" use="required"/>
  <!-- parameterValue is in CDATA, but schema definition makes no
distinction between this and string attribute. Therefore,
don't specify any constraints on the attribute, thereby allowing
it to be unbounded in length.
-->
  <xsd:attribute name="parameterValue" type="xsd:string" use="required"/>
</xsd:complexType>

<xsd:complexType name="ReportWideParametersT">
  <xsd:attribute name="parameterName" type="ms:String100Def" use="required"/>
  <!-- parameterValue is in CDATA, but schema definition makes no

```

```

        distinction between this and string attribute. Therefore,
        don't specify any constraints on the attribute, thereby allowing
        it to be unbounded in length.
-->
        <xsd:attribute name="parameterValue" type="xsd:string" use="required"/>
</xsd:complexType>

<xsd:simpleType name="String64Def">
  <xsd:restriction base="xsd:string">
    <xsd:minLength value="1"/>
    <xsd:maxLength value="64"/>
    <xsd:whiteSpace value="preserve"/>
  </xsd:restriction>
</xsd:simpleType>

<xsd:simpleType name="String100Def">
  <xsd:restriction base="xsd:string">
    <xsd:minLength value="1"/>
    <xsd:maxLength value="100"/>
    <xsd:whiteSpace value="preserve"/>
  </xsd:restriction>
</xsd:simpleType>

<xsd:simpleType name="String500Def">
  <xsd:restriction base="xsd:string">
    <xsd:minLength value="1"/>
    <xsd:maxLength value="500"/>
    <xsd:whiteSpace value="preserve"/>
  </xsd:restriction>
</xsd:simpleType>

<xsd:simpleType name="String256Def">
  <xsd:restriction base="xsd:string">
    <xsd:minLength value="1"/>
    <xsd:maxLength value="256"/>
    <xsd:whiteSpace value="preserve"/>
  </xsd:restriction>
</xsd:simpleType>

<xsd:simpleType name="BooleanDef">
  <xsd:restriction base="xsd:integer">
    <xsd:enumeration value="0"/>
    <xsd:enumeration value="1"/>
  </xsd:restriction>
</xsd:simpleType>

<xsd:simpleType name="NameDef">
  <xsd:restriction base="xsd:string">
    <xsd:minLength value="1"/>
    <xsd:maxLength value="64"/>
  </xsd:restriction>
</xsd:simpleType>

</xsd:schema>

```

Example 4-7 NamedSQL.xsd

```

<xsd:schema xmlns:xsd="http://www.w3.org/2001/XMLSchema"
  xmlns:ms="http://www.example.com/DataCenter/NamedSQL"
  targetNamespace="http://www.example.com/DataCenter/NamedSQL"

```

```
        elementFormDefault="qualified" attributeFormDefault="unqualified">
    <xsd:annotation>
        <xsd:documentation>
            <strong>This is the schema definition, used by metadata services. It is
used to fully specify a list of named sql that can be used in report
definitions</strong>
        </xsd:documentation>
    </xsd:annotation>

    <xsd:element name="NamedSQLStatements">
        <xsd:complexType>
            <xsd:sequence>
                <xsd:element name="NamedSQL" minOccurs="0" maxOccurs="unbounded">
                    <xsd:complexType>
                        <xsd:attribute name="sqlName" type="xsd:string"/>
                        <xsd:attribute name="sqlValue" type="xsd:string"/>
                    </xsd:complexType>
                </xsd:element>
            </xsd:sequence>
        </xsd:complexType>
    </xsd:element>
</xsd:schema>
```

4.7 About Enterprise Manager Command Line Interface (EM CLI) Verbs

The following EM CLI verbs are used exclusively for report definition creation and administration.

Example 4-8 EM CLI Verbs

```
emcli get_reports
    [-owner="<report-owner>"]
```

Description:

This verb returns a list of reports owned by or viewable by the user logged into the cli.

Options:

-owner The optional argument allows listing of viewable reports owned by a specific EM user.

Output:

The output of this report will be space separated quoted strings for the report title and owner with each report on its own line.

```
emcli export_report
    -title="<report-title>"
    -owner="<report-owner>"
    -output_file="<file>"
```

Description:

This verb exports a report definition and all its element definitions given its title and owner.

Options:

-title The title of the report to export.

-owner The owner of the report to export. The logged-in emcli user must have view privilege for the report. Target names will not be exported. The report is uniquely defined using title/owner so both must be supplied.

```

-output_file
    The name of the exported file.
Examples:
emcli export_report \
    -title="maintenance report" \
    -owner="SHIFT1_OPERATOR" \
    -output_file="$HOME/reports/maint_report.xml"

emcli import_report
    [-force]
    -files="file1;file2;..."
Description:
    This verb imports a report definition from a XML file using
    the title in the xml file and the currently logged-in cli user
    as the owner of the report. If the report/owner already exists,
    the operation fails for that report with an appropriate error
    message. The report will be changed to a just-in-time report with
    the target type from the exported report. In addition, schedules
    and access privileges will need to be edited using the UI. The
    system enforces title/owner uniqueness, so an error will be thrown
    if there is already a report with the same title and owner.
Options:
    -force
        If report with same title/owner exists, first delete it
        (and all jobs and saved copies)
    -files
        List of Path/file name(s) of XML file(s), which contains
        valid Report definition(s).
Examples:
emcli import_report \
    -files="$HOME/reports/maint_report1.xml;$HOME/reports/file2.xml"

```

4.8 About Development Guidelines

Oracle recommends adhering to the following guidelines when defining a report definition file:

The Component Name Must be Set to the Target Type

The component name must be set to the target type in order for Enterprise Manager to associate specific report definitions with a particular metadata plug-in. For example,

```
<ReportDefinition component_name="oracle_orgchart" ... ..
```

When Using Chart from SQL and Table from SQL Elements

- If your element accepts a single non-aggregate target (only), which is the case for most metadata plug-in target types, you can take advantage of automatic time zone date adjustment built into the Chart from SQL and Table from SQL elements by setting the `oracle.sysman.eml.ip.render.elem.adjustTimes` parameter on your element to `'true'`. When this parameter is set, the start and end dates bound to your SQL query will be adjusted from the report time zone to the target time zone. Conversely, dates returned from the query will be adjusted from the target time zone to the report time zone.
- If your element accepts multiple targets or aggregate targets, you are responsible for handling time zone adjustment for your date values. You can obtain the report time zone from the `??EMIP_BIND_TIMEZONE_REGION??` bind variable. In order for the report viewer to understand the dates shown, dates displayed in a report

must either conform to the report time zone or explicitly display the time zone associated with each date. The following examples illustrate common use cases.

Example 4–9 Adjusting a Date Returned in your Select Statement from the Time Zone of a Given Target to the Report Time Zone

```
select mgmt_view_util.adjust_tz(tbl.date, tgt.timezone_region,  
??EMIP_BIND_TIMEZONE_REGION??)  
from mgmt$target tgt, sometable tbl  
where <your where clause here>
```

Example 4–10 Adjusting a Report Time Period Start and End Dates Used in the WHERE Clause of Your SELECT Statement from the Report Time Zone to your Targets Time Zone

```
select <your selected columns here>  
from mgmt$target tgt, sometable tbl  
where  
    tgt.target_guid = ??EMIP_BIND_TARGET_GUID?? and  
    tbl.Mydate > MGMT_VIEW_UTIL.ADJUST_TZ(  
    ??EMIP_BIND_START_DATE??,  
    ??EMIP_BIND_TIMEZONE_REGION??,  
    tgt.TIMEZONE_REGION)  
and  
    tbl.Mydate < MGMT_VIEW_UTIL.ADJUST_TZ(  
    ??EMIP_BIND_END_DATE??,  
    ??EMIP_BIND_TIMEZONE_REGION??,  
    tgt.TIMEZONE_REGION)
```

Developing BI Publisher Reports

Oracle Enterprise Manager Cloud Control 12c is integrated with the Oracle Business Intelligence Publisher 11g reporting product. Oracle Business Intelligence Publisher (BI Publisher) is the recommended approach to creating reports for Enterprise Manager Cloud Control 12c.

BI Publisher is a strategic enterprise reporting product from Oracle that provides the ability to create and manage highly formatted reports from a wide range of data sources. BI Publisher report formats can be designed using Microsoft Word or Adobe Acrobat and you can create reports from different types of data sources.

Release 11g of Oracle BI Publisher has an improved user interface (UI), many enhanced features, and new functions. These new functions include the Data Model Editor, which is a graphical user interface for building data models within the BI Publisher interface, and the Layout Editor, which is a design tool that enables you to create report layouts within the BI Publisher interface.

This chapter includes the following topics:

- [Introduction to Oracle BI Publisher](#)
- [Training and Resources](#)
- [About the Report Data Source](#)
- [Developing a Report](#)
- [Using the Enterprise Manager EDK for Staging and Deploying BI Publisher Reports](#)

5.1 Introduction to Oracle BI Publisher

As a plug-in developer, you are responsible for the following steps:

1. Developing the report using data models and report templates.

For more information on developing a report, see [Developing a Report](#)

2. Staging and deploying reports

For more information on staging and deploying reports, see [Using the Enterprise Manager EDK for Staging and Deploying BI Publisher Reports](#)

5.1.1 Assumptions and Prerequisites

This chapter assumes you are familiar with the following:

- Management repository views against which you can write your own queries.

- Familiarity with BI Publisher.

5.2 Training and Resources

Before you start to develop a BI Publisher report, you should take advantage of the training and reference resources available from Oracle:

- Getting Started with Oracle BI Publisher 11g
<http://st-curriculum.oracle.com/obe/fmw/bi/bip/bip11g/gettingstarted/gettingstarted.htm>
- Oracle BI Publisher on Oracle Technology Network (OTN)
<http://www.oracle.com/technology/products/xml-publisher/index.html>
- Oracle By Example Tutorial
http://www.oracle.com/technology/obe/obe_bi/bipub/index.html
- Oracle BI Publisher Blog
<http://blogs.oracle.com/xmlpublisher/>
- Oracle BI Publishing Consulting Blog
<http://bipconsulting.blogspot.com/>

5.3 About the Report Data Source

The EMREPOS data source is available from BI Publisher server configured for use with Enterprise Manager reports. The EMREPOS data source connects to the MGMT_VIEW account in the Management Repository and establishes the proper security context (VPD) for the Enterprise Manager user logged on to BI Publisher.

As a security measure, BI Publisher reports that use the EMREPOS data source have read-only access to the public MGMT\$VIEW and GC\$ views, and not to the underlying Enterprise Manager tables; this model also supports sharing of report queries with Enterprise Manager users who might want to use the Enterprise Manager-provided reports as a basis for their own reports.

5.4 Developing a Report

By default the reports and data models in the Enterprise Manager Cloud Control folder are read-only. Develop your own reports in your local folders and then have a BI Publisher System Administrator put the finished reports in a shared folder, outside of the Enterprise Manager Cloud Control folder.

To develop a BI Publisher report, complete the following:

1. Develop your data model, based on SQL queries against the EM repository.

To develop a BI Publisher report, the following components are required:

- Data Model
- Report Template
- Sub Template

First, you develop your data model, based on SQL queries against the EM repository. You then design the layout of your report (the template) using MS

Word on Windows. Your template refers to one of the two common Oracle subtemplates, namely, Portrait or Landscape.

You can use the following reports as examples of this:

Enterprise Manager Cloud Control -> EM Sample Reports -> Targets of Specified Type

Enterprise Manager Cloud Control -> EM Datamodels -> Targets of Specified Type

2. Develop and test the SQL queries for report data and input parameters.
3. Create a data model in BI Publisher for your data queries.
4. Create parameters in BI Publisher for your report parameters.
5. Create a report layout for your report.

Start with the sample landscape or portrait layout RTF file provided with the Extensibility Development Kit (EDK).

6. Create and test your report.

Download the report and data model using the BI Publisher UI. The download option is located on the 'more...' link under the name of each report/data model.

7. Export the report (.xdoz) and data models (.xdmz) from BI Publisher into local files.

5.5 Using the Enterprise Manager EDK for Staging and Deploying BI Publisher Reports

BI Publisher reports are deployed from a plug-in (that is, ADDON or MP) to a BI Publisher web application in a series of steps:

1. Create a metadata file that adheres to the following schema:

```
emcore/source/oracle/sysman/emSDK/ip/bipublisherreport/BIPublisherReport.xsd
```

The following is an example of a Reports metadata file:

```
<?xml version = '1.0' encoding = 'UTF-8' ?>
<BIPublisherReports
  xmlns="http://www.example.com/DataCenter/BIPublisherReport">
  <ReportFile relativePath="emreports" fileName="tvmlrb104a.jar"/>
  <ReportFile relativePath="emreports" fileName="tvmlrb104b.jar"/>
</BIPublisherReports>
```

2. Stage the BI Publisher Reports, which are ZIP files with the extension .xdoz (report definition) or .xdmz (report data model, that is, SQL) into one or more JAR files. These files are referenced in the previous metadata file.

The \$ORACLE_HOME/sysman/jlib/emreports.jar file provides an example.

```
$ unzip -l emreports.jar
Archive:  emreports.jar
Label:  EMGC_MAIN_LINUX_110220
  Length      Date    Time    Name
-----
0  02-20-2011  23:08  META-INF/
71 02-20-2011  23:08  META-INF/MANIFEST.MF
0  02-20-2011  23:08  bipublisherreports/
0  02-20-2011  23:08  bipublisherreports/EM Datamodels/
```

```

4776 02-20-2011 23:08 bipublisherreports/EM Datamodels/Usage Trend
Report.xdmz
4854 02-20-2011 23:08 bipublisherreports/EM Datamodels/Usage Summary
Report.xdmz
5008 02-20-2011 23:08 bipublisherreports/EM Datamodels/Charge Trend
Report.xdmz
7344 02-20-2011 23:08 bipublisherreports/EM Datamodels/Consolidation
Reports.xdmz
5043 02-20-2011 23:08 bipublisherreports/EM Datamodels/Charge Summary
Report.xdmz
0 02-20-2011 23:08 bipublisherreports/Chargeback/
52291 02-20-2011 23:08 bipublisherreports/Chargeback/Charge Trend
Report.xdoz
66994 02-20-2011 23:08 bipublisherreports/Chargeback/Charge Summary
Report.xdoz
26505 02-20-2011 23:08 bipublisherreports/Chargeback/Usage Trend
Report.xdoz
112150 02-20-2011 23:08 bipublisherreports/Chargeback/Usage Summary
Report.xdoz
0 02-20-2011 23:08 bipublisherreports/Consolidation Planner/
50114 02-20-2011 23:08 bipublisherreports/Consolidation
Planner/Consolidation Reports.xdoz
-----
335150 16 files

```

3. Create your plug-in JAR files.

The first directory in each JAR file must be `bipublisherreports`. All the data models for the reports must be in the same subdirectory, `EM_Datamodels`, just under the `bipublisherreports` directory. For example:

- a. The plug-in, which contains all the metadata files and BI Publisher report JAR files, is installed in an Oracle home directory (for Addons) or installed dynamically (for Metadata Plug-ins) using the plug-in environment.
- b. The BI Publisher report JAR files are placed in a subdirectory of the `metadata/bipublisherreport` directory and referenced in the metadata file (`emreports`).
- c. The BI Publisher reports for both platform and plug-ins are deployed to a BI Publisher web application (either at plug-in installation time or sometime later when BI Publisher is installed).

Plug-in reports are deployed when BI Publisher is integrated with Enterprise Manager using the `configureBIP` script or sometime later using one of the following EMCLI verbs:

- * `emcli setup_bipublisher` (see help for usage details)
- * `emcli deploy_bipublisher_reports [-force]`

This last verb deploys the EM System Reports (and optionally Extensibility Development Kit plug-in loaded reports) to a previously setup EM to BI Publisher relationship (using `setup_bipublisher`). It can also be used to upload a reports JAR file (located on the OMS filesystem). The operation will not overwrite existing BI Publisher Reports in the EM Reports folder unless `-force` is used in the command.

The following options are available:

Option	Description
[-pluginid]	In addition to EM system reports, you can use this option to deploy any subsequently loaded plug-in based BI Publisher reports.
[-pluginversion]	Limit the plug-ins to a specific version
[-reportsjarfile]	Deploy a single EM Reports JAR file that contains one or more BI Publisher Reports. This JAR file is located relative to the OMS's \$ORACLE_HOME

For example, the syntax for the emct plug-in would be:

```
emcli deploy_bipublisher_reports -pluginid=oracle.sysman.emct  
-pluginversion=12.1.0.0.0
```

Note: Using overlapping folders from the different JAR files and PLATFORM JAR files causes reports from the different JAR files to be placed in the same BI Publisher folder. If the same report name is referenced in multiple JAR files, there is no way of knowing which one will get deployed last, so this is not supported.

Collecting Target Configuration Data

This chapter provides information about defining configuration collection tables and integrating them into the Enterprise Configuration Management framework.

This chapter contains the following sections:

- [Introduction to Collecting Target Configuration Data](#)
- [About the Configuration Definition Files](#)
- [Modeling Enterprise Configuration Management Tables](#)

6.1 Introduction to Collecting Target Configuration Data

As a plug-in developer, you are responsible for the following steps with respect to incorporating configuration-related functions into Enterprise Configuration Management for each snapshot type:

1. Ensure that every bit of the configuration data that will be collected will not change from one collection to the next unless there is an explicit action by an administrator. Configuration data must only collect data that can change due to explicit administrator actions but should stay unchanged without such actions.
2. Define an Enterprise Configuration Management-specific metadata file that defines collection tables to match configuration collection metrics established as part of the target type definition.

For more information, see [Section 6.3.1, "Defining Configuration Collection Tables"](#).

3. Name the snapshot type (the name must match the `CollectionItem` name specified on the Management Agent side) for a given target type.

For more information, see [Section 6.3.2, "Overview of Configuration Management Snapshot Metadata Elements"](#).

4. Register the metadata with Enterprise Configuration Management during plug-in deployment. This creates tables for snapshot data, subject to any constraints imposed by the Enterprise Configuration Management framework. It also registers the data tables in Enterprise Configuration Management metadata tables.

For more information, see [Section 6.3.4, "Registering Metadata With the Configuration Management Framework"](#).

5. Integrate with the standard Management Agent's collection mechanisms. Both `CollectionItem` and corresponding metrics must be defined at the Management Agent.

For more information, see [Section 6.3.7, "Modifications to Standard Collection Metrics and RAW Metrics"](#).

6. Verify that the defined configuration collection data is returned to the Management Repository by the Management Agent.

For more information, see [Section 6.3.8, "Testing the Configuration Collection Data"](#) and [Section 6.3.9, "Troubleshooting"](#).

6.1.1 Assumptions and Prerequisites

This chapter assumes you are familiar with the following:

- Enterprise Manager concepts including Management Agents, metrics, and collection items
- Plug-in development overview, including how to package a plug-in and its XML files

6.2 About the Configuration Definition Files

The metadata for configuration collection is defined in three metadata files packaged with the plug-in:

- The configuration metadata file

This file defines the tables in the Management Repository that will store collected configuration data. In addition, the EDK includes an example of a configuration metadata file:

```
\samples\plugins\HostSample\stage\oms\metadata\snapshotlive\demo_hostsample_ecmdef.xml
```

For more information, see [Section 6.3.1, "Defining Configuration Collection Tables"](#).

- The target type metadata file

Each configuration data item is collected as a metric that is defined with other metrics in the target type metadata file. In addition, the EDK includes an example of a target type metadata file:

```
\samples\plugins\HostSample\stage\oms\metadata\targetType\demo_hostsample.xml
```

For more information, see [Section 3.3, "Creating the Target Type Metadata File"](#).

- The default collection metadata file

The frequency at which configuration data is collected for each metric is defined in the default collection metadata file. Metric Alert event conditions for each metric and the messages to display for such alerts are also defined in this file. In addition, the EDK includes an example of a default collection metadata file:

```
\samples\plugins\HostSample\stage\oms\metadata\default_collection\demo_hostsample.xml
```

For more information, see [Section 3.5, "Creating the Default Collection File"](#).

6.3 Modeling Enterprise Configuration Management Tables

This section describes how to create configuration snapshot tables. Assume that an Oracle home target type definition is added to the Enterprise Manager framework.

You can define metadata and tables to store your configuration data and define metrics to collect the data at the Management Agent.

Each Enterprise Configuration Management metric defined in the target type metadata file corresponds to a table defined by the configuration metadata XML file. For information about target type metadata files, see [Chapter 3, "Creating Target Metadata Files"](#).

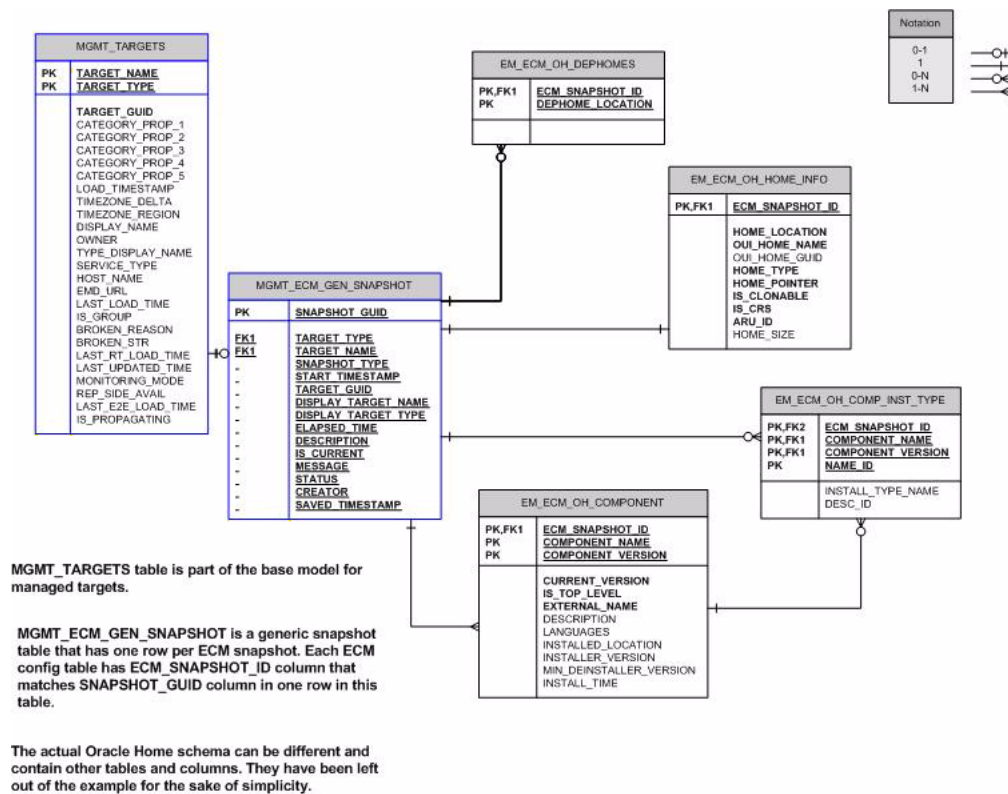
Tables can have parent-child relationships similar to foreign key constraints in the database. For every parent table row, there are n numbers of child table rows, and every child table row has exactly one parent table row. In effect, Enterprise Configuration Management allows trees of tables, where each table has at most one parent table. A configuration can consist of a set of these trees where no table is repeated.

From a database perspective, a table must include all key columns of its ancestor tables. (Internally, each Enterprise Configuration Management table includes an ECM_SNAPSHOT_ID column of type RAW(16) to identify the snapshot to which a given row belongs.)

Metadata for the collection tables must specify at a minimum:

- Table names, column names, column types, and hierarchical relationships between tables (if any).
- Key columns, that is columns that uniquely identify a row in a table. You must specify keys correctly. Enterprise Manager UI, comparison, and history use key columns extensively.
- UI display names that will in the Enterprise Manager generic UI.

[Figure 6–1](#) shows the Oracle Home Entity Relationship Diagram (ERD) of the configuration tables.

Figure 6–1 Oracle Home ERD with Tie-ins to the Framework

6.3.1 Defining Configuration Collection Tables

The configuration metadata XML file begins with the METADATA tag. The METADATA tag establishes the relationship between the target type and the snapshot type, and defines the UI display name.

```
<METADATAS>
  <METADATA SNAP_TYPE="oracle_home_config" TARGET_TYPE="oracle_home" VER="1">

    <METADATA_UI_NAME>Oracle Home Configuration</METADATA_UI_NAME>

    CONFIGURATION COLLECTION TABLE DEFINITIONS
  </METADATA>
</METADATAS>
```

The configuration collection table definitions are encapsulated within the METADATA tag.

For more information about the elements of the configuration metadata XML, see [Section 6.3.2, "Overview of Configuration Management Snapshot Metadata Elements"](#). For information about the XML Schema Definitions (XSD) that governs the configuration XML files, see the Extensibility Development Kit (EDK) specifications.

6.3.1.1 EM_ECM_OH_HOME_INFO Table

The EM_ECM_OH_HOME_INFO table stores properties related to an Oracle software installation. The metadata for this table is as follows:

```
<TABLE NAME="EM_ECM_OH_HOME_INFO" SINGLE_ROW="Y">
  <UI_NAME>Home Info</UI_NAME>
  <COLUMN NAME="HOME_LOCATION" TYPE="STRING" TYPE_FORMAT="1024">Install
```

```

Location</COLUMN>
  <COLUMN NAME="OUI_HOME_NAME" TYPE="STRING" TYPE_FORMAT="256" >OUI Home
Name</COLUMN>
  <COLUMN NAME="OUI_HOME_GUID" TYPE="STRING" TYPE_FORMAT="32" >OUI Home
GUID</COLUMN>
  <COLUMN NAME="HOME_TYPE" TYPE="STRING" TYPE_FORMAT="1" >Home
Type</COLUMN>
  <COLUMN NAME="HOME_POINTER" TYPE="STRING" TYPE_FORMAT="1024">Home
Pointer</COLUMN>
  <COLUMN NAME="IS_CLONABLE" TYPE="NUMBER" TYPE_FORMAT="1" >Is
Clonable</COLUMN>
  <COLUMN NAME="IS_CRS" TYPE="NUMBER" TYPE_FORMAT="1" >Is CRS</COLUMN>
  <COLUMN NAME="ARU_ID" TYPE="NUMBER" TYPE_FORMAT="10" >ARU ID of the
Oracle Home</COLUMN>
  <COLUMN NAME="HOME_SIZE" TYPE="NUMBER" TYPE_FORMAT="10" >Size of Oracle
Home (KB)</COLUMN>
</TABLE>

```

The corresponding database table is as follows:

COLUMN_NAME	DATA_TYPE	COLUMN_WIDTH	KEY
ECM_SNAPSHOT_ID	RAW	16	PK
HOME_LOCATION	VARCHAR2	1024	
OUI_HOME_NAME	VARCHAR2	256	
OUI_HOME_GUID	VARCHAR2	32	
HOME_TYPE	VARCHAR2	1	
HOME_POINTER	VARCHAR2	1024	
IS_CLONABLE	NUMBER	1	
IS_CRS	NUMBER	1	
ARU_ID	NUMBER	10	
HOME_SIZE (in KB)	NUMBER	10	

For this table, the Primary Key consists of ECM_SNAPSHOT_ID. The SINGLE_ROW="Y" attribute from the metadata example indicates that each snapshot will have at most a single row in this table and therefore you do not have to mark any other columns in the table as key. This implies that the ECM_SNAPSHOT_ID column, which identifies the snapshot, will be the only key column in the table.

6.3.1.2 EM_ECM_OH_DEP_HOMES Table

The EM_ECM_OH_DEP_HOMES table stores the locations of the Oracle home directories that a given Oracle home depends on. This data is used to form dependency associations between Oracle homes. The metadata for this table is as follows:

```

<TABLE NAME="EM_ECM_OH_DEP_HOMES">
  <UI_NAME>Dependee Oracle Homes</UI_NAME>
  <COLUMN NAME="DEP_HOME_LOCATION" TYPE="STRING" TYPE_FORMAT="1024" IS_
KEY="Y">Dependee Home Location</COLUMN>
</TABLE>

```

The corresponding database table is as follows:

COLUMN_NAME	DATA_TYPE	COLUMN_WIDTH	KEY
ECM_SNAPSHOT_ID	RAW	16	PK
DEP_HOME_LOCATION	VARCHAR2	1024	PK

The primary key for this table consists of ECM_SNAPSHOT_ID and DEP_HOME_LOCATION.

Note: Key columns (in addition to ECM_SNAPSHOT_ID, which is always part of the key) are marked with IS_KEY="Y" in the metadata file.

6.3.1.3 EM_ECM_OH_COMPONENT Table

The EM_ECM_OH_COMPONENT table stores information about Oracle software components in the Oracle home directory. The metadata for this table is as follows:

```
<TABLE NAME="EM_ECM_OH_COMPONENT">
  <UI_NAME>Components installed in Oracle Home</UI_NAME>
  <COLUMN NAME="COMPONENT_NAME"          TYPE="STRING" TYPE_FORMAT="128" IS_
KEY="Y">Component Name</COLUMN>
  <COLUMN NAME="COMPONENT_VERSION"        TYPE="STRING" TYPE_FORMAT="64" IS_
KEY="Y">Base Version of Component</COLUMN>
  <COLUMN NAME="CURRENT_VERSION"          TYPE="STRING" TYPE_FORMAT="64"
>Current Version of the Component</COLUMN>
  <COLUMN NAME="INSTALL_TIME"              TYPE="TIMESTAMP"
>Install Time</COLUMN>
  <COLUMN NAME="IS_TOP_LEVEL"              TYPE="NUMBER" TYPE_FORMAT="1"
>Is it a top level Component</COLUMN>
  <COLUMN NAME="EXTERNAL_NAME"             TYPE="STRING" TYPE_FORMAT="128"
>External name</COLUMN>
  <COLUMN NAME="DESCRIPTION"               TYPE="STRING" TYPE_FORMAT="1024"
>Description</COLUMN>
  <COLUMN NAME="LANGUAGES"                 TYPE="STRING" TYPE_FORMAT="1024"
>Languages</COLUMN>
  <COLUMN NAME="INSTALLED_LOCATION"        TYPE="STRING" TYPE_FORMAT="1024"
>Installed Location</COLUMN>
  <COLUMN NAME="INSTALLER_VERSION"         TYPE="STRING" TYPE_FORMAT="64"
>Installer Version</COLUMN>
  <COLUMN NAME="MIN_DEINSTALLER_VERSION"   TYPE="STRING" TYPE_FORMAT="64"
>Minimum Deinstaller Version</COLUMN>
```

Note: This metadata has no closing TABLE tag (yet) because it is a parent table to the next table ([EM_ECM_OH_COMP_INST_TYPE Table](#)), which is included as part of this TABLE tag.

The corresponding database table is as follows:

COLUMN_NAME	DATA_TYPE	COLUMN_WIDTH	KEY
ECM_SNAPSHOT_ID	RAW	16	PK
COMPONENT_NAME	VARCHAR2	128	PK
COMPONENT_VERSION	VARCHAR2	64	PK

COLUMN_NAME	DATA_TYPE	COLUMN_WIDTH	KEY
CURRENT_VERSION	VARCHAR2	64	
INSTALL_TIME	DATE	-	
IS_TOP_LEVEL	NUMBER	1	
EXTERNAL_NAME	VARCHAR2	128	
DESCRIPTION	VARCHAR2	1024	
LANGUAGES	VARCHAR2	1024	
INSTALLED_LOCATION	VARCHAR2	1024	
INSTALLER_VERSION	VARCHAR2	64	
MIN_DEINSTALLER_VERSION	VARCHAR2	64	

The primary key for this table consists of ECM_SNAPSHOT_ID, COMPONENT_NAME, and COMPONENT_VERSION.

6.3.1.4 EM_ECM_OH_COMP_INST_TYPE Table

The EM_ECM_OH_COMP_INST_TYPE table stores the installation type of Oracle software components (Oracle Universal Installer (OUI) only). This is a child table of [EM_ECM_OH_COMPONENT Table](#). The metadata for this table is as follows:

```
<TABLE NAME="EM_ECM_OH_COMP_INST_TYPE">
  <UI_NAME>Install Types of Components</UI_NAME>
  <COLUMN NAME="NAME_ID"                                TYPE="STRING" TYPE_FORMAT="128" IS_
KEY="Y">Install Type Name ID</COLUMN>
  <COLUMN NAME="INSTALL_TYPE_NAME"                       TYPE="STRING" TYPE_FORMAT="128"
>Install Type Name</COLUMN>
  <COLUMN NAME="DESC_ID"                                TYPE="STRING" TYPE_FORMAT="128"
>Install Type Description ID</COLUMN>
</TABLE>
</TABLE>
```

Note: The extra closing </TABLE> tag is for the EM_ECM_OH_COMPONENT parent table.

The corresponding database table is as follows:

COLUMN_NAME	DATA_TYPE	COLUMN_WIDTH	KEY
ECM_SNAPSHOT_ID	RAW	16	PK
COMPONENT_NAME	VARCHAR2	128	PK
COMPONENT_VERSION	VARCHAR2	64	PK
NAME_ID	VARCHAR2	128	PK
INSTALL_TYPE_NAME	VARCHAR2	128	
DESC_ID	VARCHAR2	128	

For this table, the primary key consists of ECM_SNAPSHOT_ID, COMPONENT_NAME, COMPONENT_VERSION, and NAME_ID.

Note: COMPONENT_NAME and COMPONENT_VERSION are not listed here but are inherited from the key columns of the parent table, and the values in these columns must match a row in EM_ECM_OH_COMPONENT with the same values for ECM_SNAPSHOT_ID, COMPONENT_NAME, and COMPONENT_VERSION.

6.3.1.5 Additional Information About the Configuration Metadata

Note the following when you are creating the configuration metadata XML file:

- Each table definition must specify its name and at least one column specification:

Example 6–1 Table Definition

```
<TABLE NAME="..." ...>  
Optional UI name  
Column definitions  
Optional Indexes definitions: starting 12c PS1 platform release  
Optional Child Table definitions  
</TABLE>
```

[Table 6–1](#) provides the attributes of the TABLE element.

- Table names must be specific to a given snapshot type and cannot be shared by multiple snapshot types. While internal Enterprise Manager names can start with EM_, for plug-ins and addons, Oracle recommends that you start table names with (upper-cased) plug-in tag (third part of a plugin ID) followed by underscore. For example, plugin “oracle.sysman.xyz” should define tables starting with XYZ_, such as XYZ_CONFIG.
- Each column definition must specify its name and type at least:

Example 6–2 Column Definition

```
<COLUMN NAME="..." TYPE="..." ...>...</COLUMN>
```

[Table 6–1](#) provides the attributes of the COLUMN element.

- Columns can be of type STRING, NUMBER, TIMESTAMP, or RAW. TYPE_FORMAT is optional; its meaning derives from the value of type. For a string, it is the maximum string length. For a number, it is precision and scale, as for an Oracle database (for example, TYPE_FORMAT="4, 9").
- Specify parent child relationships between tables by nesting the TABLE tags.
- Table and view names cannot exceed 25 characters.
- Table and column order is significant. The UI display replicates the order. Import and export operations preserve the order. Delete operations (on table data) occur in inverse order. Child rows are removed before parent rows.
- COLUMN tags contain the UI display name of the column.
- Tables require key columns that uniquely identify rows in the table. Oracle recommends that the total size of the key columns is not too large (4,000 should be the maximum but much smaller sizes are acceptable).

The key columns of all ancestor tables are automatically assumed to be inherited by the child tables and are not repeated in the child table specification.

However, a table does not require a key if it has at most one row per parent row, or per snapshot in the case of a top-level table. Tables that do not have a key must specify the `SINGLE_ROW="Y"` attribute, which is set to "N" by default.

For information about the key elements of the configuration metadata, see [Section 6.3.2, "Overview of Configuration Management Snapshot Metadata Elements"](#).

6.3.2 Overview of Configuration Management Snapshot Metadata Elements

[Table 6–1](#) describes the key elements that define configuration management.

Table 6–1 Key Elements of a Configuration Metadata XML File

Element	Description
METADATAS	<p>The configuration metadata XML file starts with the METADATAS tag.</p> <pre><METADATAS> One or more snapshot specifications </METADATAS></pre>
METADATA	<p>The snapshot specification corresponds to a METADATA tag.</p> <p>The snapshot specification corresponds to a METADATA tag and includes at least one table specification. It also defines the snapshot UI display name. It includes the following attributes:</p> <ul style="list-style-type: none"> ■ VER: Specifies the plug-in developer-defined metadata version and must be an integer (beginning with 1). Typically, each release increments the version (if there are changes). Only the latest ECM metadata version must be registered within a release (although Management Agents can upload older compatible versions) ■ SNAP_TYPE: Names the type of snapshot. Snapshot types are defined in the context of target types. The name should begin with company name followed by an underscore. For example, oracle_. ■ TARGET_TYPE: Target type for which the snapshot type metadata is defined. ■ UI_IGNORE: Optional. Determines whether data is displayed as part of the configuration browser. Values are N (default), which uses UI functionality or Y, which ignores UI functionality. ■ COMPARE_IGNORE: Optional. Determines whether to perform comparisons on the data. Values are N (default), which uses Compare functionality or Y, which ignores Compare functionality. ■ COMPARE_UI_IGNORE: Optional. Determines whether to display data in the comparison UI. Values are N (default), which uses Compare UI functionality or Y, which ignores Compare UI functionality. ■ HISTORY_IGNORE: Optional. Determines whether to track history on the data. Values are N (default), which uses History functionality or Y, which ignores History functionality. ■ HISTORY_UI_IGNORE: Optional. Determines whether to display data in the history UI. Values are N (default), which uses History UI functionality or Y, which ignores History UI functionality. <p>Note: If you specify a value for a <i>name_IGNORE</i> attribute, then it is specified for all tables unless overridden at a lower (TABLE or COLUMN) level. Inheritance flows from metadata to tables, parent tables to child tables, and from all tables to their columns.</p>

Table 6–1 (Cont.) Key Elements of a Configuration Metadata XML File

Element	Description
TABLE	<p>Specifies the table name and at least one column. It includes the following attributes:</p> <ul style="list-style-type: none"> ■ NAME: Required. Identifies the table uniquely ■ SINGLE_ROW: Values are N (default) or Y, which indicates tables that have at most one row per parent row or at most one row per snapshot in case of top-level tables. In this latter case, no key is required. All key columns (if any) are inherited from a parent table. ■ UI_IGNORE: Optional. Determines whether data is displayed as part of the configuration browser. Values are N (default), which uses UI functionality or Y, which ignores UI functionality. ■ COMPARE_IGNORE: Optional. Determines whether to perform comparisons on the data. Values are N (default), which uses Compare functionality or Y, which ignores Compare functionality. ■ COMPARE_UI_IGNORE: Optional. Determines whether to display data in the comparison UI. Values are N (default), which uses Compare UI functionality or Y, which ignores Compare UI functionality. ■ HISTORY_IGNORE: Optional. Determines whether to track history on the data. Values are N (default), which uses History functionality or Y, which ignores History functionality. ■ HISTORY_UI_IGNORE: Optional. Determines whether to display data in the history UI. Values are N (default), which uses History UI functionality or Y, which ignores History UI functionality. <p>Note: If you specify a value for a <i>name_IGNORE</i> attribute, then it is specified for all columns and a child table unless overridden at a lower (TABLE or COLUMN) level.</p>

Table 6–1 (Cont.) Key Elements of a Configuration Metadata XML File

Element	Description
COLUMN	<p>Each column definition must a <code>NAME</code> and <code>TYPE</code> attributes. It includes the following attributes:</p> <ul style="list-style-type: none"> ■ <code>NAME</code>: Required. Identifies the table column uniquely. ■ <code>TYPE</code>: Required. Values are <code>STRING</code>, <code>NUMBER</code>, <code>TIMESTAMP</code>, <code>CLOB</code>, or <code>RAW</code> ■ <code>TYPE_FORMAT</code>: Optional. Value depends on the value in the <code>TYPE</code> attribute ■ <code>IS_KEY</code>: Required. Specifies whether it is part of a key that uniquely identifies a row in the table. The key columns of all ancestor tables are automatically assumed to be inherited by the child tables. <p>Note: A table does not require a key (unless that of ancestor table keys, if any) if it has only one row per parent row, or per snapshot in the case of a top-level table. You must specify the <code>SINGLE_ROW="Y"</code> attribute (which is set to <code>"N"</code> by default) for tables that do not have a key.</p> <ul style="list-style-type: none"> ■ <code>UI_IGNORE</code>: Optional. Determines whether data is displayed as part of the configuration browser. Values are <code>N</code> (default), which uses UI functionality or <code>Y</code>, which ignores UI functionality. ■ <code>COMPARE_IGNORE</code>: Optional. Determines whether to perform comparisons on the data. Values are <code>N</code> (default), which uses Compare functionality or <code>Y</code>, which ignores Compare functionality. ■ <code>COMPARE_UI_IGNORE</code>: Optional. Determines whether to display data in the comparison UI. Values are <code>N</code> (default), which uses Compare UI functionality or <code>Y</code>, which ignores Compare UI functionality. ■ <code>HISTORY_IGNORE</code>: Optional. Determines whether to track history on the data. Values are <code>N</code> (default), which uses History functionality or <code>Y</code>, which ignores History functionality. ■ <code>HISTORY_UI_IGNORE</code>: Optional. Determines whether to display data in the history UI. Values are <code>N</code> (default), which uses History UI functionality or <code>Y</code>, which ignores History UI functionality.
INDEXES	<p>Optional indexes are useful when a table is used in derived associations, compliance rules, reports, or other queries where performance is important in accessing table rows based on columns other than <code>ECM_SNAPSHOT_ID</code>.</p> <p>An optional index definitions element <code><INDEXES></code> should specify at least one <code><INDEX></code> element, and each <code><INDEX></code> element should specify name and columns:</p> <pre> <INDEXES> <INDEX NAME="..." COLUMNS="..." /> <INDEX NAME="..." COLUMNS="..." /> ... </INDEXES> </pre> <ul style="list-style-type: none"> ■ <code>NAME</code>: Required. Identifies the table column uniquely and should be unique among all index names. Oracle recommends that the name should be the same as table name followed by <code>_IDX1</code>, <code>_IDX2</code>, and so on. ■ <code>COLUMN</code>: Required. Value should be a comma-separated list of column names.

Note: A predefined ECM_METADATA_ID column is always added as the last column for each index.

Example 6–3 provides an example of a configuration metadata XML file for an oracle_home_config snapshot. All *_IGNORE attributes keep their default values (N) because they are not specified in the file.

Example 6–3 Example of a Configuration Metadata XML File

```
<METADATAS>
  <METADATA SNAP_TYPE="oracle_home_config" TARGET_TYPE="oracle_home" VER="1"

  <METADATA_UI_NAME>Oracle Home Configuration</METADATA_UI_NAME>

  <TABLE NAME="EM_ECM_OH_HOME_INFO" SINGLE_ROW="Y">
    <UI_NAME>Home Info</UI_NAME>
    <COLUMN NAME="HOME_LOCATION" TYPE="STRING" TYPE_FORMAT="1024">Install
Location</COLUMN>
    <COLUMN NAME="OUI_HOME_NAME" TYPE="STRING" TYPE_FORMAT="256">OUI Home
Name</COLUMN>
    <COLUMN NAME="OUI_HOME_GUID" TYPE="STRING" TYPE_FORMAT="32">OUI Home
GUID</COLUMN>
    <COLUMN NAME="HOME_TYPE" TYPE="STRING" TYPE_FORMAT="1">Home Type</COLUMN>
    <COLUMN NAME="HOME_POINTER" TYPE="STRING" TYPE_FORMAT="1024">Home
Pointer</COLUMN>
    <COLUMN NAME="IS_CLONABLE" TYPE="NUMBER" TYPE_FORMAT="1">Is Clonable</COLUMN>
    <COLUMN NAME="IS_CRS" TYPE="NUMBER" TYPE_FORMAT="1">Is CRS</COLUMN>
    <COLUMN NAME="ARU_ID" TYPE="NUMBER" TYPE_FORMAT="10">ARU ID of the Oracle
Home</COLUMN>
    <COLUMN NAME="HOME_SIZE" TYPE="NUMBER" TYPE_FORMAT="10">Size of Oracle Home
(KB) </COLUMN>
  </TABLE>

  <TABLE NAME="EM_ECM_OH_DEP_HOMES">
    <UI_NAME>Dependee Oracle Homes</UI_NAME>
    <COLUMN NAME="DEP_HOME_LOCATION" TYPE="STRING" TYPE_FORMAT="1024" IS_
KEY="Y">Dependee Home Location</COLUMN>
  </TABLE>

  <TABLE NAME="EM_ECM_OH_COMPONENT">
    <UI_NAME>Components installed in Oracle Home</UI_NAME>
    <COLUMN NAME="COMPONENT_NAME" TYPE="STRING" TYPE_FORMAT="128" IS_
KEY="Y">Component Name</COLUMN>
    <COLUMN NAME="COMPONENT_VERSION" TYPE="STRING" TYPE_FORMAT="64" IS_
KEY="Y">Base Version of Component</COLUMN>
    <COLUMN NAME="CURRENT_VERSION" TYPE="STRING" TYPE_FORMAT="64">Current Version
of the Component</COLUMN>
    <COLUMN NAME="INSTALL_TIME" TYPE="TIMESTAMP" >Install Time</COLUMN>
    <COLUMN NAME="IS_TOP_LEVEL" TYPE="NUMBER" TYPE_FORMAT="1">Is a top level
Component</COLUMN>
    <COLUMN NAME="EXTERNAL_NAME" TYPE="STRING" TYPE_FORMAT="128">External
name</COLUMN>
    <COLUMN NAME="DESCRIPTION" TYPE="STRING" TYPE_
FORMAT="1024">Description</COLUMN>
    <COLUMN NAME="LANGUAGES" TYPE="STRING" TYPE_FORMAT="1024" >Languages</COLUMN>
    <COLUMN NAME="INSTALLED_LOCATION" TYPE="STRING" TYPE_FORMAT="1024">Installed
Location</COLUMN>
    <COLUMN NAME="INSTALLER_VERSION" TYPE="STRING" TYPE_FORMAT="64">Installer
```

```

Version</COLUMN>
  <COLUMN NAME="MIN_DEINSTALLER_VERSION" TYPE="STRING" TYPE_FORMAT="64">Minimum
Deinstaller Version</COLUMN>

  <TABLE NAME="EM_ECM_OH_COMP_INST_TYPE">
    <UI_NAME>Install Types of Components</UI_NAME>
    <COLUMN NAME="NAME_ID" TYPE="STRING" TYPE_FORMAT="128" IS_KEY="Y">Install
Type Name ID</COLUMN>
    <COLUMN NAME="INSTALL_TYPE_NAME" TYPE="STRING" TYPE_FORMAT="128" >Install
Type Name</COLUMN>
    <COLUMN NAME="DESC_ID" TYPE="STRING" TYPE_FORMAT="128">Install Type
Description</COLUMN>
  </TABLE>

</TABLE>

</METADATA>
</METADATAS>

```

Note: Use the information in this file to define a metric in the target type metadata file. For more information, see [Section 6.3.7, "Modifications to Standard Collection Metrics and RAW Metrics"](#) and [Example 6-6](#).

6.3.3 Packaging Configuration Metadata

When you have completed your configuration metadata XML file, save the file in the following directory in your plug-in staging directory:

```
oms/metadata/snapshotlive
```

After you save the configuration metadata file in the plug-in staging directory, it is available for automatic registration during plug-in deployment.

For information about the plug-in staging directory or plug-in deployment, see [Chapter 13, "Validating, Packaging, and Deploying the Plug-in"](#).

6.3.4 Registering Metadata With the Configuration Management Framework

Once you complete the ECM metadata file, you must register it with the ECM framework. The Metadata XML file should be placed in the oms/metadata/snapshotlive directory in your plug-in stage directory. Registration will automatically take place during the install/upgrade of the plug-in and the necessary ECM schema objects will be created and the ECM metadata registered with the ECM framework.

To manually run the registration service during development of XML-only plug-ins, execute the following:

```

emctl register oms metadata
-sysman_pwd <sysman password>
-pluginId <three-part plugin ID>
-service LiveSnapshotRegistration
-file <snapshot metadata XML file>

```

The <three-part plugin ID> is ID of the plug-in to which the snapshot metadata target types belong. The plug-in should be imported into the repository before the command is run. The LiveSnapshotRegistration is the metadata service ID for

the snapshot metadata registration for metadata-only plug-ins. This will register the snapshot configuration file (<snapshot metadata XML file>) using SYSMAN credentials in your development environment.

ECM provides the following utility that can be run optionally to generate some additional files:

```
<EDK_DIR>/bin/generate_ecm_resources.sh  
<snapshot metadata XML file>  
<destination directory>
```

where <EDK_DIR> is the directory where EDK package is unzipped. Note that generate_ecm_resources.bat is available in the same location for a Windows based environment.

Given an input ECM metadata <name>.xml file, generated files will be located in the destination directory, with all file names starting with lowercase <name> as the prefix. The following files are generated:

- <lowercase name>_metric.xml and <lowercase name>_collection.xml can be used as starting templates for Agent-side integration.
- <lowercase name>.dlf is a file for translations of snapshot type name, table name, and column name. A generated DLF file and its translated versions should be placed in the <stage>/oms/rsc/ecm directory of the management plug-in.
- Generated <lowercase name>*.sql files should be ignored.

6.3.5 Supporting Translation

Data Loading Format (DLF) translation files are used to support internationalization for the display strings in plug-in metadata files so that the UI can display them in the language of the end user. Usually, you provide an original (such as English) DLF file to translators who then create similar files for other locales. All such files are loaded into the Management Repository by Enterprise Manager during the installation of a plug-in.

Note: Generating DLF files is optional and is required only if you require translation.

To generate DLF files for translation of the snapshot type name, table, and column names, run the following command:

```
$ORACLE_HOME/emcore/pdk/partner/bin/empdk generate_metadata_resource  
-stage_dir staging directory  
-service LiveSnapshotRegistration  
-input_dir input directory containing snapshot metadata XML file  
-file_extension extension to use for generated DLF files  
-out_dir output directory to which to generate the DLF files  
[-debug debug output file]
```

You must place the generated DLF files and their corresponding translated versions in the following directory of the plug-in:

```
plugin_stage/oms/rsc/ecm
```

The following manual additions are required in each DLF file:

```
<table xml:lang="en" name="MGMT_MESSAGES">  
<lookup-key>
```

```

    <column name="MESSAGE_ID" />
    <column name="SUBSYSTEM" />
    <column name="LANGUAGE_CODE" />
    <column name="COUNTRY_CODE" />
  </lookup-key>
</columns>
  <column name="MESSAGE_ID" type="string" maxsize="256" />
  <column name="SUBSYSTEM" type="string" maxsize="64" />
  <column name="LANGUAGE_CODE" type="string" language="%l" />
  <column name="COUNTRY_CODE" type="string" language="%Cs" />
  <column name="MESSAGE" type="string" maxsize="1000" translate="yes" />
</columns>
<dataset>
  GENERATED FILE CONTENT
</dataset>
</table>

```

6.3.6 Upgrading Configuration Data

Note: You must increment your integer metadata version (VER attribute) whenever you release a new version of metadata to your customers. The version should be incremented in the ECM XML metadata file (VER attribute in the METADATA element) as well as the corresponding Agent collection file.

For more information about the VER attribute, see [Table 6–1, "Key Elements of a Configuration Metadata XML File"](#).

When you are upgrading existing snapshot metadata, the following changes are supported:

- New tables
- New non-key columns (these should appear after previously existing columns)
- New list of indexes (replaces any prior indexes)
- Increasing length of STRING type columns
- Values of UL_IGNORE, COMPARE_IGNORE, COMPARE_UL_IGNORE, HISTORY_IGNORE, HISTORY_UI_IGNORE and UI_NAME attributes

Note: ■ No key columns can be added or removed

- No columns formats except STRING can be altered
 - The length of the STRING columns only can be increased
 - New table cannot be added as a parent for an existing table
 - Tables or columns cannot be removed from the existing snapshot metadata
-

Take the following metadata example:

Example 6–4 Original Metadata Definition

```
<METADATAS xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
```

```

<METADATA VER="1"
  SNAP_TYPE="sn_dbconfig"
  TARGET_TYPE="sn_oracle_database"
  UI_IGNORE="Y"
  HISTORY_IGNORE="N"
  COMPARE_IGNORE="Y"
  COLLECTION_GROUP="COL_GRP_0">
<METADATA_UI_NAME>Database Configuration</METADATA_UI_NAME>
<TABLE NAME="TABLESPACES" DATA_SOURCE="R">
  <UI_NAME>Tablespaces</UI_NAME>
  <COLUMN NAME="TABLESPACE_NAME" TYPE="STRING" TYPE_FORMAT="30"
    IS_KEY="Y">Tablespace Name</COLUMN>
  <COLUMN NAME="SIZE" TYPE="NUMBER">Size</COLUMN>
  <COLUMN NAME="STATUS" TYPE="STRING" TYPE_FORMAT="10">Status</COLUMN>
<TABLE NAME="DATAFILES">
  <UI_NAME>Datafiles</UI_NAME>
  <COLUMN NAME="FILE_NAME" TYPE="STRING" TYPE_FORMAT="1024"
    IS_KEY="Y">File Name</COLUMN>
  <COLUMN NAME="FILE_SIZE" TYPE="NUMBER" HISTORY_IGNORE="Y">Size</COLUMN>
  <COLUMN NAME="STATUS" TYPE="STRING" TYPE_FORMAT="9">Status</COLUMN>
</TABLE>
</TABLE>
</METADATA>
</METADATAS>

```

Example 6-5 provides an example of an upgrade to the metadata definition described in **Example 6-4**. The changes are highlighted in bold text.

Example 6-5 Upgraded Metadata Definition

```

<METADATAS xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
  <METADATA VER="2"
    SNAP_TYPE="sn_dbconfig"
    TARGET_TYPE="sn_oracle_database"
    UI_IGNORE="Y"
    HISTORY_IGNORE="N"
    COMPARE_IGNORE="Y"
    COLLECTION_GROUP="COL_GRP_0">
  <METADATA_UI_NAME>Database Configuration</METADATA_UI_NAME>
  <TABLE NAME="TABLESPACES" DATA_SOURCE="R">
    <UI_NAME>Tablespaces</UI_NAME>
    <COLUMN NAME="TABLESPACE_NAME" TYPE="STRING" TYPE_FORMAT="30"
      IS_KEY="Y">Tablespace Name</COLUMN>
    <COLUMN NAME="SIZE" TYPE="NUMBER">Size</COLUMN>
    <COLUMN NAME="STATUS" TYPE="STRING" TYPE_FORMAT="10">Status</COLUMN>
  <TABLE NAME="DATAFILES">
    <UI_NAME>Datafiles</UI_NAME>
    <COLUMN NAME="FILE_NAME" TYPE="STRING" TYPE_FORMAT="1024"
      IS_KEY="Y">File Name</COLUMN>
    <COLUMN NAME="FILE_SIZE" TYPE="NUMBER" HISTORY_IGNORE="Y">Size</COLUMN>
    <COLUMN NAME="STATUS" TYPE="STRING" TYPE_FORMAT="10">Status</COLUMN>
    <COLUMN NAME="DESC" TYPE="STRING" TYPE_FORMAT="128">Status</COLUMN>
  </TABLE>
</TABLE>
  <TABLE NAME="TABLESPACES_1" DATA_SOURCE="R" >
    <UI_NAME>Tablespaces_1</UI_NAME>
    <COLUMN NAME="TABLESPACE_NAME" TYPE="STRING" TYPE_FORMAT="30"
      IS_KEY="Y">Tablespace Name</COLUMN>
    <COLUMN NAME="SIZE" TYPE="NUMBER">Size</COLUMN>
    <COLUMN NAME="STATUS" TYPE="STRING" TYPE_FORMAT="10">Status</COLUMN>

```

```

<COLUMN NAME="NOTES" TYPE="STRING" TYPE_FORMAT="128">Status</COLUMN>
<TABLE NAME="DATAFILES_1">
  <UI_NAME>Datafiles_1</UI_NAME>
  <COLUMN NAME="FILE_NAME" TYPE="STRING" TYPE_FORMAT="1024"
    IS_KEY="Y">File Name</COLUMN>
  <COLUMN NAME="FILE_SIZE" TYPE="NUMBER" HISTORY_IGNORE="Y">Size</COLUMN>
  <COLUMN NAME="STATUS" TYPE="STRING" TYPE_FORMAT="10">Status</COLUMN>
</TABLE>
</TABLE>
</METADATA>
</METADATAS>

```

Older Enterprise Manager releases allowed for the "DROP_EXISTING_DATA" attribute in the METADATA element, but this should be removed as it is no longer supported.

Note that if you must have non-backward compatible changes in your ECM metadata, you must create a new snapshot type (that will not be comparable with old snapshot type). ECM only supports backward compatible changes.

6.3.7 Modifications to Standard Collection Metrics and RAW Metrics

Enterprise Configuration Management data is collected by regular metrics, collections and Management Agent mechanisms. This data is collected through regular RAW metrics with the following modifications:

- Add a CONFIG="TRUE" attribute to all Metric and CollectionItem tags that collect configuration snapshot information.
- Ensure that the CollectionItem NAME attribute is the same as the snapshot type name (for example, oracle_home_config).

Note: Do not include ECM_SNAPSHOT_ID as a column in any RAW metric table descriptor

When ancestor key columns are included in child tables, you can populate a hierarchical set of tables one at a time, without having to express the hierarchical relationships during the collection. You should list parent tables before corresponding child tables in the Collection Item.

[Example 6-6](#) provides a metric definition from the target type metadata XML file for the EM_ECM_OH_HOME_INFO table defined in [Example 6-3](#). The information highlighted in bold font is provided by the configuration metadata XML file ([Example 6-3](#)).

Note: [Example 6-6](#) is just an example and additional nonconfiguration-specific Management Agent attributes may be required for your situation.

For more information about the target type metadata XML file, see [Section 3.3, "Creating the Target Type Metadata File"](#).

Example 6-6 Defining a Metric

```

<Metric NAME="EM_ECM_OH_HOME_INFO" TYPE="RAW" CONFIG="TRUE">
  <Display>

```

```

        <Label NLSID="...">Home Info</Label>
    </Display>
    <TableDescriptor TABLE_NAME="EM_ECM_OH_HOME_INFO">
        <ColumnDescriptor NAME="HomeLocation" COLUMN_NAME="HOME_LOCATION"
TYPE="S">
            <Display>
                <Label NLSID="...">Install Location</Label>
            </Display>
        </ColumnDescriptor>
        <ColumnDescriptor NAME="OuiHomeName" COLUMN_NAME="OUI_HOME_NAME"
TYPE="S">
            <Display>
                <Label NLSID="...">OUI Home Name</Label>
            </Display>
        </ColumnDescriptor>
        <ColumnDescriptor NAME="OuiHomeGuid" COLUMN_NAME="OUI_HOME_GUID"
TYPE="S">
            <Display>
                <Label NLSID="...">OUI Home GUID</Label>
            </Display>
        </ColumnDescriptor>
        <ColumnDescriptor NAME="HomeType" COLUMN_NAME="HOME_TYPE" TYPE="S">
            <Display>
                <Label NLSID="...">Home Type</Label>
            </Display>
        </ColumnDescriptor>
        <ColumnDescriptor NAME="HomePointer" COLUMN_NAME="HOME_POINTER"
TYPE="S">
            <Display>
                <Label NLSID="...">Home Pointer</Label>
            </Display>
        </ColumnDescriptor>
        <ColumnDescriptor NAME="IsClonable" COLUMN_NAME="IS_CLONABLE"
TYPE="N">
            <Display>
                <Label NLSID="...">Is Clonable</Label>
            </Display>
        </ColumnDescriptor>
        <ColumnDescriptor NAME="IsCrS" COLUMN_NAME="IS_CRS" TYPE="N">
            <Display>
                <Label NLSID="...">Is CRS</Label>
            </Display>
        </ColumnDescriptor>
        <ColumnDescriptor NAME="AruId" COLUMN_NAME="ARU_ID" TYPE="N">
            <Display>
                <Label NLSID="...">ARU ID of the Oracle Home</Label>
            </Display>
        </ColumnDescriptor>
        <ColumnDescriptor NAME="HomeSize" COLUMN_NAME="HOME_SIZE" TYPE="N">
            <Display>
                <Label NLSID="...">Size of Oracle Home (KB)</Label>
            </Display>
        </ColumnDescriptor>
    </TableDescriptor>
    <!-- TODO : EDIT: Edit the QueryDescriptor Element Template as required
-->
    <QueryDescriptor FETCHLET_ID=" " AGENT_MODE=" ">
        <Property NAME=" " SCOPE=" "></Property>
    </QueryDescriptor>
</Metric>

```


[Example 6-7](#) provides a definition of the metric collection from the default collection metadata file for the metric defined in [Example 6-6](#). The information highlighted in bold font is provided by the configuration metadata XML file ([Example 6-3](#)).

For more information about the default collection metadata file, see [Section 3.5](#), "Creating the Default Collection File".

Example 6-7 Defining Metric Collection

```
<CollectionItem NAME="oracle_home_config_test" CONFIG="TRUE">
  <Schedule OFFSET_TYPE="INCREMENTAL">
    <!-- Configuration Collection is done at most once every 24 Hours -->
    <IntervalSchedule INTERVAL="24" TIME_UNIT="Hr" />
  </Schedule>
  <MetricColl NAME="EM_ECM_OH_HOME_INFO" />
  <MetricColl NAME="EM_ECM_OH_DEP_HOMES" />
  <MetricColl NAME="EM_ECM_OH_CRS_NODES" />
  <MetricColl NAME="EM_ECM_OH_CLONE_PROPS" />
  <MetricColl NAME="EM_ECM_OH_COMPONENT" />
  <MetricColl NAME="EM_ECM_OH_COMP_INST_TYPE" />
  <MetricColl NAME="EM_ECM_OH_COMP_DEP_RULE" />
</CollectionItem>
```

For a plug-in, do not forget to include the metadata and default_collection XML files at both oms and agent directories within the opar file.

6.3.8 Testing the Configuration Collection Data

After integrating the configuration collection tables into the configuration management framework, you can test the configuration collection by completing the following steps:

1. Restart the Management Agent.

```
AGENT_HOME/agent/bin/emctl stop agent
AGENT_HOME/agent/bin/emctl start agent
```

In the preceding command, *AGENT_HOME* represents the Management Agent home directory.

2. From Enterprise Manager Cloud Control, select **Targets**, then select the required target.
3. Right-click the target and select **Configuration**, then select **Last Collected** to view the most recent data collection.
4. Check that the required collected data is visible.

6.3.9 Troubleshooting

If you are having problems with your configuration collections, do the following:

1. Check that your snapshot type is registered in the MGMT_ECM_SNAPSHOT_METADATA table:

```
select * from mgmt_ecm_snapshot_metadata
where target_type = your_target_type
and snapshot_type = your_snapshot_type;
```

You should see two rows. If not, check if there are any errors during registration in the regular log files for MRS in the following directory.

`$ORACLE_HOME/cfgtoollogs/pluginca`

2. Define the corresponding metrics to verify that collections begin and data accumulates. If you are looking at the latest collection in the UI, make sure first that the `UI_IGNORE` flags in the metadata are not Y for the data you are checking. If the collections are not happening, then check the following:

- Make sure that your collection item name is the same as snapshot type and `CONFIG="TRUE"` is specified for both the collection item and all its metrics.

For more information, see [Section 6.3.7, "Modifications to Standard Collection Metrics and RAW Metrics"](#).

- Make sure your metrics are defined as RAW metrics and table descriptor corresponds to your ECM tables.
- Check if the collection arrives to the Management Repository but is not updated as "current".

In the `MGMT_ECM_GEN_SNAPSHOT` table, check the `is_current` status for your target and snapshot type. If there are no rows, then the collection did not progress. The `IS_CURRENT` flag should be set to Y to indicate the latest snapshot of data. Rows with other `IS_CURRENT` values are possible for internal purposes. For example, if there are no rows with the Y value, then `IS_CURRENT` values of T and D would indicate a snapshot started loading but did not finish.

- Check the value of `META_VER` in agent target type metadata and default_collection XML files.

During development, when any new metric or collection item is added, `META_VER` may need to be bumped up in these files for registration of these new entries to succeed. Check the latest instructions for Enterprise Management development regarding the `META_VER` value.

For example, while during your development you may need to increase the version in order to register your changes, only one increase of the version per release is required, and therefore, while merging the code, `META_VER` may need to be the same as before if it were already incremented for the current release.

From a collection perspective, you have to make sure that both new target type metadata and default_collection XML files are successfully registered and that the agent is restarted with latest files. The following commands can be used to register target type metadata and default_collection XML files:

```
emctl register oms metadata -service targetType -file <target type XML filename> [-core | -pluginId <Plugin Id>] [ sysman_pwd "sysman password"]
```

```
emctl register oms metadata -service default_collection -file <default collection XML filename> [-core | -pluginId <Plugin Id>] [ sysman_pwd "sysman password"]
```

Finally, for a given target instance you are testing, make sure that its `META_VER` matches the `META_VER` of the loaded snapshot type in order to see the latest collected data that is based on your latest `META_VER`.

- Check the Valid-Ifs defined for the target type, snapshot type, and the metric to see if the category properties of the target instance match the Valid-Ifs. If not, the target would not show corresponding data since the configuration would not be applicable to such target.

- On the OMS Repository, check the `mgmt_system_error_log` table and `emoms_pbs.trc/log` for the snapshot type.

You can also examine the `mgmt_metric_errors` table using the following command:

```
select * from mgmt_metric_errors
where target_guid = '<your target guid>'
      and coll_name = '<your snapshot type>'
```

At the agent, check `gcagent.log` and other files in agent log directory for the same string.

- If that does not help, turn on agent backup file feature (add `backupUploadedFiles=true` to `emd.properties` and restart the agent).

Search for your snapshot type in the following location to ensure the agent is sending data to OMS and to see what it is sending:

```
agentStateDir/sysman/emd/upload/upload/succbkup/
```

- One potential problem you may run into relates to configuration diffing feature.

If for some reason the configuration did not load, but the agent thinks that it has, and if the configuration does not change from that point on, the agent will keep sending short "nothing-changed" files types and the loader will keep disregarding them. To clear this issue (or just eliminate it as a potential issue while debugging), clear the agent log of config by running the following comand (no spaces around comma):

```
emctl clearstate agent -incrementalconfig targetName,targetType
```

Then, to initiate the collection, run the following command:

```
emctl control agent runCollection targetName:targetType snapshot type
```

For example:

```
emctl clearstate agent -incrementalconfig myOracleHomeTargetName,oracle_
home
emctl control agent runCollection myOracleHomeTargetName:oracle_home
oracle_home_config
```

3. At Oracle Management Service and the Management Repository, check the `MGMT_SYSTEM_ERROR_LOG` table and the `emoms_pbs.trc` file for the snapshot type. Also check the `MGMT_METRIC_ERRORS` table as follows:

At the Management Agent, check the `gcagent.log` file and other files in agent log directory for the same string.

4. If you still have problems, turn on the Management Agent backup file feature:

- a. Open the `emd.properties` file.
- b. Add the following line to the file:

```
backupUploadedFiles=true
```

- c. Restart the Management Agent.

```
AGENT_HOME/agent/bin/emctl stop agent
AGENT_HOME/agent/bin/emctl start agent
```

In the preceding command, *AGENT_HOME* represents the Management Agent home directory.

- d. Search for your snapshot type in the following directory to ensure that the Management Agent is sending data to the OMS:

```
agentStateDir/sysman/emd/upload/upload/succbkup/
```

5. A potential issue can arise relating to the configuration difference feature. If the configuration did not load but the Management Agent interprets that the configuration did load, (and if the configuration does not change), then the Management Agent sends short files indicating that nothing changed and the loader will continue to disregard the files.

To clear or eliminate this potential issue, clear the Management Agent log as follows:

```
emctl clearstate agent -incrementalconfig targetName,targetType
```

For example:

```
emctl clearstate agent -incrementalconfig myOracleHomeTargetName,oracle_home
```

Then, to initiate the collection, run the following command:

```
emctl control agent runCollection targetName:targetType snapshot_type
```

For example:

```
emctl control agent runCollection myOracleHomeTargetName:oracle_home oracle_
home_config
```

6. From the Cloud Control console, test the history and comparison features to see how the results look or if any flags should be tweaked.
 - a. From Enterprise Manager Cloud Control, select **Targets**, then select the required target.
 - b. Right-click the target and select **Configuration**, then select **History** to view the configuration history or select **Compare** to test the comparison feature.

Note: For more information about these pages, see the Cloud Control online help.

Adding Job Types

By defining new job types, you can extend the utility and flexibility of the Enterprise Manager job system. Adding new job types also allows you to enhance corrective actions. This chapter assumes that you are already familiar with the Enterprise Manager job system.

This chapter includes the following topics:

- [About Job Types](#)
- [Introducing New Job Types](#)
- [Specifying a New Job Type in XML](#)
- [Using Commands](#)
- [About Command Error Codes](#)
- [Executing Long-Running Commands at the Oracle Management Service](#)
- [Specifying Parameter Sources](#)
- [Specifying Credential Information](#)
- [Specifying Security Information](#)
- [Specifying Lock Information](#)
- [Suspending a Job or Step](#)
- [Restarting a Job](#)
- [Adding Job Types to the Job Activity and Job Library Pages](#)
- [Examples: Specifying Job Types in XML](#)
- [About Performance Issues](#)
- [Adding a Job Type to Enterprise Manager](#)

7.1 Introduction to Adding Job Types

As a plug-in developer, you are responsible for the following steps with regard to adding job types:

1. Defining Job Types

You define a job type by using an XML specification that defines the steps in a job, the work (command) that each step performs, and the relationships between the steps.

For more information, see "[About Job Types](#)" on page 7-2.

2. Executing long-running commands

The job system allows plug-in developers to write commands that perform their work at the Management Service level.

For more information, see ["Executing Long-Running Commands at the Oracle Management Service"](#) on page 7-14.

3. Specifying parameter sources

By default, the job system expects plug-in developers to provide values for all job parameters, either when the job is submitted or at execution time (by adding/updating parameters dynamically).

For more information, see ["Specifying Parameter Sources"](#) on page 7-14.

4. Specifying credential information

for more information, see ["Specifying Credential Information"](#) on page 7-20.

5. Specifying security information

For more information, see ["Specifying Security Information"](#) on page 7-23.

6. Specifying lock information

For more information, see ["Specifying Lock Information"](#) on page 7-25.

7. Suspending a job or step

For more information, see ["Suspending a Job or Step"](#) on page 7-28.

8. Restarting a job

For more information, see ["Restarting a Job"](#) on page 7-28.

7.2 About Job Types

Enterprise Manager allows you to define jobs of different types that can be executed using the Enterprise Manager job system, thereby extending the number and complexity of the tasks you can automate.

By definition, a job type is a specific category of job that carries out a well-defined unit of work. A job type is uniquely identified by a string. For example, OSCommand may be a job type that runs a remote command. You define a job type by using an XML specification that defines the steps in a job, the work (command) that each step performs, and the relationships between the steps.

[Table 7–1](#) shows some of the Enterprise Manager job types and functions.

Table 7–1 Example of Job Types

Job Type	Purpose
Backup	Backs up a database.
Backup Management	Performs management functions such as crosschecks and deletions on selected backup copies, backup sets, or files.
CloneHome	Clones an Oracle Home directory.
DBClone	Clones an Oracle Database instance.
DBConfig	Configures monitoring for database releases earlier than release 10g.
Export	Exports database contents or objects within an Enterprise Manager user's schemas and tables.

Table 7–1 (Cont.) Example of Job Types

Job Type	Purpose
GatherStats	Generates and modifies optimizer statistics.
OSCommand	Runs an operating system command or script.
HostComparison	Compares the configurations of multiple hosts.
Import	Imports the content of objects and tables.
Load	Loads data from a non-Oracle Database into an Oracle Database.
Move Occupant	Moves occupants of the SYSAUX tablespace to another tablespace.
Patch	Patches an Oracle product.
Recovery	Restores or recovers a database, tablespaces, data files, or archived logs.
RefreshFromMetalink	Allows Enterprise Manager to download patches and critical patch advisory information from My Oracle Support (https://support.oracle.com).
Reorganize	Rebuilds fragmented database indexes or tables, moves objects to a different tablespace, or optimizes the storage attributes of specified objects.
Multi-Task	Runs a composite job consisting of multiple tasks.
SQLScript	Runs a SQL or PL/SQL script using SQL*Plus.

7.3 Introducing New Job Types

An Enterprise Manager job consists of a set of steps and each step runs a command or script. The job type defines how the steps are assembled. For example, which steps run serially, which ones execute in parallel, step order, and dependencies. You can express a job type, the steps, and commands in XML (for more information, see "[Specifying a New Job Type in XML](#)"). The job system then constructs an execution plan from the XML specification that allows it to run the steps in the specified order.

7.4 Specifying a New Job Type in XML

A new job type is specified in XML. The job type specification provides information to the job system on the following:

- Steps that make up the job.
- Commands or scripts to run in each step.
- How steps relate to each other. For example, whether steps run in parallel or serially, or whether one step depends on another step.
- User credentials to authenticate the job (typically, the owner of the job must provide these). The job type author must also declare these credentials in the job type XML.
- How specific job parameters should be computed (optional).
- What locks, if any, a running job execution should attempt to acquire and what happens if the locks are not available.
- What privileges users must have in order to submit a job.

The XML job type specification is then added to a metadata plug-in archive. After the metadata plug-in is added to Enterprise Manager, the job system has enough information to schedule the steps of the job, as well as what to run in each step.

7.4.1 Understanding Job Type Categories

A job type can have one of the following categories depending on how it performs tasks on the targets to which it is applied:

- **Single-Node**

A single-node job type is a job type that runs the same set of steps in parallel on every target on which the job is run. Typically, the target lists for these job types is not fixed. They can take any number of targets. The following are examples of single-node job types:

- **OSCommand**

Runs an OS command or script on all of its targets.

- **SQL**

Runs a specified SQL script on all of its targets.

- **Multi-Node/Combination**

A multi-node job type is a job type that performs different, possibly inter-related tasks on multiple targets. Such job types typically operate on a fixed set of targets. For example, a Clone job that clones an application schema might require two targets, a source database and a target database.

Note: Iterative stepsets may be used for multi-node and combination job types to repeat the same activity over multiple targets.

7.4.2 Using Agent-Bound Job Types

An Agent-bound job type is one whose jobs cannot be run unless the Agent of one or more targets in the target list is functioning and responding. A job type that fits this category must declare itself to be Agent-bound by setting the `agentBound` attribute of the `jobType` XML tag to true.

If a job type is Agent-bound, then the job system does not schedule any executions if one or more of the Agents corresponding to the targets in the target list of the job execution are down or not responding. The job (and all its scheduled steps) is set to a special state called `Suspended/Agent down`. The job is kept in this state until the Enterprise Manager repository tier detects that the emd has come back up.

At this point, the job and its steps are set to scheduled status again and the job can now execute. By declaring their job types to be Agent-bound, a job-type writer can ensure that the job system will not schedule the job when it has detected that the Agent is down.

Note: Single-node job types are Agent-bound by default while multi-node job types are not.

If an Agent-bound job has multiple targets in its target list, it is marked as `Suspended` even if one of the Agents goes down.

A good example of an Agent-bound job type is the OSCCommand job type, which executes an OSCCommand using the Agent of a specified target. However, not all job types are Agent-bound. For example, a job type that executes SQL in the Management Repository is not Agent-bound.

Enterprise Manager has a heartbeat mechanism that enables the repository tier to quickly determine when a remote emd goes down. After an emd is marked as Down, all Agent-bound job executions that have this emd in their target list are marked Suspended/Agent Down. However, there is still a possibility that the job system will try to dispatch some remote operations during the time the emd went down and when the Management Repository detects the fact. In cases when agent cannot be contacted and the step executes, the step is set back to a SCHEDULED state and is retried by the job system. The series of retries continues until the heartbeat mechanism marks the node as down, at which point the job is suspended.

When a job is marked as Suspended/Agent Down, by default the job system keeps the job in that state until the emd comes back up. However, there is a parameter called the grace period which, if defined, can override this behavior. The grace period is the maximum amount of time (in minutes) that a job's execution is allowed to start executing within. If the job cannot start within this grace period, the job execution is skipped for that schedule.

The only way that a job execution in a Suspended/Agent Down state can resume is for the Agents to come back up. The `resume_execution()` APIs cannot be used to resume the job.

7.4.3 About Job Steps

The unit of execution in a job is called a step. A step has a command, which determines what work the step will be doing. Each command has a Java class, called a command executor, that implements the command. A command also has a set of parameters, which will be interpreted by the command executor.

The job system will offer a fixed set of pre-built commands, such as the remote operation command (which executes a command remotely), the file transfer command that transfers a file between two Agents, and a get file command that streams a log file produced on the Agent tier into the Management Repository).

Steps are grouped into sets called stepsets. Stepsets can contain steps or other stepsets and can be categorized into the following types:

- **Serial Stepsets**

Serial stepsets are stepsets whose steps execute serially. Steps in a serial stepset can have dependencies on their execution. For example, a job can specify that step S2 executes only if step S1 completes successfully, or that step S3 executes only if S1 fails.

Steps in a serial stepset can have dependencies only on other steps or stepsets within the same stepset. By default, a serial stepset is considered to complete successfully if the last step in the stepset completed successfully. It is considered to have aborted/failed if the last step in the stepset was aborted. This behavior may be overridden by using the *stepsetStatus* attribute. Overriding is allowed only when the step is not a dependent on another (no *successOf*/*failureOf*/*abortOf* attribute).

- **Parallel Stepsets**

Parallel stepsets are stepsets whose steps execute in parallel (execute simultaneously). Steps in a parallel stepset cannot have dependencies. A parallel

stepset is considered to have succeeded if all the parallel steps have completed successfully. It is considered to have aborted if any step within it was aborted. By default, a parallel stepset is considered to have failed if one or more of its constituent steps failed, and no steps were aborted. This behavior can be overridden by using the `stepsetStatus` attribute.

- **Iterative Stepsets**

Iterative stepsets are special stepsets that iterate over a vector parameter. The target list of a job is available using special, implicit parameters named `job_target_names` and `job_target_types`. An iterative stepset iterates over the target list or vector parameter and essentially executes the stepset N times; once for each value of the target list or vector parameter.

Iterative stepsets can execute in parallel (N stepset instances execute at simultaneously), or serially (N stepset instances are scheduled serially, one after another). An iterative stepset is said to have succeeded if all its N instances have succeeded. Otherwise, it is said to have aborted if at least one of the N stepsets aborted. It is said to have failed if at least one of the N stepsets failed and none were aborted. An abort will always cause an iterative stepset to stop processing further.

Steps within each iterative stepset instance execute serially and can have serial dependencies similar to those within serial stepsets. Iterative serial stepsets have an attribute called `iterateHaltOnFailure` (not applicable for `iterativeParallel` stepsets). If this is set to true, the stepset halts at the first failed or aborted child iteration. By default, all iterations of an iterative serial stepset execute, even if some of them fail (`iterateHaltOnFailure=false`).

- **Switch Stepsets**

Switch stepsets are stepsets where only one of the steps in the stepset is executed based on the value of a specified job parameter. A switch stepset has an attribute called `switchVarName`, which is a job (scalar) parameter whose value will be examined by the job system to determine which of the steps in the stepset should be executed. Each step in a switch stepset has an attribute called `switchCaseVal`, which is one of the possible values the parameter specified by `switchVarName` can have.

The step in the switch stepset that is executed is the one whose `switchCaseVal` parameter value matches the value of the `switchVarName` parameter of the switch stepset. Only the selected step in the switch stepset is executed. Steps in a switch stepset cannot have dependencies with other steps or stepsets within the same stepset or outside.

By default, a switch stepset is considered to complete successfully if the selected step in the stepset completed successfully. It is considered to have aborted/failed if the selected step in the stepset was aborted/failed. Also, a switch stepset will succeed if no step in the stepset was selected.

For example, if there is a switch stepset with two steps, `S1` and `S2` and you specify that `switchVarName` is `sendEmail` and that `switchCaseVal` for `S1` is `true` and `false` for `S2`. If the job is submitted with the job parameter `sendEmail` set to `true`, then `S1` will be executed. If the job is submitted with the job parameter `sendEmail` set to `false`, then `S2` will be executed. If the value of `sendEmail` is anything else, the stepset will still succeed but do nothing.

- **Nested Jobs**

One of the steps in a stepset may itself be a reference to another job type. A job type can therefore include other job types within itself. However, a job type cannot reference itself.

Nested jobs are a convenient way to reuse blocks of functionality. For example, performing a database backup could be a job in its own right, with a complicated sequence of steps. However, other job types (such as patch and clone) might use the backup facility as a nested job. With nested jobs, the job type writer can choose to pass all the targets of the containing job to the nested job, or only a subset of the targets. Likewise, the job type can specify whether the containing job should pass all its parameters to the nested job or whether the nested job has its own set of parameters (derived from the parent job's parameters). The status of a nested job is determined by the status of the individual steps and stepsets (and possibly other nested jobs) within the nested job.

7.4.3.1 Affecting the Status of a Stepset

The default algorithm by which the status of a stepset is computed from the status of its steps can be altered by the job type, using the `stepsetStatus` attribute of a stepset. By setting `stepsetStatus` to the name (ID) of a step, stepset, or job contained within it, a stepset can indicate that the status of the stepset depends on the status of the specific step, stepset, or job named in the `stepStatus` attribute. This feature is useful if the author of a job type wishes a stepset to succeed, even if certain steps within it fail.

A good example is a step that runs as the final step in a stepset in a job that sends e-mail about the status of the job to a list of administrators. The actual status of the job should be set to the status of the step (or steps) that performs the work, not the status of the step that sent the e-mail. Only steps that are unconditionally executed can be named in the `stepsetStatus` attribute. A step, stepset, or job that is executed as a `successOf` or `failureOf` dependency cannot be named in the `stepsetStatus` attribute.

7.4.3.2 Passing Job Parameters

The parameters of the job can be passed to steps by enclosing the parameter name in a placeholder (contained within two % symbols). For example, `%patchNo%` would represent the value of a parameter named `patchNo`. The job system will substitute the value of this parameter when it is passed to the command executor of a step.

Placeholders can also be defined for vector parameters by using the `[]` notation. For example, the first value of a vector parameter called `patchList` is referenced as `%patchList%[1]`, the second is `%patchList%[2]`.

The job system provides a predefined set of placeholders that can be used. These are always prefixed by `job_`. The following placeholders are provided:

- `job_iterate_index`
The index of current value of the parameter in an iterative stepset, when iterating over any vector parameter. The index refers to the closest enclosing stepset only. In case of nested iterative stepsets, the outer iterate index cannot be accessed.
- `job_iterate_param`
The name of the parameter being iterated over, in an iterative stepset.
- `job_target_names[n]`
The job target name at position `n`. For single-node jobs, the array would always be only of size 1 and refer only to the current node the job is execution on, even if the job was submitted against multiple nodes.
- `job_target_types[n]`

The type of the job target at position *n*. For single-node jobs, the array would always only be of size one and refer only to the current node the job is executing on, even if the job was submitted against multiple nodes.

- `job_name`
The name of the job.
- `job_type`
The type of the job.
- `job_owner`
The Enterprise Manager user that submitted the job.
- `job_id`
The job id. This is a string representing a globally unique identifier (GUID).
- `job_execution_id`
The execution id. This is a string representing a GUID.
- `job_step_id`
The step id. This is an integer.

In addition to the above placeholders, the following target-related placeholders are also supported:

- `emd_root`: The root location of the emd install
- `perlbin`: The location of the (Enterprise Manager) Perl install
- `scriptsdir`: The location of emd-specific scripts

The above placeholders are not interpreted by the job system, but by the Management Agent. For example, when `%emd_root%` is used in the `remoteCommand` or `args` parameters of the `remoteOp` command, or in any of the file names in the `putFile`, `getFile` and `transferFile` commands, the Management Agent substitutes the actual value of the Management Agent root location for this placeholder.

7.4.3.3 About Job Step Output and Errors

A step consists of a status (indicates whether it succeeded, failed, or terminated), some output (the log of the step), and an error message. If a step failed, the command executed by the step could indicate the error in the error message column. By default, the standard output and standard error of an asynchronous remote operation is set to be the output of the step that requested the remote operation.

A step can choose to insert error messages by either using the `getErrorWriter()` method in `CommandManager` (synchronous), or by using the `insert_step_error_message` API in the `mgmt_jobs` package (typically, this is called by a remotely executing script in a command channel).

7.5 Using Commands

This section describes available commands and associated parameters. Targets of any type can be provided for the target names and target type parameters described in the following sections. The job system will automatically identify and contact the Agent that is monitoring the specified targets.

7.5.1 About the remoteOp Command

The remote operation command has the identifier `remoteOp`. The command accepts a credential usage with name as `defaultHostCred`. These credentials are required to perform the operation on the host of the target. The binding can be performed as follows:

```
<step ID="Step_2" command="remoteOp">
  <credList>
    <cred usage="defaultHostCred" reference="osCreds"/>
  </credList>
  <paramList>
    <param name="targetName">%job_target_names%[1]</param>
    <param name="targetType">%job_target_types%[1]</param>
    <param name="remoteCommand">%remoteCommand%</param>
    <param name="args">%args%</param>
    <param name="executeSynchronous">false</param>
  </paramList>
</step>
```

Here `defaultHostCred` is the credential usage which is understood by the command. For example, the java code in the command would make a call for credential with this string, whereas the `osCreds` is the credential usage declared in the job type at the top level.

It takes the following parameters:

- `remoteCommand`: The path name to the executable/script (for example, `/usr/local/bin/perl`).
- `args`: A comma-separated list of arguments to the `remoteCommand`.
- `targetName`: The name of the target that the command is executed on. Note that placeholders can be used to represent targets.
- `targetType`: The target type that the command is executed on.
- `executeSynchronous`: This option defaults to false whereby a remote command always executes asynchronously on the Agent and the status of the step is updated after the command is done executing.

If set to true, the command executes synchronously, waiting until the Agent completes the process. Typically, this parameter is set to true for quick, short-lived remote operations (such as starting up a listener). For remote operations that take a long time to execute, this parameter should always be set to false.

- `successStatus`: A comma-separated list of integer values that determines the success of the step. If the remote command returns any of these numbers as the exit status, the step is considered successful. The default is zero. These values are only applicable when `executeSynchronous` is set to true.
- `failureStatus`: A comma-separated list of integer values that determines the failure of the step. If the remote command returns any of these numbers as the exit status, the step is considered to have failed. The default is all non-zero values. These values are only applicable when `executeSynchronous` is set to true.
- `input`: If specified, this is passed as standard input to the remote program.
- `outputType`: Specifies the type of output the remote command is expected to generate. This can have two values, `normal` (the default) or `command`. Normal output is output that is stored in the log corresponding to this step and is not interpreted in any way. Command output is output that could contain one or more command blocks, which are XML sequences that map to pre-registered SQL

procedure calls. By using the command output option, a remote command can generate command blocks that can be directly loaded into schema in the Enterprise Manager repository database.

The standard output generated by the executed command is stored by the job system as the output corresponding to this step.

7.5.2 Using the fileTransfer Command

The `fileTransfer` command transfers a file from one Agent to another. It can also execute a command on the source Agent and transfer its standard output as a file to the destination Agent or as standard input to a command on the destination Agent. The `fileTransfer` command is always asynchronous and it takes the following parameters:

```
<step ID="S1" command="fileTransfer">
  <credList>
    <cred usage="srcReadCreds" reference="mySourceReadCreds"/>
    <cred usage="dstWriteCreds" reference="myDestWriteCreds"/>
  </credList>
  <paramList>
    <param name="sourceTargetName">%job_target_names%[1]</param>
    <param name="sourceTargetType">%job_target_types%[1]</param>
    <param name="destTargetName">%job_target_names%[2]</param>
    <param name="destTargetType">%job_target_types%[2]</param>
    <param name="sourceFile">%sourceFile%</param>
    <param name="sourceCommand">%sourceCommand%</param>
    <param name="sourceArgs">%sourceArgs%</param>
    <param name="sourceInput">%sourceInput%</param>
    <param name="destFile">%destFile%</param>
    <param name="destCommand">%destCommand%</param>
    <param name="destArgs">%destArgs%</param>
  </paramList>
</step>
```

The following command uses two credentials. The `srcReadCreds` credential is used to read the file from the source and the `dstWriteCreds` credential is used to write the file to the destination. The binding can be performed as follows:

```
<step ID="S1" command="fileTransfer">
  <credList>
    <cred usage="srcReadCreds" reference="mySourceReadCreds"/>
    <cred usage="dstWriteCreds" reference="myDestWriteCreds"/>
  </credList>
  <paramList>
    <param name="sourceTargetName">%job_target_names%[1]</param>
    <param name="sourceTargetType">%job_target_types%[1]</param>
    <param name="destTargetName">%job_target_names%[2]</param>
    <param name="destTargetType">%job_target_types%[2]</param>
    <param name="sourceFile">%sourceFile%</param>
    <param name="sourceCommand">%sourceCommand%</param>
    <param name="sourceArgs">%sourceArgs%</param>
    <param name="sourceInput">%sourceInput%</param>
    <param name="destFile">%destFile%</param>
    <param name="destCommand">%destCommand%</param>
    <param name="destArgs">%destArgs%</param>
  </paramList>
</step>
```

- `sourceTargetName`: The target name corresponding to the source Agent.

- `destTargetName`: The target name corresponding to the destination Agent.
- `destTargetType`: The target type corresponding to the destination Agent.
- `sourceFile`: The file to be transferred from the source Agent.
- `sourceCommand`: The command to be executed on the source Agent. If this is specified, then the standard output of this command is streamed to the destination Agent. Both `sourceFile` and `sourceCommand` parameters cannot be specified.
- `sourceArgs`: A comma-separated set of command-line parameters for the `sourceCommand`.
- `destFile`: The location or file name where the file is to be stored on the destination Agent.
- `destCommand`: The command to be executed on the destination emd. If this is specified, then the stream generated from the source emd (whether from a file or from a command) is sent to the standard input of this command. Both `destFile` and `destCommand` parameters cannot be specified.
- `destArgs`: A comma-separated set of command-line parameters for the `destCommand`.

The `fileTransfer` command succeeds (and returns a status code of 0) if the file was successfully transferred between the Agents. If there was an error, it returns error codes appropriate to the reason for failure.

7.5.3 About the `putFile` Command

The `putFile` command affords the capability to transfer large amounts of data from the Management Repository to a file on the Management Agent. The data transferred can come from a blob in the Management Repository, a file on the file system, or be embedded in the specification (inline).

If a file is being transferred, the location of the file must be accessible from the Management Repository installation. If a blob in a database is being transferred, it must be in a table in the Management Repository database that is accessible to the Management Repository schema user (typically `mgmt_rep`).

The command accepts a credential usage with name as `defaultHostCred`. These credentials are required to write the file at the host of the target. The binding can be performed as follows:

```
<step ID="S1" command="putFile">
  <credList>
    <cred usage="defaultHostCred" reference="osCreds"/>
  </credList>
  <paramList>
    <param name="sourceType">file</param>
    <param name="targetName">%job_target_names%[1]</param>
    <param name="targetType">%job_target_types%[1]</param>
    <param name="sourceFile">%oms_root%/myfile</param>
    <param name="destFile">%emd_root%/yourfile</param>
  </paramList>
</step>
```

The `putFile` command requires the following parameters:

- `sourceType`: The type of the source data. This may be `sql`, `file`, or `inline`.
- `targetName`: The name of the target where the file is to be transferred (destination Agent).

- `targetType`: The type of the destination target.
- `sourceFile`: The file to be transferred from the Management Repository, if the `sourceType` is set to `fileSystem`. This must be a file that is accessible to the Management Repository installation.
- `sqlType`: The type of SQL data (if the `sourceType` is set to `sql`). Valid values are CLOB and BLOB.
- `accessSql`: A SQL statement that is used to retrieve the blob data (if the `sourceType` is set to `sql`). For example, "select output from my_output_table where blob_id=%blobid%".
- `destFile`: The location or file name where the file is to be stored on the destination Agent.
- `contents`: If the `sourceType` is set to "inline", this parameter contains the contents of the file. Note that the text could include placeholders for parameters in the form `%param%`.

The `putFile` command succeeds if the file was transferred successfully and the status code is set to 0. On failure, the status code is set to an integer appropriate to the reason for failure.

7.5.4 Using the `getFile` Command

The `getFile` command transfers a file from a Management Agent to the Management Repository. The file is stored as the output of the step that executed this command.

The command accepts a credential usage with name as `defaultHostCred`. These credentials are required to read the file at the host of the target. The binding can be performed as follows:

```
<step ID="S1" command="getFile">
  <credList>
    <cred usage="defaultHostCred" reference="osCreds"/>
  </credList>
  <paramList>
    <param name="targetName">%job_target_names%[1]</param>
    <param name="targetType">%job_target_types%[1]</param>
    <param name="sourceFile">%sourceFile%</param>
    <param name="destType">%destType%</param>
    <param name="destFile">%destFile%</param>
    <param name="destParam">%destParam%</param>
  </paramList>
</step>
```

The `getFile` command has the following parameters:

- `sourceFile`: The location of the file to be transferred on the Agent.
- `targetName`: The name of the target whose Agent will be contacted to get the file.
- `targetType`: The type of the target.

The `getFile` command succeeds if the file was transferred successfully and the status code is set to 0. On failure, the status code is set to an integer appropriate to the reason for failure.

7.5.5 Using the execAndSuspend Command

The `execAndSuspend` command accepts a credential usage with name as `defaultHostCred`. These credentials are required to perform the operation on target host of the target. The binding can be performed as follows:

```
<step ID="Ta_Sl_suspend" command="execAndSuspend">
  <credList>
    <cred usage="defaultHostCred" reference="osCreds"/>
  </credList>
  <paramList>
    <param name="remoteCommand">%command%/param>
    <param name="args">%args%/param>
    <param name="targetName">%job_target_names%[1]/param>
    <param name="targetType">%job_target_types%[1]/param>
    <param name="suspendTimeout">2</param>
  </paramList>
</step>
```

Here `defaultHostCred` is the credential usage which is understood by the command. For example, the java code in the command would make a call for credential with this sting, whereas the `osCreds` is the credential usage declared in the job type at the top level.

7.6 About Command Error Codes

The `remoteOp`, `putFile`, `fileTransfer` and `getFile` commands return the error codes listed in [Table 7-2, "Command Error Codes"](#). In the messages below, "command process" refers to a process that the Agent executes that actually executes the specified remote command and grabs the standard output and standard error of the executed command.

On a UNIX install, this process is called `nmo` and lives in `$EMD_ROOT/bin`. It must be `SETUID` to root before it can be used successfully. This does not pose a security risk since `nmo` will not execute any command unless it has a valid username and password.

Table 7-2 Command Error Codes

Error Code	Description
0	No error.
1	Could not initialize core module. Most likely, something is wrong with the install or environment of the Agent.
2	The Agent ran out of memory.
3	The Agent could not read information from its input stream.
4	The size of the input parameters was too large for the Agent to handle.
5	The command process was not setuid to root. (Every UNIX Agent installation has an executable called <code>nmo</code> , which must be setuid root).
6	The specified user does not exist on this system.
7	The password was incorrect.
8	Could not run as the specified user.
9	Failed to fork the command process (<code>nmo</code>).
10	Failed to execute the specified process.

Table 7–2 (Cont.) Command Error Codes

Error Code	Description
11	Could not obtain the exit status of the launched process.
12	The command process was interrupted before exit.
13	Failed to redirect the standard error stream to standard output.

7.7 Executing Long-Running Commands at the Oracle Management Service

The job system allows plug-in developers to write commands that perform their work at the Management Service level. For example, a command that reads two LOBs from the database and performs various transformations on them and writes them back. The job system expects such commands to implement an (empty) interface called `LongRunningCommand`, which is an indication that the command executes synchronously on the middle tier, and could potentially execute for a long time. This will allow a component of the job system called the dispatcher to schedule the long-running command as efficiently as possible, in such a way as to not degrade the throughput of the system.

7.7.1 Configuring the Job Dispatcher to Handle Long-Running Commands

The dispatcher is a component of the job system that executes the various steps of a job when they are ready to execute. The command class associated with each step is called and any asynchronous operations requested by it are dispatched; a process referred to as dispatching a step. The dispatcher uses thread-pools to execute steps. A thread-pool is a collection of a specified number of worker threads, any one of which can dispatch a step.

The job system dispatcher uses two thread-pools, namely, a short-command pool for dispatching asynchronous steps and short synchronous steps, and a long-command pool for dispatching steps that have long-running commands. Typically, the short-command pool will have a larger number of threads (for example, 25) compared to the long-running pool (for example, 10).

The theory is that long-running middle-tier steps will be few compared to more numerous, short-running commands. However, the sizes of the two pools will be fully configurable in the dispatcher to suit the job mix at a particular site. Since multiple dispatchers can be run on different nodes, the site administrator will be able to even dedicate a dispatcher to only dispatch long-running or short-running steps.

7.8 Specifying Parameter Sources

By default, the job system expects plug-in developers to provide values for all job parameters, either when the job is submitted or at execution time (by adding/updating parameters dynamically). Typically, an application supplies these parameters in one of the following ways:

- Asking the user of the application at the time of submitting the job.
- Fetching parameter values from application-specific data (such as a table) and then inserting them into the job parameter list.
- Generating new parameters dynamically through the command blocks in the output of a remote command. These could be used by subsequent steps.

The job system offers the concept of parameter sources so that plug-in developers can simplify the amount of application-specific code they have to write to fetch and populate job or step parameters (such as the second category above). A parameter source is a mechanism that the job system uses to fetch a set of parameters, either when a job is submitted or when it is about to start executing.

The job system supports SQL (a PL/SQL procedure to fetch a set of parameters), credential (retrieval of username and password information from the Enterprise Manager credentials table) and user sources. Plug-in developers can use these pre-built sources to fetch a wide variety of parameters. When the job system has been configured to fetch one or more parameters using a parameter source, the parameters need not be specified in the parameter list to the job when a job is submitted. The job system will automatically fetch the parameters and add them to the parameter list of the job.

A job type can embed information about the parameters that need to be fetched by having an optional paramInfo section in the XML specification. The following is a snippet of a job type that executes a SQL query on an application-specific table to fetch three parameters, a, b, and c.

```
<jobType version="1.0" name="OSCommand" >
  <paramInfo>
    <!-- Set of scalar params -->
    <paramSource paramNames="a,b,c" sourceType="sql" overrideUser="true">
      select name, value from name_value_pair_table where
        name in ('a', 'b', 'c');
    </paramSource>
  </paramInfo>
  .... description of job type follows ....
</jobType>
```

As can be seen from the example, the paramInfo section consists one or more paramSource tags. Each paramSource tag references a parameter source that can be used to fetch one or more parameters. The paramNames attribute is a comma-separated set of parameter names that the parameter source is expected to fetch. The sourceType attribute indicates the source that will be used to fetch the parameters (one of sql, credential or user).

The overrideUser attribute, if set to true, indicates that this parameter-fetching mechanism will always be used to fetch the value of the parameters, even if the parameter was specified by the user (or application) at the time the job was submitted. The default for the overrideUser attribute is false, meaning that the parameter source mechanism will be disabled if the parameter was already specified when the job was submitted. A parameter source could have additional source-specific properties that describe the fetching mechanism in greater detail and these will be described in the following sections.

The evaluateOnRetry attribute is an optional attribute, applicable for all. The default setting is false for all, except credentials (credentials ignores the value set and forces true). It indicates whether the parameter source must be re-run when a failed execution of this job type is retried.

7.8.1 Understanding SQLParameter Source

The SQL parameter source allows the plug-in developer to specify a SQL query or a PL/SQL procedure that will fetch a set of parameters.

7.8.1.1 Using a PL/SQL Procedure to Fetch Scalar and Vector Parameters

The job type XML syntax is as follows:

```
<paramSource sourceType="sql" paramNames="param1, param2, ...">
  <sourceParam name="procName" value="MyPackage.MyPLSQLProc"/>
  <sourceParam name="procParams" value="%a%, %b%[1], ..."/>
</paramSource>
```

The values specified in `paramNames` are the names of the parameters that are expected to be returned by the PL/SQL procedure specified in `procName`. The values in `procParams` specify the list of values to be passed to the PL/SQL procedure.

PL/SQL Procedure Definition

The definition of the PL/SQL procedure must adhere to the following guidelines:

- The PL/SQL procedure must be accessible from the SYSMAN schema
- The PL/SQL procedure must have the following signature:

```
PROCEDURE MySQLProc(p_param_names MGMT_JOB_VECTOR_PARAMS,
                    p_proc_params  MGMT_JOB_VECTOR_PARAMS,
                    p_param_list   OUT MGMT_JOB_PARAM_LIST)
```

The list of parameters specified in `paramNames` are passed as parameter `p_param_names` to the procedure.

The comma-separated list of values specified in `procParams` allows the plug-in developer to pass a list of scalar (string/VARCHAR2) values as parameters to the procedure. These values are substituted with job parameter references (if used), bundled into an array (in the order specified in the XML) and passed to the PL/SQL procedure as the second parameter (`p_proc_params`)

The third parameter is an OUT parameter that should contain the list of parameters fetched by the procedure. The names of the parameters returned by this OUT parameter must match the names specified in `p_param_names`.

Note: Although this check is not currently enforced, plug-in developers are strongly advised to ensure that the names of the parameters returned by `p_param_list` matches or is a subset of the list of parameter names passed in `p_param_names`.

Example

The following SQL parameter source creates a parameter named `db_role_suffix` based on an existing parameter named `db_role`. It also preserves the type (scalar/vector) of the original parameter and therefore looks up the parameter from the internal tables rather than have its value passed (`db_role` is passed as a literal rather than as a substituted value). The values of `job_id` and `job_execution_id` are passed substituted.

```
<paramSource sourceType="sql" paramNames="db_role_suffix">
  <sourceParam name="procName" value="MGMT_JOB_FUNCTIONS.get_dbrole_
    prefix"/>
  <sourceParam name="procParams" value="%job_id%, %job_execution_id%, db_
    role"/>
</paramSource>
```

Within the PL/SQL procedure `MGMT_JOB_FUNCTIONS.get_dbrole_prefix`, the `p_proc_params` list contains the values corresponding to the `job_id` at index 1 and the

execution_id at index 2, while the element at index 3 corresponds to the literal text db_role.

Available SQL Paramsource Procedures

The following PL/SQL procedures have been provided by the job system team for use in job types across Enterprise Manager:

- is_null

Checks whether the passed job variable is null. A missing variable is also considered to be null. For each variable passed, the procedure creates a corresponding variable with the scalar value true if the passed variable is non-existent or is null. For all other cases, the scalar value false is set. A vector of zero elements is considered non-null.

Example:

```
<paramSource sourceType="sql" paramNames="a_is_null, b_is_null, c_is_null">
  <sourceParam name="procName" value="MGMT_JOB_FUNCTIONS.is_null"/>
  <sourceParam name="procParams" value="%job_id%, %job_execution_id%, a, b,
  c"/>
</paramSource>
```

In this example, the job variables a, b, and c are checked for null-ness and the variables a_is_null, b_is_null, and c_is_null are assigned the values of true or false correspondingly.

- add_dbrole_prefix

For every variable passed, the procedure prefixes the string AS if the value is not null or Normal (case-insensitive), otherwise it returns null. Therefore, a variable with value SYSDBA would result in a value of AS SYSDBA, but a value of Normal would return null. If the passed variable corresponds to a vector, the same logic is applied to each individual element of the vector. This is useful while using DB credentials to connect to a SQL*Plus session.

Example:

```
<paramSource sourceType="sql" paramNames="db_role_suffix1, db_role_
suffix2">
  <sourceParam name="procName" value="MGMT_JOB_FUNCTIONS.get_dbrole_
  prefix"/>
  <sourceParam name="procParams" value="%job_id%, %job_execution_id%, db_
  role1, db_role2"/>
</paramSource>
```

Here, the values of the variables db_role1 and db_role2 are prefixed with AS as necessary and saved into variables db_role_suffix1 and db_role_suffix2 respectively.

7.8.2 About the User Parameter Source

The job system also offers a special parameter source called "user" which indicates that a set of parameters must be supplied when a job of that type is submitted. If a parameter is declared to be of source "user" and the "required" attribute is set to "true", the job system will validate that all specified parameters in the source are provided when a job is submitted.

The user source can be evaluated at job submission time or job execution time. When evaluated at submission time, it causes an exception to be thrown if any required

parameters are missing. When evaluated at execution time, it causes the execution to abort if there are any missing required parameters.

```
<paramInfo>
  <!-- Indicate that parameters a, b and c are required params -->
  <paramSource paramNames="a, b, c" required="true" sourceType="user" />
</paramInfo>
```

The user source can also be used to indicate that a pair of parameters are target parameters. For example:

```
<paramInfo>
  <!-- Indicate that parameters a, b, c, d, e, f are target params -->
  <paramSource paramNames="a, b, c, d, e, f" sourceType="user" >
    <sourceParam name="targetNameParams" value="a, b, c" />
    <sourceParam name="targetTypeParams" value="d, e, f" />
  </paramSource>
</paramInfo>
```

The example shown above indicates that parameters (a,d), (b,e), (c,f) are parameters that hold target information. Parameter "a" holds target names and "d" holds the corresponding target types. Similarly with parameters "b" and "e", and "c" and "f". For each parameter that holds target names, there must be a corresponding parameter that holds target types. The parameters may be either scalar or vector.

7.8.3 About the Inline Parameter Source

The `inline` parameter source allows job types to define parameters in terms of other parameters. It is a convenient mechanism to construct parameters that can be reused in other parts of the job type. For example, the section below creates a parameter called `filename` based on the job execution id, presumably for use in other parts of the job type.

```
<jobType>
  <paramInfo>
    <!-- Indicate that value for parameter filename is provided inline -->
    <paramSource paramNames="fileName" sourceType="inline" >
      <sourceParam name="paramValues" value="%job_execution_id%.log" />
    </paramSource>
  </paramInfo>
  .....
  <stepset ID="main" type="serial">
    <step command="putFile" ID="S1">
      ...
      <param name="destFile">%fileName%</param>
      ...
    </step>
  </stepset>
</jobType>
```

The following example sets a vector parameter called `vparam` to be a vector of the values `v1`, `v2`, `v3`, and `v4`. Only one vector parameter at a time can be set using the inline source.

```
<jobType>
  <paramInfo>
    <!-- Indicate that value for parameter vparam is provided inline -->
    <paramSource paramNames="vparam" sourceType="inline" >
      <sourceParam name="paramValues" value="v1,v2,v3,v4" />
      <sourceParam name="vectorParams" value="vparam" />
    </paramSource>
  </paramInfo>
</jobType>
```

```

        </paramSource>
    </paramInfo>
    ....

```

7.8.4 Using the checkValue Parameter Source

The `checkValue` parameter source allows job types to have the job system check that a specified set of parameters has a specified set of values. If a parameter does not have the specified value, the job system will either terminate or suspend the job.

```

<paramInfo>
    <!-- Check that the parameter halt has the value true. If not, suspend the
job
-->
    <paramSource paramNames="halt" sourceType="checkValue" >
        <sourceParam name="paramValues" value="true" />
        <sourceParam name="action" value="suspend" />
    </paramSource>
</paramInfo>

```

The following example checks whether a vector parameter `v` has the values `v1,v2,v3`, and `v4`. Only one vector parameter at a time can be specified in a `checkValue` parameter source. If the vector parameter does not have those values, in that order, then the job is terminated.

```

<paramInfo>
    <!-- Check that the parameter halt has the value true. If not, suspend the job
-->
    <paramSource paramNames="v" sourceType="checkValue" >
        <sourceParam name="paramValues" value="v1,v2,v3,v4" />
        <sourceParam name="action" value="abort" />
        <sourceParam name="vectorParams" value="v" />
    </paramSource>
</paramInfo>

```

7.8.5 About the properties Parameter Source

The `properties` parameter source fetches a named set of target properties for each of a specified set of targets and stores each set of property values in a vector parameter.

The example below fetches the properties "OracleHome" and "OracleSID" for the specified set of targets (dlsun966 and ap952sun), into the vector parameters `ohomes` and `osids`, respectively. The first vector value in the `ohomes` parameter will contain the OracleHome property for dlsun966, and the second will contain the OracleHome property for ap952sun. Likewise with the OracleSID property.

```

<paramInfo>
    <!-- Fetch the OracleHome and OracleSID property into the vector params ohmes,
osids -->
    <paramSource paramNames="ohomes,osids" overrideUser="true"
sourceType="properties">
        <sourceParams>
            <sourceParam name="propertyNames" value="OracleHome,OracleSID" />
            <sourceParam name="targetNames" value="dlsun966,ap952sun" />
            <sourceParam name="targetTypes" value="host,host" />
        </sourceParams>
    </paramSource>
</paramInfo>

```

As with the credentials source, vector parameter names can be provided for the target names and types.

```
<paramInfo>
  <!-- Fetch the OracleHome and OracleSID property into the vector params ohmes,
osids -->
  <paramSource paramNames="ohomes,osids" overrideUser="true"
sourceType="properties">
    <sourceParams>
      <sourceParam name="propertyNames" value="OracleHome,OracleSID" />
      <sourceParam name="targetNamesParam" value="job_target_names" />
      <sourceParam name="targetTypes" value="job_target_types" />
    </sourceParams>
  </paramSource>
</paramInfo>
```

7.8.6 Understanding Parameter Sources and Parameter Substitution

Parameter sources are applied in the order they are specified. Parameter substitution (of the form %param%) can be used inside sourceParam tags, but the parameter that is being substituted must exist when the parameter source is evaluated. Otherwise, the job system will substitute an empty string in its place.

7.8.7 About Parameter Encryption

The job system offers the facility of storing specified parameters in encrypted form. Parameters that contain sensitive information, such as passwords, must be stored encrypted. A job type can indicate that parameters fetched through a parameter source be encrypted by setting the encrypted attribute to true in a parameter source.

For example:

```
<paramInfo>
  <!-- Fetch params from the credentials table into vector parameters; store
them encrypted -->
  <paramSource paramNames="vec_usernames,vec_passwords" overrideUser="true"
sourceType="credentials" encrypted="true">
    <sourceParams>
      <sourceParam name="credentialType" value="patch" />
      <sourceParam name="credentialColumns" value="node_username,node_
password" />
      <sourceParam name="targetNames" value="dlsun966,ap952sun" />
      <sourceParam name="targetTypes" value="host,host" />
      <sourceParam name="credentialScope" value="system" />
    </sourceParams>
  </paramSource>
</paramInfo>
```

A job type can also specify that parameters supplied by the user be stored encrypted:

```
<paramInfo>
  <!-- Indicate that parameters a, b and c are required params -->
  <paramSource paramNames="a, b, c" required="true" sourceType="user"
encrypted="true" />
</paramInfo>
```

7.9 Specifying Credential Information

Until Oracle Enterprise 11g release 1, credentials were represented as two parameters, one for username and one for password. The job type owner can either have a

credential parameter source to extract these parameters or define these as user parameters, and then pass on the parameters to the various steps that require the parameters.

This required knowledge about the credential set, credential types, and their columns, along with knowledge about various authentication mechanisms, must be supported by the job type, irrespective of the pool of authentication schemes that could be supported by the Enterprise Manager. This restricted the freedom of the job type owner to model just the job type and ignore the authentication required to perform the operations. To overcome these issues and to evolve a unified mechanism in the job type to specify the credentials, Oracle has introduced a new concept called credential usage.

7.9.1 About Credential Usage

A credential usage is the point where the credential is required to perform an operation. The various tags present in the credential usage are required to mainly paint the credential selector UI component provided by the project 28263. Credential submissions should be made against these usages only.

7.9.2 Overview of Credential Binding

A credential binding is a reference of a credential by a step. Each step exposes its credential usage which needs to be fulfilled in the metadata. Therefore, each credential binding refers to a reference credential usage that is defined in the credential usage section of the metadata. When the step requests its own credential usage, a binding helps resolve which credential submission in a particular automation entity (Job or DP instance) should be passed to that step.

In earlier releases, the job types would have a credential parameter source to extract the username and password from the credentials (JobCredRecord) passed to the job and then these were available as parameters to the entire job type. This behavior has been deprecated with no support and is being superseded by the new credential usage structure.

The following Job type example shows the use of credentials declaration in the job type:

```
<jobType version="1.0" name="OSCommandNG"
  singleTarget="true" targetTypes="all"
  defaultTargetType="host" editable="true"
  restartable="true" suspendable="true" >
  <credentials>
    <credential usage="hostCreds" authTargetType="host"
      defaultCredentialSet="HostCredsNormal" />
  </credentials>
  <paramInfo>
    <paramSource sourceType="user" paramNames="command"
      required="true" evaluateAtSubmission="true" />
    <paramSource sourceType="inline"
      paramNames="TargetName,TargetType"
      overrideUser="true"
      evaluateAtSubmission="true">
      <sourceParam name="paramValues"
        value="%job_target_names%[1],
          %job_target_types%[1]" />
    </paramSource>
    <paramSource sourceType="properties"
      overrideUser="true"
```

```

        evaluateAtSubmission="false" >
        <sourceParam name="targetNamesParam"
            value="job_target_names" />
        <sourceParam name="targetTypesParam"
            value="job_target_types" />
    </paramSource>
    <paramSource sourceType="substValues"
        paramNames="host_command,host_args,os_script"
        overrideUser="true" evaluateAtSubmission="false">
        <sourceParam name="sourceParams"
            value="command,args,os_script" />
    </paramSource>
</paramInfo>
<stepset ID="main" type="serial" >
    <step ID="Command" command="sampleRemoteOp">
        <credList>
            <cred usage="OS_CRED" reference="hostCreds" />
        </credList>
        <paramList>
            <param name="remoteCommand">%host_command%</param>
            <param name="args">%host_args%</param>
            <param name="input"><![CDATA[%os_script%]]></param>
            <param name="largeInputParam">large_os_script</param>
            <param name="substituteLargeParam">true</param>
            <param name="targetName">%job_target_names%[1]</param>
            <param name="targetType">%job_target_types%[1]</param>
            <param name="executeSynchronous">false</param>
        </paramList>
    </step>
</stepset>
</jobType>

```

The first set of three lines declares a credential usage in the job type. The next set of lines binds the credential usage to that of the step. Note that user name and password cannot be extracted by the jobs system and therefore can no longer be exposed as parameters.

7.9.3 XSD Elements – Credential Usage and Credential Binding

The XSD element credential usage and credentials binding are explained in [Table 7–3](#) and [Table 7–4](#).

Table 7–3 Credential Usage (credential)

Attribute	Required (Y/N)	Description
usage	Y	Name of the credential through which it will be referred in the job type. All credential submissions are to be made for this name.
authTargetType	Y	Target type against which authentication is to be performed for any operation. For example, running “ls” any target means authentication against the host.
defaultCredentialSet	Y	Name of the credential set to be picked up as a credential if no submissions are found for the credential usage when required.

Table 7–3 (Cont.) Credential Usage (credential)

Attribute	Required (Y/N)	Description
credentialTypes	N	Name of the credential types which can only be used for specifying the credentials. This is to facilitate filtering of credentials in the credential selector UI component.
displayName	N	Name that is intended to be shown in the credential selector UI.
description	N	Description that is intended to be shown in the credential selector UI.

Table 7–4 Credential Binding (cred)

Attribute / sub element	Required (Y/N)	Description
usage	Y	Credential usage understood by the step.
reference	Y	Credential usage referred to and present in the declarations of the job type or DP metadata.

Note: The Credential Binding element can only be used inside the step or job elements in the job type xml.

7.10 Specifying Security Information

Typically a job type will tend to perform actions that can be considered to be "privileged". For example, patching a production database or affecting the software installed in an Oracle home directory or appltop. Accordingly, such job types should only be submitted by Enterprise Manager users that have the appropriate level of privileges to perform these actions.

The job system provides a section called securityInfo, which the author of a job type can use to specify the minimum level of privileges (system and target) that the submitter of a job of this type must have.

Having a securityInfo section allows the author of a job type to encapsulate the security requirements associated with submitting a job in the job type itself. No further code needs to be written to enforce security. Also, it ensures that Enterprise Manager users cannot directly submit jobs of a specific type (using the job system APIs and bypassing the application) unless they have the set of privileges defined by the job type author.

Example 1

The following example shows what a typical securityInfo section looks like. Suppose you are writing a job type that clones a database. This job type requires two targets, namely, a source database and a destination node on which the destination database will be created. This job type will probably require that the user submitting a clone job have a CLONE FROM privilege on the source (database) and a MAINTAIN privilege on the destination (node).

In addition, the user will require the CREATE TARGET system privilege to be able to introduce a new target into the system. Assuming that the job type is written so that the first target in the target list is the source and the second target in the target list is

the destination, the security requirements for such a job type could be addressed as shown below:

```
<jobType>
  <securityInfo>
    <privilege name="CREATE TARGET" type="system" />
    <privilege name="CLONE FROM" type="target" evaluateAtSubmission="false" >
      <target name="%job_target_names[1]" type="%job_target_types[1]" />
    </privilege>
    <privilege name="MAINTAIN" type="target" evaluateAtSubmission="false">
      <target name="%job_target_names[2]" type="%job_target_types[2]" />
    </privilege>
  </securityInfo>
  <!-- An optional <paramInfo> section will follow here, followed by the stepset
        definition of the job
  -->
  <paramInfo>
    ....
  </paramInfo>
  <stepset ...>
  </stepset>
</jobType>
```

The securityInfo section is a set of <privilege> tags. Each privilege could be a system or target privilege, as indicated by the type attribute of the tag. If the privilege is a target privilege, the targets that the privilege is attached to should be explicitly enumerated, or the target_names_param and target_types_param attributes should be used as shown in the example below. The usual %param% notation can be used to indicate job parameter and target placeholders.

By default, all <privilege> directives in the securityInfo section are evaluated at job submission time, after all submit-time parameter sources have been evaluated. The job system throws an exception if the user does not have any of the privileges specified in the securityInfo section.

Note that execution-time parameter sources will not have been evaluated at job submission time, so care should be taken to not use job parameters that may not have been evaluated yet. You could also direct the job system to evaluate a privilege directive at job execution time by setting the evaluateAtSubmission parameter to false.

The only reason you might want to do this is if the exact set of targets that the job is operating on is unknown until job execution time (for example, it is computed using an execution-time parameter source). Execution-time privilege directives are evaluated after all execution-time parameter sources are evaluated.

Example 2

Assume that you are writing a job type that requires MODIFY privilege on each one of its targets, but the exact number of targets is unknown at the time of writing. The target_names_param and target_types_param attributes could be used for this purpose. These specify vector parameters that the job system will get the target names and the corresponding target types from. These could be any vector parameters. This example uses the job target list (job_target_names and job_target_types).

```
<securityInfo>
  <privilege name="MODIFY" type="target" target_names_param="job_target_names"
    target_types_param="job_target_types" />
</securityInfo>
```

7.11 Specifying Lock Information

Often executing jobs will need to acquire resources. For example, a job applying a patch to a database may need a mechanism to ensure that other jobs (submitted by other users in the system) on the database are prevented from running while the patch is being applied. In other words, it may wish to acquire a lock on the database target so that other jobs that try to acquire the same lock block (or terminate). This will allow a patch job, once it starts, to perform its work without disruption.

Sometimes, locks could be at more than one level. A hot backup of a database, for example, can allow other hot backups to proceed (since they do not bring down the database), but cannot allow cold backups or database shutdown jobs to proceed (since they will end up shutting down the database, thereby causing the backup to fail).

A job execution can indicate that it is reserving a resource on a target by acquiring a lock on the target. A lock is really a proxy for reserving some part of the functionality of a target. When an execution acquires a lock, it will block other executions that try to acquire the same lock on the target. A lock is identified by a name and a type and can be of the following types:

- **Global:** These are locks that are not associated with a target. An execution that holds a global lock will block other executions that are trying to acquire the same global lock (such as a lock with the same name).
- **Target Exclusive:** These are locks that are associated with a target. An execution that holds an exclusive lock on a target will block executions that are trying to acquire any named lock on the target, as well as executions trying to acquire an exclusive lock on the target. Target exclusive locks have no name: there is exactly one exclusive lock per target.
- **Target Named:** A named lock on a target is analogous to obtaining a lock on one particular functionality of the target. A named lock has a user-specified name. An execution that holds a named lock will block other executions that are trying to acquire the same named lock, as well as executions that are trying to acquire an exclusive lock on the target.

Example

Locks that a job type wishes to acquire can be obtained by specifying a lockInfo section in the job type. This example lists the locks that the job is to acquire, their types, as well as the targets that it wishes to acquire the locks on:

```
<lockInfo action="suspend">
  <lock type="targetExclusive">
    <targetList>
      <target name="%backup_db%" type="oracle_database" />
    </targetList>
  </lock>
  <lock type="targetNamed" name="LOCK1" >
    <targetList>
      <target name="%backup_db%" type="oracle_database" />
      <target name="%job_target_names%[1]" type="%job_target_types%[1]" />
      <target name="%job_target_names%[2]" type="%job_target_types%[2]" />
    </targetList>
  </lock>
  <lock type="global" name="GLOBALLOCK1" />
</lockInfo>
```

This example shows a job type that acquires a target-exclusive lock on a database target whose name is given by the job parameter backup_db. It also acquires a named

target lock named "LOCK1" on three targets, namely, the database whose name is stored in the job parameter `backup_db`, and the first two targets in the target list of the job. Finally, it acquires a global lock named "GLOBALLOCK1". The "action" attribute specifies what the job system should do to the execution if any of the locks in the section cannot be obtained (presumably because some other execution is holding them). Possible values are `suspend` (all locks are released and the execution state changes to "Suspended:Lock") and `abort` (the execution terminates). The following points can be made about executions and locks:

- An execution can only attempt to obtain locks when it starts (although it is possible to override this by using nested jobs).
- An execution can acquire multiple locks. Locks are always acquired in the order specified. Because of this, executions can potentially deadlock each other if they attempt to acquire locks in the wrong order.
- Target locks are always acquired on targets in the same order as they are specified in the `<targetList>` tag.
- If a target in the target list is null or does not exist, the execution will terminate.
- If an execution attempts to acquire a lock it already holds, it will succeed.
- If an execution cannot acquire a lock (usually because some other execution is holding it), it has a choice of suspending itself or terminating. If it chooses to suspend itself, all locks it has acquired so far will be released, and the execution is put in the state `Suspended/Lock`.
- All locks held by an execution will be released when an execution finishes (whether it completes, aborts, or is stopped). There may be several waiting executions for each released lock and these are sorted by time, with the earliest request getting the lock.

When jobs that have the `lockInfo` section are nested inside each other, the nested job's locks are obtained when the nested job first executes, not when an execution starts. If the locks are not available, the parent execution could be suspended or terminated, possibly after a few steps have already executed.

lockInfo Example 1

In this example, two job types called `HOTBACKUP` and `COLDBACKUP` perform hot backups and cold backups, respectively, on the database. The difference is that the cold backup brings the database down, but the hot backup leaves it up. Only one hot backup can execute at a time and it should keep out other hot backups as well as cold backups.

When a cold backup is executing, no other job type can execute (since it shuts down the database as part of its execution). A third job type called `SQLANALYZE` performs scheduled maintenance activity that results in modifications to database tuning parameters (two `SQLANALYZE` jobs cannot run at the same time).

Table 7–5 shows the incompatibilities between the job types. An 'X' indicates that the job types are incompatible. An 'OK' indicates that the job types are compatible.

Table 7–5 Job Type Incompatibilities

Job Type	HOTBACKUP	COLDBACKUP	SQLANALYZE
HOTBACKUP	X	X	OK
COLDBACKUP	X	X	X
SQLANALYZE	OK	X	X

The lockInfo sections for the three job types are shown below. The cold backup obtains an exclusive target lock on the database. The hot backup job does not obtain an exclusive lock, but only the named lock "BACKUP_LOCK". Likewise, the SQLANALYZE job obtains a named target lock called "SQLANALYZE_LOCK".

Assuming that the database that the jobs operate on is the first target in the target list of the job, the lock sections of the two jobs look as follows:

```
<jobType name="SQLANALYZE">
  <lockInfo action="abort">
    <lock type="targetNamed" name="SQLANALYZE_LOCK" >
      <targetList>
        <target name="%job_target_names%[1]" type="%job_target_names%[1]" />
      </targetList>
    </lock>
  </lockInfo>
  ..... Rest of the job type follows
</jobType>
```

Since a named target lock blocks all target exclusive locks, executing hot backups will suspend cold backups, but not analyze jobs (since they try to acquire different named locks). Executing SQL analyze jobs will abort other SQL analyze jobs and suspend cold backups, but not hot backups. Executing cold backups will suspend hot backups and abort SQL analyze jobs.

lockInfo Example 2

A job type called PATCHCHECK periodically checks a patch stage area and downloads information about newly staged patches into the Management Repository. Two such jobs cannot run at the same time; however, the job is not really associated with any target. The solution is for the job type to attempt to grab a global lock:

```
<jobType name="PATCHCHECK">
  <lockInfo>
    <lock type="global" name="PATCHCHECK_LOCK" />
  </lockInfo>
  ..... Rest of the job type follows
</jobType>
```

lockInfo Example 3

A job type that nests the SQLANALYZE type within itself is shown below. Note that the nested job executes after the first step (S1) executes.

```
<jobType name="COMPOSITEJOB">
  <stepset ID="main" type="serial">
    <step ID="S1" ...>
      ....
    </step>
    <job name="nestedsql" type="SQLANALYZE">
      ....
    </job>
  </stepset>
</jobType>
```

In the previous example, the nested job tries to acquire locks when it executes (since the SQLANALYZE has a lockInfo section). If the locks are currently held by other executions, then the nested job terminates (as specified in the lockInfo), which will in turn end up terminating the parent job.

7.12 Suspending a Job or Step

Suspended is a special state that indicates that steps in the job will not be considered for scheduling and execution. A step in an executing job can suspend the job, through the `suspend_job` PL/SQL API. This suspends both the currently executing step, as well as the job itself.

Suspending a job means that all steps in the job that are currently in a "scheduled" state will be marked as "suspended" and will thereafter not be scheduled or executed. All currently executing steps (this could happen, for example, in parallel stepsets) will continue to execute. However, when any currently executing step completes, the next steps in the job will not be scheduled. Instead they will be put in suspended state. When a job is suspended on submission, the above applies to the first steps in the job that would have been scheduled.

Suspended jobs may be restarted at any time by calling the `restart_job()` PL/SQL API. However, jobs that are suspended because of serialization (locking) rules are not restartable manually. The job system will restart such jobs automatically when currently executing jobs of that job type complete. Restarting a job will effectively change the state of all suspended steps to scheduled and job execution will proceed normally thereafter.

7.13 Restarting a Job

If a job has been suspended, failed or terminated, it is possible to restart it from any given step (typically, the stepset that contains a failed or terminated step). For failed or terminated jobs, what steps actually get scheduled again when a job is restarted depends on which step the job is restarted from.

7.13.1 Restarting Versus Resubmitting

If a step in a job is resubmitted, it means that it executes regardless of whether the original execution of the step completed or failed. If a stepset is resubmitted, then the first step/stepset/job in the stepset is resubmitted, recursively. Therefore, when a job is resubmitted the entire job is executed again by recursively resubmitting its initial stepset. The parameters and targets used are the same that were used when the job was first submitted. Essentially, the job executes as if it were submitted for the first time with the specified set of parameters and targets. A job can be resubmitted by using the `resubmit_job` API in the `mgmt_jobs` package. Jobs can also be resubmitted even if the earlier executions completed successfully.

Job executions that were aborted or failed can be restarted. Restarting a job generally refers to resuming job execution from the last failed step (although the job type can control this behavior using the `restartMode` attribute of steps/stepsets/jobs). In the common case, steps from the failed job execution that actually succeeded are not re-executed. A failed or terminated job can be restarted by calling the `restart_job` API in the `mgmt_jobs` package. A job that completed successfully cannot be restarted.

7.13.2 Default Restart Behavior

Restarting a job creates a new execution called the restart execution. The original, failed execution of the job is called the source execution. All parameters and targets are copied over from the source execution to the restart execution. Parameter sources are not reevaluated, unless the original job aborted because of a parameter source failure.

To restart a serial (or iterative stepset), the job system first examines the status of the serial stepset. If the status of the serial stepset is "Completed", then all the entries for its

constituent steps are copied over from the source execution to the restart execution. If the status of the stepset is "Failed" or "Aborted", then the job system starts top down from the first step in the stepset.

If the step previously completed successfully in the source execution, it is copied to the restart execution. If the step previously failed or aborted, it is rescheduled for execution in the restart execution. After such a step has finished executing, the job system determines the next steps to execute. These could be successOf or failureOf dependencies, or simply steps/stepsets/jobs that execute after the current step.

If the subsequent step completed successfully in the source execution, then it will not be scheduled for execution again and the job system merely copies the source execution status to the restart execution for that step. It continues in this fashion until it reaches the end of the stepset. It then recomputes the status of the stepset based on the new executions.

To restart a parallel stepset, the job system first examines the status of the parallel stepset. If the status of the stepset is "Completed", then all the entries for its constituent steps are copied over from the source execution to the restart execution. If the status of the stepset is "Failed" or "Aborted", the job system copies over all successful steps in the steps from the source to the restart execution. It reschedules all steps that failed or aborted in the source execution, in parallel. After these steps have finished executing, the status of the stepset is recomputed.

To restart a nested job, the restart algorithm is applied recursively to the first (outer) stepset of the nested job.

Note that in the previous paragraphs, if one of the entities being considered is a stepset or a nested job, the restart mechanism is applied recursively to the stepset or job. When entries for steps are copied over to the restart execution, the child execution entries point to the same output CLOB entries as the parent execution.

7.13.3 Using the restartMode Directive

A job type can affect the restart behavior of each step, stepset, or job within it by the use of the restartMode attribute. This can be set to "failure" (the default) or "always". When set to failure and the top-down copying process described in the previous section occurs, the step, stepset, or job is copied without being re-executed if it succeeded in the source execution. If it failed or terminated in the source execution, it is restarted recursively at the last point of failure.

When the restartMode attribute is set to "always" for a step, the step is always re-executed in a restart, regardless of whether it succeeded or failed in the source execution. The use of this attribute is useful when certain steps in a job must always be re-executed in a restart (for example, a step that shuts down a database prior to backing it up).

For a stepset or nested job, if the restartMode attribute is set to "always", then all steps in the stepset/nested job are restarted, even if they completed successfully in the source execution. If it is set to "failure", then restart is attempted only if the status of the stepset or nested job was set to Failed or Aborted in the source execution.

Note that individual steps inside a stepset or nested job may have their restartMode set to "always" and such steps are always re-executed.

Restart Examples

The following sections discuss a range of scenarios related to restarting stepsets.

Example 1

Consider the serial stepset with the sequence of steps below:

```
<jobtype ...>
<stepset ID="main" type="serial" >
  <step ID="S1" ...>
    ...
  </step>
  <step ID="S2" ...>
    ...
  </step>
  <step ID="S3" failureOf="S2"...>
    ...
  </step>
  <step ID="S4" successOf="S2"...>
    ...
  </step>
</stepset>
</jobtype>
```

In the above stepset, assume the source execution had S1 execute successfully and step S2 and S3 (the failure dependency of S2) fail. When the job is restarted, steps S1 is copied to the restart execution from the source execution without being re-executed (since it successfully completed in the source execution). Step S2, which failed in the source execution, is rescheduled and executed. If S2 completes successfully, then S4, its success dependency (which never executed in the source execution), is scheduled and executed. The status of the stepset (and the job) is the status of S4. If S2 fails, then S3 (its failure dependency) is rescheduled and executed (since it had failed in the source execution), and the status of the stepset (and the job) is the status of S3.

Assume that step S1 succeeded, S2 failed, and S3 (its failure dependency) succeeded in the source execution. As a result, the stepset (and therefore the job execution) succeeded. This execution cannot be restarted since the execution completed successfully although one of its steps failed.

Finally, assume that steps S1 and S2 succeed, but S4 (S2's success dependency) failed. Note that S3 is not scheduled in this situation. When the execution is restarted, the job system copies over the executions of S1 and S2 from the source to the restart execution, and reschedules and executes S4. The job succeeds if S4 succeeds.

Example 2

Consider the following:

```
<jobtype ...>
<stepset ID="main" type="serial" stepsetStatus="S2" >
  <step ID="S1" restartMode="always" ...>
    ...
  </step>
  <step ID="S2" ...>
    ...
  </step>
  <step ID="S3" ...>
    ...
  </step>
</stepset>
</jobtype>
```

In the previous example, assume that step S1 completes and S2 fails. S3 executes (since it does not have a dependency on S2) and succeeds. The job, however, fails, since the

stepset main has its stepsetStatus set to S2. When the job is restarted, S1 is executed all over again, although it completed the first time, since the restartMode of S1 was set to "always". Step S2 is rescheduled and executed, since it failed in the source execution. After S2 executes, step S3 is *not* rescheduled for execution again, since it executed successfully in the source execution. If the intention is that S3 must execute in the restart execution, its restartMode must be set to "always".

If, in the above example, S1 and S2 succeeded and S3 failed, the stepset main would still succeed (since S2 determines the status of the stepset). In this case, the job would succeed, and cannot be restarted.

Example 3

Consider the following example:

```
<jobtype ...>
<stepset ID="main" type="serial" >
  <stepset type="serial" ID="SS1" stepsetStatus="S1">
    <step ID="S1" ...>
      ...
    </step>
  <stepset ID="S2" ...>
    ...
  </step>
</stepset>
<stepset type="parallel" ID="PS1" successOf="S1" >
  <step ID="P1" ...>
    ...
  </step>
  <step ID="P2" ...>
    ...
  </step>
  <step ID="P3" ...>
    ...
  </step>
</stepset>
</stepset>
</jobtype>
```

In the above example, let us assume that steps S1 and S2 succeeded (and therefore, stepset SS1 completed successfully). Thereafter, the parallel stepset PS1 was scheduled, and let us assume that P1 completed, but P2 and P3 failed. As a result, the stepset "main" (and the job) failed. When the execution is restarted, the steps S1 and S2 (and therefore the stepset SS1) will be copied over without execution. In the parallel stepset PS1, both the steps that failed (P2 and P3) will be rescheduled and executed.

Now assume that S1 completed and S2 failed in the source execution. Note that stepset SS1 still completed successfully since the status of the stepset is determined by S1, not S2 (because of the stepsetStatus directive). Now, assume that PS1 was scheduled and P1 failed, and P2 and P3 executed successfully. When this job is rescheduled, the step S2 will not be re-executed (since the stepset SS1 completed successfully). The step P1 will be rescheduled and executed.

Example 4

Consider a slightly modified version of the XML in ["Example 3"](#):

```
<jobtype ...>
<stepset ID="main" type="serial" >
  <stepset type="serial" ID="SS1" stepsetStatus="S1" restartMode="always" >
    <step ID="S1" ...>
```

```
...
</step>
<stepset ID="S2" ...>
...
</step>
</stepset>
<stepset type="parallel" ID="PS1" successOf="S1" >
  <step ID="P1" ...>
    ...
  </step>
  <step ID="P2" ...>
    ...
  </step>
  <step ID="P3" ...>
    ...
  </step>
</stepset>
</stepset>
</jobtype>
```

In the previous example, assume that S1 and S2 succeeded (and therefore, stepset SS1 completed successfully). Thereafter, the parallel stepset PS1 was scheduled, and let us assume that P1 completed, but P2 and P3 failed. When the job is restarted, the entire stepset SS1 is restarted (since the restartMode is set to "always"). This means that steps S1 and S2 are successively scheduled and executed. Now the stepset PS1 is restarted, and since the restartMode is not specified (it is always "failure" by default), it is restarted at the point of failure, which in this case means that the failed steps P2 and P3 are re-executed, but not P1.

7.14 Adding Job Types to the Job Activity and Job Library Pages

In order to make a new job type accessible from the Enterprise Manager console Job Activity, or Job Library page, or both, you need to modify specific XML tag attributes.

To display the job type on Job Activity page, set useDefaultCreateUI to "true" as shown in the following example.

```
<displayInfo useDefaultCreateUI="true"/>
```

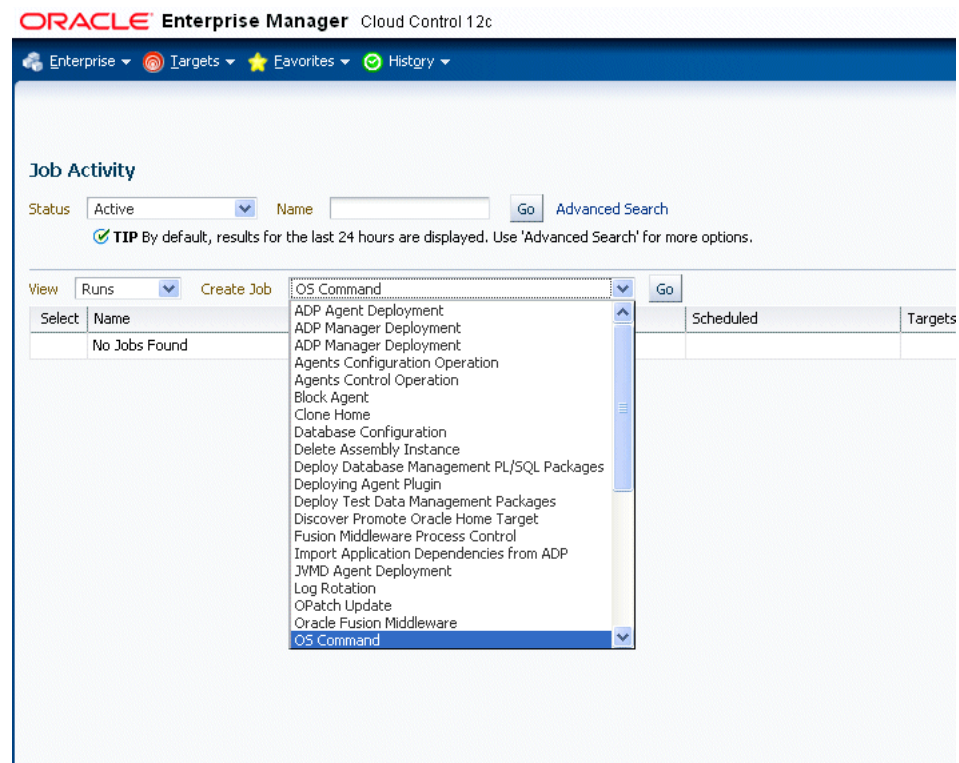
To display the job type on the Job Library page, in addition to setting useDefaultCreateUI attribute, you must also set the jobtype editable attribute to "true."

```
<jobtype name="jobType1" editable="true">
```

If only useDefaultCreateUI="true" and editable="false", then the job type will only be displayed on the Job Activity page and not on Job Library page. Also the job definition will be not editable.

7.14.1 Adding a Job Type to the Job Activity Page

As shown in [Figure 7-1](#), setting the useDefaultCreateUI attribute to "true" allows users creating a job to select the newly added job type from the Create Job menu.

Figure 7–1 Available Job Types from the Job Activity Page

Making the job type available from the Job Activity page also permits access to the default Create Job user interface when a user attempts to create a job using the newly added job type.

Adding the displayInfo Tag

The `displayInfo` tag can be added to the job definition file at any point after the `</stepset>` tag and before the `</jobtype>` tag at the end of the job definition file, as shown in the following example.

```
<jobtype ...>
<stepset ID="main" type="serial" >
  <stepset type="serial" ID="S1" stepsetStatus="S1">
    <step ID="S1" ...>
      ...
    </step>
  <stepset ID="S2" ...>
    ...
  </step>
</stepset>
<stepset type="parallel" ID="PS1" successOf="S1" >
  <step ID="P1" ...>
    ...
  </step>
  <step ID="P2" ...>
    ...
  </step>
  <step ID="P3" ...>
    ...
  </step>
</stepset>
</jobtype>
```

```
<displayInfo useDefaultCreateUI="true"/>
</jobtype>
```

7.14.2 Adding a Job Type to the Job Library Page

To make the job type available from the Job Library page, you must also set the `jobType` tag's `editable` attribute to "true" in addition to adding the `displayInfo` tag. This makes the newly added job type a selectable option from the Create Library Job menu.

Making the Job Type Editable

The `editable` attribute of the `jobtype` tag is set at the beginning of the job definition file, as shown in the following example.

```
<jobtype name="jobType1" editable="true">
<stepset ID="main" type="serial" >
  <stepset type="serial" ID="S1" stepsetStatus="S1">
    <step ID="S1" ...>
      ...
    </step>
  <stepset ID="S2" ...>
    ...
  </step>
</stepset>
<stepset type="parallel" ID="PS1" successOf="S1" >
  <step ID="P1" ...>
    ...
  </step>
  <step ID="P2" ...>
    ...
  </step>
  <step ID="P3" ...>
    ...
  </step>
</stepset>
</stepset>
<displayInfo useDefaultCreateUI="true"/>
</jobtype>
```

7.15 Examples: Specifying Job Types in XML

The following sections provide examples of specifying job types in XML.

Example 1

The following XML describes a job type called `jobType1` that defines four steps, `S1`, `S2`, `S3`, and `S4`. It executes `S1` and `S2` serially, one after another. It executes step `S3` only if step `S2` succeeds, and step `S4` only if `S2` fails. Note that all the steps execute within an iterative subset, so these actions are performed in parallel on all targets in the job target list of type database.

Note also, the use of % signs to indicate parameters, `%patchno%`, `%username%`, `%password%`, and `%job_target_name%`. The job system will substitute the value of a job parameter named "patchno" in place of the `%patchno%`. Likewise, it will substitute the values of the corresponding parameters for `%username%` and `%password%`. `%job_target_name%` and `%job_target_type%` are "pre-built" placeholders that will substitute the name of the target that the step is currently executing against.

The steps S2, S3, and S4 illustrate how the remoteOp command can be used to execute a SQL*Plus script on the Agent.

The status of a job is failed if any of the following occurs:

- S2 fails and S4 fails
- S2 succeeds and S3 fails

Note that since S2 executes after S1 (regardless of whether S1 succeeds or fails), the status of S1 does not affect the status of the job in any way.

```
<jobtype name="jobType1" editable="true" version="1.0">
<credentials>
  <credential usage="defaultHostCred" authTargetType="host"
    defaultCredentialSet="DBHostCreds"/>
  <credential usage="defaultDBCred" authTargetType="oracle_database"
    credentialTypes="DBCreds"
    defaultCredentialSet="DBCredsNormal"/>
</credentials>
<stepset ID="main" type="iterativeParallel" iterate_param="job_target_types"
iterate_param_filter="oracle_database" >
  <step ID="s1" command="remoteOp">
    <credList>
      <cred usage="defaultHostCred" reference="defaultHostCred"/>
    </credList>
    <paramList>
      <param name="remoteCommand">myprog</param>
      <param name="targetName">%job_target_names[%job_iterate_
index%]
      </param>
      <param name="targetType">%job_target_types[%job_iterate_
index%]
      </param>
      <param name="args">-id=%patchno%</param>
      <param name="successStatus">3</param>
      <param name="failureStatus">73</param>
    </paramList>
  </step>
  <step ID="s2" command="remoteOp">
    <credList>
      <cred usage="defaultHostCred" reference="defaultHostCred"/>
    </credList>
    <paramList>
      <param name="remoteCommand">myprog2</param>
      <param name="targetName">%job_target_names[%job_iterate_
index%]</param>
      <param name="targetType">%job_target_types[%job_iterate_
index%]</param>
      <param name="args">-id=%patchno%</param>
      <param name="successStatus">3</param>
      <param name="failureStatus">73</param>
    </paramList>
  </step>
  <step ID="s3" successOf="s2" command="remoteOp">
    <credList>
      <cred usage="defaultHostCred" reference="defaultHostCred"/>
      <cred usage="defaultDBCred" reference="defaultDBCred">
        <map toParam="db_username" credColumn="DBUserName"/>
        <map toParam="db_passwd" credColumn="DBPassword"/>
        <map toParam="db_alias" credColumn="DBRole"/>
      </cred>
    </credList>
  </step>
</stepset>
</jobtype>
```

```

</credList>
  <paramList>
    <param name="command">prog1</command>
    <param name="script">
      <![CDATA[
        select * from MGMT_METRICS where target_name=%job_target_type%[%job_
          iterate_param_index%]
      ]]>
    </param>
    <param name="args">%db_username%/%db_passwd%/%db_alias%</param>
    <param name="targetName">%job_target_names%[%job_iterate_
      index%]</param>
    <param name="targetType">%job_target_types%[%job_iterate_
      index%]</param>
    <param name="successStatus">0</param>
    <param name="failureStatus">1</param>
  </paramList>
</step>
  <step ID="s4" failureOf="s2" command="remoteOp">
    <credList>
      <cred usage="defaultHostCred" reference="defaultHostCred"/>
    </credList>
    <paramList>
      <param name="input">
        <![CDATA[
          This is standard input to the executed progeam. You can use placeholders
          for parameters, such as
          %job_target_name%[%job_iterate_param_index%]
        ]]>
      </param>
      <param name="remoteCommand">prog2</param>
      <param name="targetName">%job_target_names%[%job_iterate_
        index%]</param>
      <param name="targetType">%job_target_types%[%job_iterate_
        index%]</param>
      <param name="args"></param>
      <param name="successStatus">0</param>
      <param name="failureStatus">1</param>
    </paramList>
  </step>
</stepset>
<displayInfo useDefaultCreateUI="true"/>
</jobtype>

```

Example 2

The following XML describes a job type that has two steps, S1 and S2, that execute in parallel (within a parallel stepset ss1) and a third step, S3, that will execute only after both S1 and S2 have completed successfully. This is achieved by placing the step S3 in a serial stepset ("main") that also contains the parallel stepset ss1. This job type is a "multi-node" job. Note that use of %job_target_name%[1], %job_target_name%[2] in the parameters to the commands. In stepsets other than an iterative stepset, job targets can only be referred to using their position in the targets array (which is ordered).

So, %job_target_name%[1] refers to the first target, %job_target_name%[2] to the second, and so on. The assumption is that most multi-node jobs will expect their targets to be in some order. For example, a clone job might expect the source database to be the first target, and the target database to be the second target. This job fails if any of the following occurs:

- The parallel stepset SS1 fails (either one of S1 or S2, or both fail)
- Both S1 and S2 succeed, but S3 fails

Also note that the job type has declared itself to be Agent-bound. This means that the job will be set to Suspended/Agent Down state if either emd (corresponding to the first target or the second target) goes down.

```
<jobtype name="jobType2" version="1.0" agentBound="true" >
  <stepset ID="main" type="serial" editable="true">
    <!-- All steps in this stepset ss1 execute in parallel -->
    <credentials>
      <credential usage="hostCreds" authTargetType="host"
        defaultCredentialSet="HostCredsNormal"/>
    </credentials>
    <stepset ID="ss1" type="parallel" >
      <step ID="s1" command="remoteOp" >
        <credList>
          <cred usage="defaultHostCred" reference="defaultHostCred"/>
        </credList>
        <paramList>
          <param name="remoteCommand">myprog</param>
          <param name="targetName">%job_target_names%[1]</param>
          <param name="targetType">%job_target_types%[1]</param>
          <param name="args">-id=%patchno%</param>
          <param name="successStatus">3</param>
          <param name="failureStatus">73</param>
        </paramList>
      </step>
      <step ID="s2" command="remoteOp" >
        <credList>
          <cred usage="defaultHostCred" reference="hostCreds"/>
        </credList>
        <paramList>
          <param name="remoteCommand">myprog</param>
          <param name="targetName">%job_target_names%[2]</param>
          <param name="targetType">%job_target_types%[2]</param>
          <param name="args">-id=%patchno%</param>
          <param name="successStatus">3</param>
          <param name="failureStatus">73</param>
        </paramList>
      </step>
    </stepset>
    <!-- This step executes after stepset ss1 has executed, since it is inside the
    serial subset "main"
    -->
    <step ID="s3" successOf="ss1" command="remoteOp" >
      ...
    </step>
  </stepset>
  <displayInfo useDefaultCreateUI="true"/>
</jobtype>
```

Example 3

The following example defines a new job type called jobType3 that executes jobs of type jobType1 and jobType2, one after another. The job job2 of type jobType2 is executed only if the first job fails. In order to execute another job, the target list and the param list must be passed. The targetList tag has a parameter called allTargets, which when sent to true passes along the entire target list passed to this job. By setting

allTargets to false, a job type has the option of passing along a subset of its targets to the other job type.

In the example below, jobType3 passes along all its targets to the instance of the job of type jobType1, but only the first two targets in its target list (in that order) to the job instance of type jobType2. There is another attribute called allParams (associated with paramList) that performs a similar function with respect to parameters. If allParams is set to true, then all parameters of the parent job are passed to the nested job. More typically though, the nested job will have a different set of parameters (with different names).

If allParams is set to false (the default), then the job type can name the nested job parameters explicitly and they need not have the same names as those in the parent job. Parameter substitution can be used to express the nested job parameters in terms of the parent job parameters, as shown in this example.

Note that dependencies can be expressed between nested jobs just as if they were steps or stepsets. In this example, a job of type jobType3 succeeds if either the nested job job1 succeeds or if job1 fails and job2 succeeds.

```
<jobType name="jobType3" editable="true" version="1.0">
  <stepset ID="main" type="serial">
    <job type="jobType1" ID="job1" >
      <target_list allTargets="true" />
      <paramList>
        <param name="patchno">%patchno%</param>
      </paramList>
    </job>
    <job type="jobType2" ID="job2" failureOf="job1" >
      <targetList>
        <target name="%job_target_names%[1]" type="%job_target_types%[1]" />
        <target name="%job_target_names%[2]" type="%job_target_types%[2]" />
      </targetList>
      <paramList>
        <param name="patchno">%patchno%</param>
      </paramList>
    </job>
  </stepset>
<displayInfo useDefaultCreateUI="true"/>
</jobType>
```

Example 4

This example illustrates the use of the generateFile command. Let us assume that you are executing a sequence of scripts, all of which need to source a common file that sets up some environment variables, which are known only at runtime. One way to do this is to generate the variables in a file with a unique name. All subsequent scripts are passed this file name as one of their command-line arguments, which they read to set the needed environment or shell variables.

The first step, S1, in this job uses the generateFile command to generate a file named <app-home>/<execution-id>.env. Since the execution id of a job is always unique, this ensures a unique file name. It generates three environment variables, ENVVAR1, ENVVAR2, and ENVVAR3, which are set to the values of the job parameters param1, param2 and param2, respectively. These parameters must be set to the right values when the job is submitted. Note that %job_execution_id% is a placeholder provided by the job system, while %app-home% is a job parameter which must be explicitly provided when the job is submitted.

The second step, S2, executes a script called myscript. The first command-line argument to the script is the generated filename. This script must "source" the

generated file, which will set the required environment variables, and then go about it's other tasks, in the manner shown below:

```
#!/bin/ksh
ENVFILE=$1
# Execute the generated file, sets the required environment vars
. $ENVFILE
# I can now reference the variables set in the file doSomething $ENVVAR1 $ENVVAR2
$ENVVAR3...
```

The full job type specification is given below. Note the step S3 removes the file that was created by the first step S1. It is important to clean up after yourself when using the putFile and generateFile commands to write temporary files on the Management Agent. The cleanup is done here explicitly as a separate step, but it could also be done by one of the scripts that execute on the remote host.

Additionally, note the use of the securityInfo section that specifies that the user that submits a job of this job type must have maintain privilege on both the targets that the job operates on.

```
<jobtype name="jobType4" editable="true" version="1.0">
  <securityInfo>
    <privilege name="MAINTAIN" type="target" evaluateAtSubmission="false">
      <target name="%job_target_names%[1]" type="%job_target_types%[1]" />
      <target name="%job_target_names%[2]" type="%job_target_types%[2]" />
    </privilege>
  </securityInfo>
  <credentials>
    <credential usage="hostCreds" authTargetType="host"
      defaultCredentialSet="HostCredsNormal" />
  </credentials>
  <stepset ID="main" type="serial">
    <step ID="s1" command="putFile" >
      <paramList>
        <param name="sourceType">inline</param>
        <param name="destFile">%app-home%/%job_execution_id%.env</param>
        <param name="targetName">%job_target_names%[1]</param>
        <param name="targetType">%job_target_types%[1]</param>
        <param name="contents">
          <![CDATA[#!/bin/ksh
            export ENVVAR1=%param1% export ENVVAR2=%param2% export ENVVAR3=%param3%
          ]]>
        </param>
      </paramList>
    </step>
    <step ID="s2" command="remoteOp" >
      <credList>
        <cred usage="defaultHostCred" reference="hostCreds" />
      </credList>
      <paramList>
        <param name="remoteCommand">myscript</param>
        <param name="targetName">%job_target_names%[2]</param>
        <param name="targetType">%job_target_types%[2]</param>
        <param name="args">%app-home%/%job_execution_id%.env</param>
        <param name="successStatus">3</param>
        <param name="failureStatus">73</param>
      </paramList>
    </step>
    <step ID="s3" command="remoteOp" >
      <credList>
        <cred usage="defaultHostCred" reference="hostCreds" />
```

```

</credList>

<paramList>
  <param name="remoteCommand">rm</param>
  <param name="targetName">%job_target_names%[2]</param>
  <param name="targetType">%job_target_types%[2]</param>
  <param name="args">-f, %app-home%/ %job_execution_id%.env</param>
  <param name="successStatus">0</param>
</paramList>
</step>
</stepset>
<displayInfo useDefaultCreateUI="true"/>
</jobtype>

```

Example 5

This example illustrates the use of the `repSQL` command to execute SQL statements and anonymous PL/SQL blocks against the Management Repository. The job type specification below calls a simple SQL statement in the first step S1, and a PL/SQL procedure in the second step. Note the use of the variables `%job_id%` and `%job_name%`, which are special job-system placeholders. Other job parameters can be similarly escaped as well. Also note the use of bind parameters in the SQL queries. The parameters `sqlinparam[n]` can be used to specify bind parameters. There must be one parameter of the form `sqlinparam[n]` for each bind parameter. Bind parameters must be used as far as possible to make optimum use of database resources.

```

<jobtype name="repSQLJob" editable="true" version="1.0">
  <stepset ID="main" type="serial">
    <step ID="s1" command="repSQL" >
      <paramList>
        <param name="sql">update mytable set status='executed' where
          name=?</param>
        <param name="sqlinparam1">%job_name%</param>
      </paramList>
    </step>
    <step ID="s2" command="repSQL" >
      <paramList>
        <param name="sql">begin mypackage.job_done(?,?,?); end;</param>
        <param name="sqlinparam1">%job_id%</param>
        <param name="sqlinparam2">3</param><param name="sqlinparam3">mgmt_rep</param>
      </paramList>
    </step>
  </stepset>
  <displayInfo useDefaultCreateUI="true"/>
</stepset>
</jobtype>

```

Example 6

This example illustrates the use of the switch stepset. The main stepset of this job is a switch stepset whose `switchVarName` is a job parameter called `stepType`. The possible values (`switchCaseVal`) that this parameter can have are "simpleStep", "parallel", and "OSJob", which will end up selecting, respectively, the step `SWITCHSIMPLESTEP`, the parallel stepset `SWITCHPARALLELSTEP`, or the nested job `J1`.

```

<jobType version="1.0" name="SwitchSetJob" editable="true">
  <stepset ID="main" type="switch" switchVarName="stepType" >
    <credentials>
      <credential usage="hostCreds" authTargetType="host"
        defaultCredentialSet="HostCredsNormal"/>
    </credentials>

```

```

<step ID="SWITCHSIMPLESTEP" switchCaseVal="simpleStep" command="remoteOp">

  <credList>
    <cred usage="defaultHostCred" reference="hostCreds" />
  </credList><paramList>
    <param name="remoteCommand">%command%</param>
    <param name="args">%args%</param>
    <param name="targetName">%job_target_names%[1]</param>
    <param name="targetType">%job_target_types%[1]</param>
  </paramList>
</step>
<stepset ID="SWITCHPARALLELSTEP" type="parallel" switchCaseVal="parallelStep">
  <step ID="P11" command="remoteOp" >
    <credList>
      <cred usage="defaultHostCred" reference="hostCreds" />
    </credList>
    <paramList>
      <param name="remoteCommand">%command%</param>
      <param name="args">%args%</param>
      <param name="targetName">%job_target_names%[1]</param>
      <param name="targetType">%job_target_types%[1]</param>
    </paramList>
  </step>
  <step ID="P12" command="remoteOp" >
    <credList>
      <cred usage="defaultHostCred" reference="hostCreds" />
    </credList>
    <paramList>
      <param name="remoteCommand">%command%</param>
      <param name="args">%args%</param>
      <param name="targetName">%job_target_names%[1]</param>
      <param name="targetType">%job_target_types%[1]</param>
    </paramList>
  </step>
</stepset>
<job ID="J1" type="OSCommandSerial" switchCaseVal="OSJob" >
  <paramList>
    <param name="command">%command%</param>
    <param name="args">%args%</param>
  </paramList>
  <targetList>
    <target name="%job_target_names%[1]" type="%job_target_types%[1]" />
  </targetList>
</job>
</stepset>
<displayInfo useDefaultCreateUI="true" />
</jobType>

```

Example 7

This example shows the use of the <securityInfo> tag to ensure that only users that have CLONE FROM privilege over the first target and MAINTAIN privilege over the second target will be able to submit jobs of the following type:

```

<jobType name="Clone" editable="true" version="1.0" >
  <securityInfo>
    <privilege name="CREATE TARGET" type="system" />
    <privilege name="CLONE FROM" type="target" evaluateAtSubmission="false" >
      <target name="%job_target_names%[1]" type="%job_target_types%[1]" />
    </privilege>
  </securityInfo>

```

```
<privilege name="MAINTAIN" type="target" evaluateAtSubmission="false">
  <target name="%job_target_names%[2]" type="%job_target_types%[2]" />
</privilege>
</securityInfo>
<!-- An optional <paramInfo> section will follow here, followed by the stepset
definition of the job
-->
<paramInfo>
  ....
</paramInfo>
<stepset ...>
  .....
</stepset>
<displayInfo useDefaultCreateUI="true"/>
</jobType>
```

Example 8

The following shows an example of a scenario where credentials are passed to a nested job in the job type xml:

```
<jobType version="1.0" name="SampleJobType001" singleTarget="true" editable="true"
defaultTargetType="host" targetTypes="all">
  <credentials>
    <credential usage="osCreds" authTargetType="host"
      defaultCredentialSet="HostCredsNormal" credentialTypes="HostCreds">
      <displayName nlsid="LABEL_NAME">OS Credentials</displayName>
      <description nlsid="LABEL_DESC">Please enter credentials.</description>
    </credential>
  </credentials>
  <stepset ID="main" type="serial">
    <step ID="Step" command="remoteOp">
      <credList>
        <cred usage="defaultHostCred" reference="osCreds" />
      </credList>
      <paramList>
        <param name="targetName">%job_target_names%[1]</param>
        <param name="targetType">%job_target_types%[1]</param>
        <param name="remoteCommand">/bin/sleep</param>
        <param name="args">1</param>
      </paramList>
    </step>
    <job ID="Nested_Job" type="OSCommand">
      <credList>
        <cred usage="defaultHostCred" reference="osCreds" />
      </credList>
      <targetList allTargets="true" />
      <paramList>
        <param name="command">/bin/sleep</param>
        <param name="args">1</param>
      </paramList>
    </job>
  </stepset>
</jobType>
```

7.16 About Performance Issues

This section provides a brief discussion on issues you should consider when designing your job type. These issues may impact the performance of your job type as well as the overall job system.

7.16.1 Using Parameter Sources

The following issues are important in relation to the use of parameter sources:

- Parameter sources are a convenient way to obtain needed parameters from known sources, such as the Management Repository or the credentials table. The parameter sources must be used only for quick queries that fetch information already stored somewhere else.
- Parameter sources that are evaluated at job execution time will, in general, effect the throughput of the job dispatcher and must be used with care. In some cases, the fetching of parameters at execution time may be unavoidable and if you do not care whether the parameters are fetched at execution time or submission time, set `evaluateAtSubmission` to `false`.
- When executing SQL queries to obtain parameters (using the SQL parameter source) the usual performance improvement guidelines apply. These include using indexes only where necessary and avoiding the joining of large tables.

7.17 Adding a Job Type to Enterprise Manager

To package a new job type with a metadata plug-in, you should adhere to the following implementation guidelines:

New job types packaged with a metadata plug-in will have two new files:

- Job type definition XML file: used by the job system during metadata plug-in deployment to define your new job type. There is one XML file for each job type.
- Job type script file: installed on selected Agents during metadata plug-in deployment. A single script may be shared amongst different jobs.

The following two properties must be set to "true" in the first line of the job type definition XML file:

- `agentBound`
- `singleTarget`

Here is an example:

```
<jobType version="1.0" name="PotatoUpDown" singleTarget="true" agentBound="true"
targetTypes="potatoserver_os">
```

Because the use of Java for a new job type is not supported for job types packaged with a metadata plug-in, new job types are `agentBound` and perform their work through a script delivered to the Agent (the job type script file). The job type definition XML file contains a reference to the job type script file and will execute it on the Agent whenever the job is run from the Enterprise Manager console.

Adding a Job Type to an Oracle Plug-in Archive (OPAR)

After you have created the job type definition XML file and modified the target type definition file, you can add your files to an Oracle Plug-in Archive (OPAR) just as you would any other target type. See [Chapter 13, "Validating, Packaging, and Deploying the Plug-in"](#) for more information.

Release 11.1 Job Types Versus Enterprise Manager Cloud Control 12c Job Types

In Oracle Enterprise Manager Cloud Control 12c, the job type parser has moved to an XSD-based parser. However, Enterprise Manager release 11.1 job types should work,

as there are no major changes required to enable an 11.1 job type to be parsed with a Cloud Control 12c parser.

The following are some of the known changes required by the Cloud Control 12c parser in the job type XML:

- `<jobtype>` should change to `<jobType>`.
- `<paramInfo>` should not contain `<stepset>`.
- `<ParameterUriSource>` tag should end like `<parameterUriSource attr1="" attr2="" />` and not like `<parameterUriSource attr1="" attr2=""> </parameterUriSource>`.
- `<paramInfo/>` should be removed.
- `stepSet` does not contain `successOf` or `failureOf` attributes.
- Make sure the ID specified in the `stepDisplayInfo` does exist in the job type (that is, a step with that ID should exist).

In Cloud Control 12c, job types can be registered through an `emctl` command, see the following command information:

```
emctl register oms metadata -service jobTypes -file <file name with absolute path> -sysman <sysman password> -pluginId <plugin id>
```

Defining a Management User Interface

Enterprise Manager can be extended to support the management of new domains through the introduction of discovery, monitoring, and automation. While the Enterprise Manager framework provides a powerful set of features related to these management capabilities, most plug-in developers need to expose management capabilities in a way that is appropriate to their domain. The Metadata Plug-in Custom User Interface (MPCUI) features of Enterprise Manager provide you with this capability.

This chapter contains the following sections:

- [Introduction to Defining a Management User Interface](#)
- [MPCUI Concepts](#)
- [UI Options for a Plug-in](#)
- [Creating the MPCUI Metadata File](#)
- [Defining Metadata](#)
- [Defining the MPCUI Application](#)
- [Packaging the MPCUI Implementation With the Plug-in](#)
- [Converting a Metadata-based UI to a Flex-based UI](#)
- [Defining System Home Pages](#)
- [Defining Navigation](#)
- [Accessing Enterprise Manager Data](#)
- [Performing Task Automation](#)
- [Defining Page Layout Components](#)
- [Including Packaged Regions](#)
- [Defining Charts](#)
- [Defining Tables](#)
- [Defining Dialogs](#)
- [Defining Trains](#)
- [Defining Information Item and Information Displays \(Label-Value Pairs\)](#)
- [Defining Links](#)
- [Displaying a Processing Cursor](#)
- [Defining a Processing Window](#)

- [Defining Icons for Target Types](#)
- [Displaying the Target Navigator](#)
- [Defining a UI for Guided Discovery](#)
- [About Logging](#)
- [Development Environment Options](#)
- [Migrating Home Page Customizations](#)
- [Accessibility Guidelines](#)
- [Localization Support](#)
- [Providing Online Help](#)

8.1 Introduction to Defining a Management User Interface

As a plug-in developer, you are responsible for the following steps for defining a custom user interface for managing your target types:

Note: In addition to this document, the Extensibility Development Kit (EDK) includes a complete sample implementation that should be used as a guide during this process.

1. Decide on the model for your target including:
 - Associations with other targets
 - Performance metrics and configuration data
 - Subcomponents of the target
 - Administrative tasks and operations
2. Familiarize yourself with the capabilities provided by the MPCUI library, such as:
 - UI components that are available (pages, charts, and so on)
 - Services that are available (metric data, SQL query, associations, task execution, and so on)
 - Difference in capabilities between a metadata-only implementation and a Flex implementation
 - Sample implementations and how they are constructed

3. Design the UI based on:
 - a. Data and tasks that are important
 - b. Capabilities provided by MPCUI

This can involve drawing the pages and describing their content, and reviewing the page with domain experts to ensure they expose the appropriate management capabilities.

4. Select the metadata-only or Flex implementation option.

For more information, see [Section 8.3, "UI Options for a Plug-in"](#).

Note: It is easy to migrate the metadata-only approach to a Flex implementation later if required. For more information, see [Section 8.8, "Converting a Metadata-based UI to a Flex-based UI"](#).

5. Create the target metadata for the items in your design (see step 1). This metadata is necessary to implement your UI later. For more information about target metadata, see the relevant chapters within this guide.
6. Develop the SQL queries required to retrieve configuration data that will be displayed in the UI. Typically, these queries reference the configuration CM\$ views. For more information about configuration data, see [Chapter 6](#).
7. Identify and define the activities that make up your UI, such as pages, wizards, and dialogs. The Integration metadata defines these activities. For more information, see [Section 8.5.2, "Defining Integration Metadata"](#).
8. If you selected the Flex implementation option, then continue with the steps in [Section 8.1.1, "Flex Implementation"](#). Otherwise, continue with the steps in [Section 8.1.2, "Metadata-only Implementation"](#).

8.1.1 Flex Implementation

If you are using the Flex implementation option, then you are responsible for the following steps:

1. Obtain a copy of Adobe Flash Builder or download the Adobe Flex Software Development Kit (SDK).

For more information, see [Section 8.30, "Development Environment Options"](#).

Note: The Adobe Flex SDK is free but it does not provide graphical editing or debugging capabilities.

2. Create a project (if using Adobe Flash Builder) to hold the source code for your custom UI. You can use the sample project included in the EDK as a template. Ensure that the project settings are correct. For more information, see [Section 8.30.2, "Developing MPCUI in Adobe Flash or Flex Builder"](#).
3. Implement an MXML class that extends the `MpApplication` class. This is the Flex application class. For more information, see [Section 8.6, "Defining the MPCUI Application"](#).
4. Implement an MXML class that extends the `Integration` class. This defines the set of activities included in the custom UI. For more information, see [Section 8.6.1, "Defining the Application Activities \(Integration Class\)"](#).
5. Develop each activity (such as page or dialog). Typically, each page includes a page class (written in MXML, extending the `Page` class) and a controller class (written in ActionScript extending the `ActivityController` class). For more information, see [Section 8.6.2, "Defining Pages"](#), [Section 8.6.3, "Defining Dialogs"](#), and [Section 8.6.4, "Defining Trains and Train Pages"](#).

6. Build and test your custom UI from Adobe Flash Builder.

Note: You must deploy at least one version of your plug-in before building and testing. The deployed plug-in must include the target metadata (such as metrics and configuration data). However, the plug-in does not have to include your MPCUI metadata for testing.

7. Create the MPCUI metadata file.

This file includes:

- SQL statements used by your custom UI
- Menu items you want to include to support navigation to different pages defined in your UI
- Reference to the Flex UI that you built

For more information, see [Section 8.4, "Creating the MPCUI Metadata File"](#).

8. Modify your plug-in to include the MPCUI metadata file and the SWF file built in Adobe Flash Builder.

Place these files in the oms/metadata/mpcui directory of the plug-in staging area.

For more information, see [Section 8.7, "Packaging the MPCUI Implementation With the Plug-in"](#).

9. Test your custom UI by accessing a target home page from the Enterprise Manager console.

This loads your custom UI in the context of the Enterprise Manager application and displays the Enterprise Manager application and target menus.

In addition to this document, additional resources for developing with Flex components are provided:

- The API reference: This is located in your partner EDK directory under doc/sdk_api_ref.html
- The HostSample example plug-in: The sample plug-in provided by Oracle provides examples of many MPCUI features. It is located in the EDK under samples/plugins/HostSample

You may also include any of the base Flex components (e.g. Button, Label, etc.) Oracle does not provide the documentation for Flex components as part of the EDK, but you can find the documentation online at the following link:

<http://livedocs.adobe.com/flex/3/html/index.html>.

8.1.2 Metadata-only Implementation

If you are using the metadata-only implementation option, then you are responsible for the following steps:

1. Create the MPCUI metadata file.

This file includes:

- SQL statements used by your custom UI
- Menu items you want to include to support navigation to different pages defined in your UI

- All the metadata definitions discussed in the following steps
For more information, see [Section 8.4, "Creating the MPCUI Metadata File"](#).
- 2. Add the integration metadata to the MPCUI metadata file.
The integration metadata defines the set of activities included in the custom UI.
For more information, see [Section 8.5.2, "Defining Integration Metadata"](#).
- 3. Add the page definitions (ActivityDefinitions) to the MPCUI metadata file.
For more information, see [Section 8.5.3, "Defining Navigation"](#).
- 4. Modify your plug-in to include the MPCUI metadata file.
Place these files in the oms/metadata/mpcui directory of the plug-in staging area.
For more information, see [Section 8.7, "Packaging the MPCUI Implementation With the Plug-in"](#).
- 5. Test your custom UI by accessing a target home page from the Enterprise Manager console.
This loads your custom UI in the context of the Enterprise Manager application and displays the Enterprise Manager application and target menus.

8.1.3 Assumptions and Prerequisites

This chapter assumes you are familiar with the following:

- Plug-in development overview, including how to package a plug-in and its XML files.
- The XML-based user interface markup language known as MXML.

8.2 MPCUI Concepts

There are several important concepts that should be understood when using the MPCUI framework. These concepts are defined briefly in this section and discussed in more detail in the subsequent sections.

8.2.1 Integration Class

The integration class is the bootstrap for your application, and is used to define the set of pages, dialogs, and trains that are included in the application. The MPCUI framework uses the information included in the integration class to drive the application including managing navigation between UI elements.

8.2.2 Activity

Top-level UI elements in the MPCUI are referred to generally as activities. Activities include pages, dialogs, trains and train pages, URLs, and jobs.

8.2.3 Page

Flex does not include the notion of a page, though this is a construct that is provided by the MPCUI framework to simplify the construction of the UI and make it fit more naturally into the larger Enterprise Manager console.

The MPCUI framework manages pages within the application, providing simple navigation between pages and integrating them into the browser history and the Enterprise Manager menu system.

8.2.4 Services

The MPCUI framework provides a series of services that can be used to retrieve data from the Management Server or to process actions (jobs or remote operations).

8.2.4.1 Data Services

The Data Services provided by MPCUI include data services to retrieve metric data, associations, target properties and so on. It includes a `SQLDataService` that can be used to run named SQL statements within the plug-in.

8.2.4.2 Operation Services

MPCUI includes a Job service and RemoteOp service that can be used to perform administrative actions against the targets managed by the plug-in code.

- The Job service requires the inclusion of job type definitions in the plug-in
- The RemoteOp service requires the registration of scripts with the plug-in framework

8.2.4.3 Asynchronous Service Request Handling

The Adobe Flex and Adobe Flash framework (and therefore the MPCUI framework) handles network requests asynchronously. This requires the use of a result handler pattern where a request is made to the server and as part of the request, a handler (or callback) is registered with the request. Upon completion of the request (or if a fault occurs), the handler is called and passed the result.

8.2.5 URL

MPCUI provides a number of different capabilities related to the generation of URLs and the ability to embed links to:

- Other Enterprise Manager pages
- Other pages within the MPCUI application
- External pages

Because MPCUI is a Flex application, it is not quite as easy as embedding a link to a URL. For more information about URLs, see [Section 8.10.2, "URL and Links"](#).

8.3 UI Options for a Plug-in

The following UI options for a metadata plug-in are:

1. Use the default home page with limited information and no customization.
2. Convert a metadata plug-in (release 10.2 or 11.1) that includes home page customizations. Provides a home page with the same information as release 11.1 but no customization. For more information, see [Section 8.31, "Migrating Home Page Customizations"](#).
3. Metadata-only implementation using MPCUI. Provides a customized UI but with restrictions. For more information, see [Section 8.3.1, "Metadata-only Implementation"](#).

Note: You can implement the UI using metadata only. Then, if you want to add features that are only available in the Flex implementation, evolve the metadata-only implementation to create a Flex implementation. For more information, see [Section 8.8, "Converting a Metadata-based UI to a Flex-based UI"](#).

4. Flex implementation using MPCUI. A more complex implementation but provides most flexibility and features. For more information, see [Section 8.3.2, "Flex Implementation"](#).

8.3.1 Metadata-only Implementation

MPCUI metadata is XML that describes the layout of the UI and the binding to Enterprise Manager services. For more information about MPCUI metadata, see [Section 8.4, "Creating the MPCUI Metadata File"](#)

Use the Demo Host Sample (demo_hostsample) as a starting point or else you can develop a new UI. The UI must include at least one page (the home page), and can optionally include other pages. Each page definition is included in the MPCUI metadata file.

In addition to the metadata description of each page, the metadata must also include an integration definition. For more information about integration definitions, see [Section 8.5.2, "Defining Integration Metadata"](#).

8.3.2 Flex Implementation

The Flex implementation option provides additional capabilities for providing a customized UI on top of administrative capabilities included in the plug-in as jobs or as Agent scripts.

While one of the goals of the MPCUI framework is to provide a simplified layer of abstraction over the Flex framework with which it is implemented, you must become familiar with the Flex framework and how to develop using the Flex framework.

- [MXML](#)
- [ActionScript](#)
- [SWF Binary File](#)

8.3.2.1 MXML

Flex includes a tag language (MXML) that can be used to lay out the user interface and bind the UI components to data elements. Much of what you do using MPCUI can be accomplished in MXML.

Note: MXML is used for the metadata-only and Flex-based implementations.

- For the metadata-only implementation, use MXML to define all the pages within the MPCUI metadata file. For more information, see [Section 8.5.2.1, "Defining Pages"](#).
 - For the Flex-based implementation, save the metadata for each page in a separate MXML file to compile into a SWF file. For more information, see [Section 8.6.2, "Defining Pages"](#).
-

8.3.2.2 ActionScript

For cases that require more complex handling of data or events, you might have to develop part of the UI using ActionScript (the ECMA-script compliant programming language). Developers familiar with Java should become comfortable with ActionScript quickly.

8.3.2.3 SWF Binary File

When building a Flex application, the MXML and ActionScript are compiled to form a binary file. This binary format (*.swf) is included in the plug-in and interpreted by the Adobe Flash Player at run time. The Enterprise Manager extensibility framework and the MPCUI framework, in particular, handle the integration of the SWF with the necessary Enterprise Manager wrapper page and handle rendering of that page at run time. You do not have to construct additional Application Development Framework (ADF), HTML, or JavaScript to enable the display of your custom UI.

8.3.2.4 Defining the Home Page

Build the UI by defining pages and custom UI using the MPCUI components and services, and building an Adobe Flex application that is shipped as part of your plug-in. This option provides flexibility and control over the UI, but also requires additional effort to understand the components and services provided by the MPCUI framework.

If the modifications to the home page template are simple, take advantage of the MPCUI framework ability to ship metadata to describe the page, thus avoiding the requirement to build and package a Flex application (MPCUI SWF).

8.4 Creating the MPCUI Metadata File

Each plug-in that includes MPCUI must include an MPCUI metadata file.

The metadata file:

- Defines SQL queries required by the MPCUI
- Defines the menu items required by the MPCUI
- Contains UI metadata (the layout of the custom UI) (Metadata-based option only)
- Specifies the SWF (file name that includes the MPCUI) (Flex-based option only)
- Specifies target icons, target navigator, and system home page options
- Specifies the Discovery SWF (file name that includes the Guided Discovery UI) (Flex-based option only)

For more information about the syntax for this file, see the XSD file located in the Extensibility Development Kit (EDK) specifications.

[Example 8-1](#) provides a summary of the metadata-based UI MPCUI metadata file and [Example 8-2](#) provides a summary of the Flex UI metadata file

Example 8-1 MPCUI Metadata File for Metadata-based UI

```
<CustomUI target_type="demo_hostsample"
          xmlns="http://www.oracle.com/EnterpriseGridControl/MpCui">

  <!-- SqlStatements defines the individual SQL statements that are used by
       the MPCUI code. Each statement is identified by a unique name and
       can only be referenced by that name from the MPCUI code itself -->
```



```

    <SqlStatements>
      <Sql name="INSTANCE_INFO">
        select * from...
      </Sql>
    </SqlStatements>

  <UIMetadata>
    <Integration>
      .....
    </Integration>

    <ActivityDefinition>
      .....
    </ActivityDefinition>

  </UIMetadata>

  <!-- MenuMetadata defines the set of menu items that should appear in the
       target menu on the homepage and specifies which of the MPCUI pages
       should be accessed from that menu item -->
  <MenuMetadata>
    <menu label="Host Sample">
      <menuItem>
        <command .. />
      </menuItem>
    </menu>
  </MenuMetadata>

  <EmuiConfig>
    <context-pane-visible>true</context-pane-visible>
    <large-icon>dhs_large.png</large-icon>
    <small-icon>dhs_small.png</small-icon>
    <use-framework-homepage>true</use-framework-homepage>
  </EmuiConfig>

</CustomUI>

```

Example 8-2 MPCUI Metadata File for Flex-based UI

```

<CustomUI target_type="demo_hostsample"
  xmlns="http://www.oracle.com/EnterpriseGridControl/MpCui">

  <!--
  SQL Statements to be used by the custom UI code. All bind variables should be
  identified using "?VAR?" type notation, and can then be referenced using
  either the SQLDataService MXML tag or using the SQL or BatchSQL services.
  -->
  <SqlStatements>
    <Sql name="INSTANCE_INFO">
      select * from...
    </Sql>
  </SqlStatements>

  <!--
  SwfFiles tag is used to register the Flex application (must extend
  MpApplication) that includes the custom UI for the plug-in.
  The SWF file registered must be included in the plug-in along with this
  meta-data in the oms/metadata/mpcui directory.
  -->

```

```

<SwfFiles>
  <Swf is_homepage="true">HostSample.swf</Swf>
  <Swf discovery_module="DemoHostSample">HostSampleDiscovery.swf</Swf>
</SwfFiles>

<!-- MenuMetadata defines the set of menu items that should appear in the
      target menu on the homepage and specifies which of the MPCUI pages
      should be accessed from that menu item -->
<MenuMetadata>
  <menu label="Host Sample">
    <menuItem>
      <command .. />
    </menuItem>
  </menu>
</MenuMetadata>

<EmuiConfig>
  <large-icon>dhs_large.png</large-icon>
  <small-icon>dhs_small.png</small-icon>
</EmuiConfig>

</CustomUI>

```

8.4.1 Overview of MPCUI Metadata Elements

[Table 8–1](#) describes the key elements that define the discovery metadata.

Table 8–1 Key Elements Used to Define Discovery Metadata

Element	Description
SqlStatements	The <code>SqlStatements</code> element contains the SQL statements that enable you to access information stored in the Management Repository. For more information about these SQL statements, see Section 8.11.4, "Packaged SQL and the Query Service" .

Table 8–1 (Cont.) Key Elements Used to Define Discovery Metadata

Element	Description
UIMetadata (Metadata-based UI only)	<p>The UIMetadata element is the top-level container for the integration and page (activity) definitions described by that metadata:</p> <pre><UIMetadata> <!-- The meta-data only definition must include an Integration element which defines the set of activities (pages, dialogs, etc.) that make up the application --> <Integration> <intg:Integration targetType='demo_hostsample' xmlns:intg="oracle.sysman.emx.intg" > .. </intg:Integration> </Integration> <!-- The meta-data only definition must include 1 or more ActivityDefinition elements each of which defines an activity (e.g. page, dialog, etc.) --> <ActivityDefintion> <intg:Page id="homePg" label="Home Page" ...> </intg:Page> </ActivityDefintion> </UIMetadata></pre>
Integration (Metadata-based UI only)	<p>The Integration element defines the integration metadata used to specify the set of pages and to define task flows between these pages (if required). For information about integration metadata, see Section 8.5.2, "Defining Integration Metadata".</p>
ActivityDefinition (Metadata-based UI only)	<p>The ActivityDefinition element defines the navigation between pages in the UI. For information about defining navigation, see Section 8.5.3, "Defining Navigation".</p>
SwfFiles (Flex UI only)	<p>The SWFfiles element that specifies the name of the SWF file.</p>
MenuMetadata	<p>The MenuMetadata element includes the menuItem elements that define navigation to activities defined in the MPCUI metadata. For more information about the MenuMetadata element, see Section 8.5.3, "Defining Navigation".</p>
EmuiConfig	<p>The EmuiConfig element includes elements to define the following</p> <ul style="list-style-type: none"> ■ Target navigator (context-pane-visible) <p>For more information, see Section 8.27, "Displaying the Target Navigator".</p> ■ Icons to represent target types in the Cloud Control console (large-icon, small-icon) <p>For more information, see Section 8.26, "Defining Icons for Target Types".</p> ■ System home page (use-framework-homepage) <p>For more information, see Section 8.9, "Defining System Home Pages".</p>

8.5 Defining Metadata

For a complete example of an MPCUI metadata implementation, see the Demo Sample implementation (demo_hostsample_uimd_fullmd.xml) provided with the Extensibility Development Kit (EDK).

8.5.1 Limitations of the Metadata Implementation

This implementation supports the definition of pages only, and does not support the definition of dialogs or trains. If the custom UI requires dialogs or trains, then you must use the Flex-based option for building your home page. For more information, see [Section 8.3.2, "Flex Implementation"](#).

The ability to perform manipulation of data for display or the ability to respond to some UI events and to invoke jobs or remote operations is not available in this implementation.

8.5.2 Defining Integration Metadata

Use the integration metadata to specify the set of pages and to define task flows between these pages (if required).

Example 8–3 Integration Metadata

```
<Integration>
  <mp:Integration targetType="demo_hostsample"
xmlns:mp="http://www.oracle.com/mpcui">
    <mp:PageActivityDef id="homePg" label="Home" isDefaultPage="true" />

    <mp:PageActivityDef id="perfPg" label="Performance" />
    <mp:PageActivityDef id="processesPg" label="Processes" />
    <mp:PageActivityDef id="adminPg" label="Configuration" />
    <mp:DialogActivityDef id="detailsDialog" label="Metrics Detail" />
    <mp:DialogActivityDef id="metricHistory" label="Metric History">
      <mp:inputParams>
        <mp:InputParam name="targetName" />
        <mp:InputParam name="targetType" />
        <mp:InputParam name="metric" />
        <mp:InputParam name="columns" />
        <mp:InputParam name="period" />
        <mp:InputParam name="title" />
      </mp:inputParams>
    </mp:DialogActivityDef>
    <mp:DialogActivityDef id="metricDetails" label="Metric Details">
      <mp:inputParams>
        <mp:InputParam name="targetName" />
        <mp:InputParam name="targetType" />
        <mp:InputParam name="metric" />
        <mp:InputParam name="columns" />
        <mp:InputParam name="period" />
        <mp:InputParam name="title" />
      </mp:inputParams>
    </mp:DialogActivityDef>
  </mp:Integration>
</Integration>
```

8.5.2.1 Defining Pages

The page metadata defines the layout of a page including the components that make up the page and the data that is displayed on the page. Each page is contained within its own XML and must be registered with the home page Metadata Registration Services (MRS) using the same activity identifier specified for the page in the integration metadata. For more information about the integration metadata, see [Section 8.5.2, "Defining Integration Metadata"](#).

Note: The tag language used to define pages is a subset of the tag language supported for the Flex-based implementation. You can take a page defined in a metadata file and turn it into a page that is part of an MPCUI application (SWF file) by:

- Changing the file extension from XML to MXML
- Adding the appropriate MXML namespace at the top of the file

For more information, see [Section 8.8, "Converting a Metadata-based UI to a Flex-based UI"](#).

[Example 8-4](#) is a partial sample of a metadata page definition:

Example 8-4 Metadata Page Definition

```
<ActivityDefinition>
  <!--
    Each page included in the plugin UI should extend the Page class and be
    coded in MXML.
    The page file specifies the layout of the page, declares some of the data
    binding (see below) and specifies handlers for events that are initiated in
    the page, when a user clicks a button or link for example. The page also
    has a controller class (that extends PageController) that is associated with
    the page. The controller loads data shown in the page and includes
    functions that are called as event handlers.
  -->
  <mp:Page id="homePg" label="Home Page"
    xmlns:mx="http://www.adobe.com/2006/mxml"
    xmlns:mp="http://www.oracle.com/mpcui" >

    <!--
      Data Services - these are sources of data that will be shown in the
      page. Data can either be bound from a data service declared here, or it
      may be loaded within the controller.
    -->
    <mp:services>
      <!--
        SQLDataService - this service allows you to execute a SQL query packaged
        with your plugin and then refer to the result set from the query
        execution. Properties passed to the query are declared as name-value
        pairs. If the properties are runtime/dynamic properties then
        you will have to use the SQL service within the controller, load the
        data there and then map the result set to the page model.
      -->
      <mp:SQLDataService id="ids" queryID="INSTANCE_INFO"
        properties="{props('TARGET_GUID', appModel.target.guid)}" />

      <mp:SQLDataService id="cht1" queryID="CHTSQL1" properties="{props('HC_
        TARGET_GUID', appModel.target.guid)}" />
```

```

<!--
    MetricValuesDataService - this service allows you to obtain data for a
    metric, for some period of time. This time period may be a historical
    time period, or it may be REALTIME which creates a data service that
    will poll for the current value of the metric through the Agent.
-->
<mp:MetricValuesDataService id="mv1"
    flattenData="true"
    targetName="{appModel.target.name}" targetType="{appModel.target.type}"
    metricName="CPUProcessorPerf"
    columns="{['CPUUser','CPUIdle']}"
    timePeriod="LAST_DAY" />

<!--
    AvailDataService - this service obtains target availability, that
    includes current status, availability for the last 24 hours, and up
    since time.
-->
<mp:AvailDataService id="ads" targetName="{appModel.target.name}"
targetType="{appModel.target.type}" />

<!--
    AssociationService - this service obtains associated targets
-->
<mp:AssociationDataService id="asc" targetName="{appModel.target.name}"
targetType="{appModel.target.type}" assocTypes="{['hosted_by']}" />
</mp:services>

<!--
    Page Content - the page should be laid out in a grid pattern using a
    combination of columns (VBox) and rows (HBox). Also to ensure property
    sizing/resizing behavior relative height/width should be used in
    percentages.
-->
<mx:VBox width="100%" height="100%">
    <!--
        1st Row - will occupy 30% of the height of the page and includes a
        Summary region, Availability region and Job Summary region. The
        Availability and Job Summary region require no parameters. The
        Summary region uses the InfoDisplay component to show a series of
        name-value pairs. Each item may also specify an optional
        destination and image.

        The example below also demonstrates the ways data may be bound to a
        UI component included in the page:

        1. Data Service Reference
        2. Global/Application Model Reference
        3. Page Model Reference
        4. Set Directly from Controller
    -->
    <mx:HBox width="100%" height="30%">
        <mp:Region title="Summary" width="25%" height="100%" >
            <mp:InfoDisplay id="summaryInfo">
                <mp:InfoItem label="CPU Model"
value="{ids.result.getString(0,'CPU Model')}" /> <!-- ref to SQLDataService -->
                <mp:InfoItem label="Target Name"
value="{appModel.target.name}" /> <!-- ref to global/application model -->
                <mp:InfoItem label="Current Status"

```

```

value="{ads.currentStatus}" image="{ads.currentStatusIcon}" /> <!-- ref to
AvailDataService -->
    <mp:InfoItem source="{ads.statusSinceItem}" /> <!-- ref to
AvailDataService -->
    <!-- <mp:InfoItem label="{model.osVersLabel}"
value="{model.osVersion}" /> --> <!-- ref to page model; model set in controller
in SQL svc handler -->
    <!-- <mp:InfoItem id="infoItem" label="Controller Set" />
--> <!-- value property set directly in controller -->
    <mp:InfoItem label="Hosted By"
value="{asc.assoc.getAssoc('hosted_by').name}" /> <!-- ref to
AssociationService -->
    </mp:InfoDisplay>
</mp:Region>

<!--
<mp:AvailabilityRegion width="33%" height="100%" daySpan="1" />
-->

    <mp:Region title="Memory Usage (Last 24 Hrs)" width="45%"
height="100%">
        <mp:AreaChart id="memHist" width="100%" height="100%"
            metricName="MemoryPerf"
            metricColumns="['Active','MemFree']"
            timePeriod="LAST_DAY" />
        <!-- <mp:Link label="Current"
click="{controller.showCpuMetricDetails(event)}" /> -->
    </mp:Region>

    <mp:Region title="Memory Used (Current)" width="30%" height="100%">
        <mp:LineChart id="memRt" width="100%" height="100%"
            metricName="MemoryPerf"
            metricColumns="['Active']"
            timePeriod="REALTIME"
            interval="15"/>
        <!-- <mp:Link label="History"
click="{invokeActivity('metricHistory',
    bean('targetName', appModel.target.name, 'targetType',
appModel.target.type,
    'metric', 'Response', 'columns', ['Load'], 'period', 'LAST_DAY',
'title', 'Metric History'))}" /> -->
    </mp:Region>

    <!-- <mp:JobSummaryRegion width="25%" height="100%" /> -->
</mx:HBox>

<!--
2nd row - will occupy 35% of the overall page height and shows three
charts and shows the ability to access other activities (pages,
dialogs, etc.)

The 1st chart shows a line chart that displays a metric in
real-time.
It will automatically start polling the metric value in the
background and will continue to update the chart on the page until
the page is not shown. The 2nd chart shows a line chart that
displays a metric history. The 3rd chart shows a barch chart
showing metric data grouped by the key in the data, in this case the
CPU #.

```

Each region also includes a Link component and shows the ability to navigate to other activities. This may be any activity (page, dialog, train, URL, job). The 1st chart shows using the `invokeActivity` method being called directly from the page and passing context to the activity using the bean method to form the input context for the `metricDetails` activity.

The 2nd link shows calling a function in the controller, and then navigating to another activity from within the controller. The final link shows the use of the `invokeActivity` method again, however shows navigating to an activity that requires no additional context (the `Processes` page).

```
-->
<mx:HBox width="100%" height="35%">

    <mp:Region title="Per Processor Idle Time (%)" width="25%"
height="100%">
        <mp:BarChart id="bchart" timePeriod="LAST_DAY" width="100%"
groupBy="byKey" metricName="CPUProcessorPerf" metricColumns="{['CPUIdle']}" />
        <mp:Link label="Show Processes"
click="{invokeActivity('processesPg')}}" />
    </mp:Region>

    <mp:Region title="CPU Utilization % (Last 24 Hrs)" width="45%"
height="100%">
        <mp:LineChart id="cpuutil" width="100%" height="100%"
            metricName="CPUPerf"
            metricColumns="['system','idle','io_wait']"
            timePeriod="LAST_DAY" />
        <mp:Link label="Current"
click="{controller.showCpuMetricDetails(event)}" />
    </mp:Region>

    <mp:Region title="CPU Load (Current)" width="30%" height="100%">
        <mp:LineChart id="cpuload" width="100%" height="100%"
            metricName="Response"
            metricColumns="['Load']"
            timePeriod="REALTIME"
            interval="15" />
        <mp:Link label="History" click="{invokeActivity('metricHistory',
            bean('targetName', appModel.target.name, 'targetType',
appModel.target.type,
            'metric', 'Response', 'columns', ['Load'], 'period', 'LAST_DAY',
'title', 'Metric History'))}" />
    </mp:Region>

    <!--
    <mp:Region title="Per Processor Idle Time (%)" width="34%"
height="100%">
        <mp:BarChart id="bchart" timePeriod="LAST_DAY" width="100%"
groupBy="byKey" metricName="CPUProcessorPerf" metricColumns="{['CPUIdle']}" />
        <mp:Link label="Show Processes"
click="{invokeActivity('processesPg')}}" />
    </mp:Region>
    <!--

</mx:HBox>

<!-- 3rd row - events region -->
<mx:HBox width="100%" height="35%">
```



```

        <mp:IncidentRegion width="75%" height="100%" />
        <!--
        <mp:Region title="Memory Details" width="25%" height="100%" >
            <mx:ComboBox id="selMemChart" dataProvider="{model.memChoices}"
labelField="choiceLabel" change="{controller.changeMemChart(event)}" />
            <mp:PieChart id="memChart" targetName="{appModel.target.name}"
targetType="{appModel.target.type}"
                metricName="MemoryPerf"
                metricColumns="{model.memoryColumns}"
                timePeriod="REALTIME" interval="15" />
        </mp:Region>
        -->
        <mp:JobSummaryRegion width="25%" height="100%" />
    </mx:HBox>

    </mx:VBox>
</mp:Page>

</ActivityDefinition>

```

8.5.2.2 Mapping Data to UI Components

Use one of the following options to specify the data to be mapped to UI components in the metadata page definitions:

- For components that support properties that specify the data to be included in the component. For example, the chart component supports properties to specify metric names and columns that are shown in the chart.

```

<c:LineChart id="memRt" width="100%" height="100%"
    metricName="MemoryPerf" metricColumns="['Active']"
    timePeriod="REALTIME" interval="15"/>

```

In the preceding example, `LineChart` includes properties that specify which metric should be displayed in the chart. The MPCUI framework retrieves the data from the Management Server to populate the chart.

- For data binding using the data service tag. The data service tag has several forms, including the following:

Note: You must declare the data services that will be used within an activity (page) at the top of the page definition:

```

<intg:Page id="homePg" label="Home Page"
    ....
    <intg:services>
        ....
    </intg:services>

```

– MetricValuesDataService

The `MetricValuesDataService` tag provides the ability to include metric data, either real-time or historical, from the Management Server. Then it binds that data to the UI components.

```

<ds:MetricValuesDataService id="procData"
    flattenData="true"
    targetName="{appModel.target.name}" targetType="{appModel.target.type}"

```

```
metricName="CPUProcessorPerf"
columns="{['CPUIdle']}"
timePeriod="REALTIME"interval="15" />
```

– SQLDataService

The `SQLDataService` tag provides the ability to run a packaged SQL statement and bind the columns included in the `resultSet` to the UI component.

```
<ds:SQLDataService id="ids" queryID="INSTANCE_INFO"
  properties="{props('TARGET_GUID',appModel.target.guid)}" />
```

After you declare a `DataService` for the page, components within the page can reference the data provided by the service:

```
<components:InfoItem label="CPU(0) Idle %"
  value="{procData.result.getString('CPU1','CPUIdle')}" />

<components:InfoItem label="CPU Model"
  value="{ids.result.getString(0,'CPU Model')}" />
```

For more information about these tags, including the structure of the data returned by each, how parameters are set on the tags, and which UI components support easy integration of the data returned from these services, see [Section 8.12, "Performing Task Automation"](#).

Note: There are a number of data services that are supported from the metadata-only or Flex-based implementations, including `MetricValuesDataService`, `SQLDataService`, `AssociationDataService`, and `AvailDataService`.

Finally, there are a number of common data items that are available to be mapped to the metadata components. These items contain properties that can be:

- Displayed directly (for example, `appModel.target.name`)
- Used as parameters to Data Services
- Used as a bean input to an activity to which it is being navigated

Common data items include the `appModel` property. The `appModel` property includes static properties associated with the application runtime including `target` (and all its properties, see `oracle.sysman.emx.model.Target`) for the target the application is being rendered.

Reference the `appModel.target` properties from UI components in either the metadata-only or Flex implementation by using notation similar to:

```
<mp:InfoItem label="Target Name"
  value="{appModel.target.name}" />
```

In this case, the following appears in the UI:

Target Name MyTargetName

8.5.3 Defining Navigation

The metadata UI definition support included in MPCUI is limited to the definition of pages, and does not support other activities such as dialogs, trains, jobs, and so on.

Therefore, the only navigation possible between pages in the UI is by one of the following:

- Defining a menu item in the metadata that can be used to access a page

The MenuMetadata item includes the menuItem elements that define navigation to activities defined in the MPCUI metadata. For example, if the metadata includes the following page definition:

```
<ActivityDefinition>
  <intg:Page id="processesPg" label="Processes" ..>
    <!-- the body of the processes page would be declared here ?
  </intg:Page>
</ActivityDefinition>
```

Specify a menuItem in the MenuMetadata element to allow navigation to the previous page:

```
<menuItem>
  <command id="processesPg" label="Processes"
    class="oracle.sysman.emSDK.pagemodel.menu.EMNavigationMenuCommand
    partialSubmit="true" >
    <property name="actionOutcome" value="goto_core-mpcustom-nav" />
  <property name="paramsMap">
  <mapEntry name="pageid" value="processesPg" />
    </property>
  </command>
</menuItem>
```

The key properties in the menuItem element are:

- label within the command element.
label specifies the label that appears in the target menu on the home page. In the example given, a menu item “Processes” would be included.
- the value specified for the actionOutcome property.
actionOutcome specifies the view ID for the page containing the SWF file.

- Navigating from within a page using the invokeActivity directive

Use the invokeActivity directive to navigate between activities defined in an MPCUI metadata implementation. Associate this directive with the click property of the Link or Button components. When the end user clicks one of these components, the click property specifies the action that should be taken.

The invokeActivity directive takes the following:

- one required parameter (the activity id)
The activity id specifies the activity to which control should be passed
- one optional parameter (a bean to provide input context to the activity)
The input context specifies information to be passed to the activity.

For example, if the implementation includes two activities, 'homePg' and 'processesPg' (two pages). Include a link in the home page that when clicked, it changes the display to the processes page.

```
<ActivityDefinition>
<intg:Page id="homePg" label="Home Page"
...
<mx:Link label="Show Processes" click="{invokeActivity('processesPg')}" />
```

Use the activity content parameter to define an activity that is parameterized and therefore can be invoked from different contexts. The requirement for parameter input must be specified as part of the activity definition included in the metadata. For example, suppose you want a dialog activity that can show a historical line chart for any number of different metrics, and possibly even different targets. In the integration metadata the activity definition would appear similar to [Example 8-5](#):

Example 8-5 Defining Activity

```
<intg:DialogActivityDef id='metricHistory' label='Metric History' >
  <intg:inputParams>
    <intg:InputParam name='targetName' />
    <intg:InputParam name='targetType' />
    <intg:InputParam name='metric' />
    <intg:InputParam name='columns' />
    <intg:InputParam name='period' />
    <intg:InputParam name='title' />
  </intg:inputParams>
</intg:DialogActivityDef>
```

The `inputParams` elements specify the input parameters to the dialog (activity). Then from another page activity, two different links can direct to the same dialog, but with different parameters:

```
<comp:Region title="Memory Used (Current)" width="30%" height="100%">
  <c:LineChart id="memRt" width="100%" height="100%"
    metricName="MemoryPerf" metricColumns="['Active']"
    timePeriod="REALTIME" interval="15"/>
  <comp:Link label="History"
    click="{invokeActivity('metricHistory',
      bean('targetName', appModel.target.name,
        'targetType', appModel.target.type,
        'metric', 'Response',
        'columns', ['Load'],
        'period', 'LAST_DAY',
        'title', 'Memory Used (History)'))}" />
</comp:Region>

<comp:Region title="CPU Used (Current)" width="30%" height="100%">
  <c:LineChart id="memRt" width="100%" height="100%"
    metricName="CPUProcessorPerf" metricColumns="['CPUUser','CPUIdle']"
    timePeriod="REALTIME" interval="15"/>
  <comp:Link label="History"
    click="{invokeActivity('metricHistory',
      bean('targetName', appModel.target.name,
        'targetType', appModel.target.type,
        'metric', 'CPUProcessorPerf',
        'columns', ['CPUUser','CPUIdle'],
        'period', 'LAST_DAY',
        'title', 'CPU Used (History)'))}" />
</comp:Region>
```

8.6 Defining the MPCUI Application

The basis for the custom UI built using the MPCUI framework requires the construction of a Flex-based application. To simplify this process, the framework

provides a series of base classes and structures. The starting point for this development is to extend the `MpApplication` class to define the application.

The application file is a simple MXML file that implements a single method, `getIntegrationClass()`, which returns an instance of the integration class associated with this application. The integration class (described in [Section 8.6.1, "Defining the Application Activities \(Integration Class\)"](#)) defines the set of activities (such as pages, dialogs, and trains) that make up the application.

When compiled, the application binary (SWF file) is called by the same name as the application source file (by default). While you can call the application anything meaningful, Oracle recommends that the application class has the same name as the target type that it supports.

Example 8–6 Application MXML

```
<?xml version="1.0" encoding="utf-8"?>

<intg:MpApplication xmlns:mx="http://www.adobe.com/2006/mxml"
  xmlns:intg="oracle.sysman.emx.intg.*"
  backgroundColor="#EFF3F7" preloader="oracle.sysman.emx.MpPreloader" >
  <mx:Script>
    <![CDATA[
      /* Must override the getIntegrationClass method and
      return the class that extends Integration */
      override public function getIntegrationClass():Class
      { return HostSampleInteg; }
    ]]>
  </mx:Script>
</intg:MpApplication>
```

8.6.1 Defining the Application Activities (Integration Class)

The integration class defines the set of UI elements that make up the application. The MPCUI framework interacts with the integration class to understand the structure of the application, allowing the framework to be the primary driver behind the display of and navigation between the UI elements that make up the application.

The application registers the integration class with the MPCUI framework in the application class through the `getIntegrationClass` method. Each application should have a single integration class only.

Example 8–7 Registering the Integration Class

```
<?xml version="1.0" encoding="utf-8"?>
<intg:Integration
  xmlns:mx="http://www.adobe.com/2006/mxml"
  xmlns:intg="oracle.sysman.emx.intg.*"
  >
  <!-- The integration class defines the pages, dialogs
       and trains included in the application -->

  <intg:activities>
    <intg:PageActivityDef id='homePg' label='Home' pageClass='{HomePage}'
pageControllerClass='{HomePageController}' isDefaultPage="true" />
    <intg:PageActivityDef id='processesPg' label='Processes'
pageClass='{ProcessesPage}' pageControllerClass='{ProcessesPageController}' />
    <intg:PageActivityDef id='adminPg' label='Administration'
```

```
pageClass='{CredentialsPage}' pageControllerClass='{CredentialsPageController}' />
-->
<intg:DialogActivityDef id='metricHistory' label='Metric History'
dialogClass='{MetricHistoryDialog}' >
  <intg:inputParams>
    <intg:InputParam name='targetName' />
    <intg:InputParam name='targetType' />
    <intg:InputParam name='metric' />
    <intg:InputParam name='columns' />
    <intg:InputParam name='period' />
    <intg:InputParam name='title' />
  </intg:inputParams>
</intg:DialogActivityDef>

<intg:DialogActivityDef id='availDialog' label='Availability'
dialogClass='{AvailabilityDialog}' />

</intg:activities>
</intg:Integration>
```

In [Example 8-7](#), the availability activities include `PageActivityDef`, `DialogActivity`, and `TrainActivityDef`. Each activity typically specifies an “id” property that will be used throughout the application to refer to this activity.

The `pageClass` and `dialogClass` properties specify the view class or the UI layout for each activity. For example, the `homePg` activity has a `pageClass` of `HomePage`. This means that included in the application should be a class called `HomePage` typically written in MXML.

Activities can specify a controller class (`pageControllerClass`). This property points to a class (often written in `ActionScript`) that will be associated with the activity and called by the MPCUI framework to initialize data in the page and respond to UI events from user interaction within the page.

8.6.2 Defining Pages

Each page must be registered with the MPCUI framework through the `Integration` class by adding a `PageActivityDef`. The `PageActivityDef` is defined by:

- Page class

The page class is the concrete implementation of the page, that is its layout and contents and is a class that must extend the `Page` class.

- Page controller

The page controller is a class that extends the `ActivityController` base class and encapsulates the set of handlers that support interacting with the Enterprise Manager services layer to obtain data and bind it to the UI components and respond to events issued by the UI on behalf of the end-user (e.g.button presses or link clicks)

Each application must include at least one page (one page activity) and you must identify one of the page activities as the default page.

Note: The default page is displayed by the MPCUI framework as the home page for the selected target

8.6.2.1 Page Class

The `Page` class is the base class for all pages defined by the end-user. `page` is the top-level UI element in the application. The framework provides integration of pages into the Enterprise Manager console by:

- integrating pages with the Enterprise Manager menu system
- performing updates of the browser history so that pages can be bookmarked
- providing simple navigation between pages

Implement page classes in MXML and extend the `Page` base class to integrate with the MPCUI framework.

The tag language that is used to describe the page includes a mix of Flex components and MPCUI-provided components for layout and data display. The description of each component and example for its use are included in subsequent sections of this document.

For examples of the page class, see the `HomePage.mxml` and `ProcessesPage.mxml` files from the Demo HostSample in the EDK.

8.6.2.2 Page Model

Components within the page display information obtained through the Enterprise Manager services layer, and typically are bound to this data through the page model. The page model is the set of data associated with the page. The framework manages the lifecycle of this data so that as pages are displayed, data is loaded. When pages are removed, the data is cleaned up and can be garbage collected by the Adobe Flash Player plug-in.

Specify the data included in the page model by:

- using data service tags
- adding data directly to the page model in the result handlers for Enterprise Manager service requests

For additional information about describing the use of the service layer and how data is added to the model, see [Section 8.12, "Performing Task Automation"](#).

Although Adobe Flex and ActionScript support the ability to inline code in an MXML file using the `Script` tag, Oracle recommends that the `Page` code is limited to the layout of the UI elements that make up the page. Delegate data binding and event handling to the controller. This ensures that the MPCUI framework can manage the lifecycle of each page and the data bound to it correctly.

8.6.2.3 Page Controller

The page controller is a class that extends the `PageController` base class and includes the code that interacts with the Enterprise Manager services layer to obtain data and to process administrative actions. Furthermore, the controller contains the set of event handlers that are called in response to events issued from the `Page` components.

Note: A page controller is not necessary if all of the data displayed in the page can be specified through the component tags or the `DataService` tags and custom event handling is not necessary.

For example, if a page is a container for a number of `Chart` components, then each component supports the specification of the metric to be displayed in the chart. The component interacts with the MPCUI framework to manage the life cycle of that data correctly.

For cases where a controller is necessary, the `init (page:Page)` method is the location in the code where you can load data to be bound to the page UI elements. For examples for interacting with Enterprise Manager services and binding using the page model, see [Section 8.12, "Performing Task Automation"](#).

In addition to the `init` method, the controller includes methods that respond to events originating in the page. In cases where it is necessary to perform some processing in response to an event (for example, a button press), you can reference a method in the controller that will be called when that event occurs.

- Within the Page:

```
<components:Link label="Show History"
  click="controller.showHistory(buffCacheChart);" />
```

- Within the Controller:

```
public function showProcessorHistory(event:MouseEvent):void
{
    // show an example of invoking an activity (a dialog in this case) and
    // getting information from the dialog when it returns (is closed)
    // create the context to be passed to the dialog
    var bean:Bean = new Bean('targetName',
        ApplicationContext.getTargetName(), 'targetType',
        ApplicationContext.getTargetType(),
        'metric', 'CPUProcessorPerf', 'columns', ['CPUIdle'],
        'period', 'LAST_DAY', 'title', 'Metric History');
    page.invokeActivity('metricHistory', bean, processorHistoryDone);
}
```

In the page code, a reference to `controller` is all that is necessary to interact with code included in the page controller. The framework manages creating the controller class when the page is loaded and provides the ability to call through into the controller to take some action.

The framework simplifies the process for taking some actions by providing convenience methods that can be called directly from the Page without requiring additional event handlers in the controller. For example, accessing another activity can be done in most cases without requiring additional controller code.

In the following example, clicking the link redirects the application to the `processPg` activity.

```
<components:Link label="Show Process" click="{invokeActivity('processPg')}" />
```

Note: For more information, see the `HomeController.as` and `ProcessesPageController.as` files from the Demo Sample.

8.6.3 Defining Dialogs

The Dialog activity extends the MPCUI Dialog class. Dialogs are popup windows that display on top of the application without navigating away from the current Page displayed. Dialogs are typically defined in MXML files and do not have separate controller classes (although they can).

```
<?xml version="1.0" encoding="utf-8"?>
<intg:Dialog
  xmlns:mx="http://www.adobe.com/2006/mxml"
  xmlns:cht="oracle.sysman.emx.components.charts.*"
  xmlns:intg="oracle.sysman.emx.intg.*"
  xmlns:ds="oracle.sysman.emx.service.util.*"
  xmlns:comp="oracle.sysman.emx.components.*"
  xmlns:tbl="oracle.sysman.emx.components.table.*"
  height="250" width="450"
  title="{model.title}"
>
  <cht:LineChart id="hchart" targetName="{model.targetName}"
targetType="{model.targetType}" timePeriod="{model.period}" interval="15"
metricName="{model.metric}"
metricColumns="{model.columns}" keys="{model.keys}" width="100%" height="100%" />
</intg:Dialog>
```

In the previous example, the dialog references `model` as the source of the properties it uses in the UI components.

Initialize the dialog model either:

- In a controller associated with the dialog
- By the MPCUI framework if the Dialog definition in the Integration class specifies input parameters

```
<intg:DialogActivityDef id='metricHistory' label='Metric History'
dialogClass='{MetricHistoryDialog}' >
  <intg:inputParams>
    <intg:InputParam name='targetName' />
    <intg:InputParam name='targetType' />
    <intg:InputParam name='metric' />
    <intg:InputParam name='columns' />
    <intg:InputParam name='period' />
    <intg:InputParam name='title' />
  </intg:inputParams>
</intg:DialogActivityDef>
```

Note: In this case, you must supply a bean as input that includes the input parameters required by the dialog.

```
<components:Link label="Show History"
click="{invokeActivity('metricHistory',
bean('targetName', appModel.target.name,
'targetType', appModel.target.type,
'metric', 'Response',
'columns', ['Load'],
'period', '',
'title', 'Metric History'))}" />
```

Note: For more examples, see the `MetricDetailsDialog.mxml` and the `AvailabilityDialog.mxml` files from the Demo Sample.

8.6.4 Defining Trains and Train Pages

The train activity enables you to define a train (a guided workflow or wizard) by stringing together a series of pages.

To define a train, include a declaration of the train itself (`TrainActivityDef`) and each of the steps (`TrainStepActivityDef`) in the Integration class:

```
<intg:TrainActivityDef id='addNewUserEmbeddedTrain' label='Add New User'>
  <intg:stepActivities>
    <mx:Array>
      <intg:TrainStepActivityDef id='anuStep1' label='User
Info'pageClass='{trainSamp.S1_UserInfo}'
pageControllerClass='{trainSamp.AddNewUserTrainStepController}'/>
      <intg:TrainStepActivityDef id='anuStep2' label='Expiry'
pageClass='{trainSamp.S2_
Expiry}'pageControllerClass='{trainSamp.AddNewUserTrainStepController}'/>
      <intg:TrainStepActivityDef id='anuStep3' label='Credentials'
pageClass='{trainSamp.S3_Credentials}'
pageControllerClass='{trainSamp.AddNewUserTrainStepController}'/>
      <intg:TrainStepActivityDef id='anuStep4' label='Schedule'
pageClass='{trainSamp.S4_Schedule}'
pageControllerClass='{trainSamp.AddNewUserTrainStepController}'/>
      <intg:TrainStepActivityDef id='anuStep5' label='Notifications'
pageClass='{trainSamp.S5_Notifications}'
pageControllerClass='{trainSamp.NotificationsTrainStepController}'/>
      <intg:TrainStepActivityDef id='anuStep6' label='Confirmation'
pageClass='{trainSamp.S6_Confirm}'
pageControllerClass='{trainSamp.AddNewUserTrainStepController}'/>
    </mx:Array>
  </intg:stepActivities>
</intg:TrainActivityDef>
```

The `TrainController` includes the following methods:

- `init(Train)`: a method that is called when the train is loaded, and enables you to control the model associated with the train.
- `trainDone`: a method that is called when the user clicks the **Finish** or **Cancel** button within the train. At that point, you can inspect the train state (whatever is stored in the train model) to do one of the following
 - Control if the train should complete and continue to the completion activity
 - Take some other action such as moving the train back to a previous step by using the `train.setStep` method or end the train and invoke another activity.

Each train step within the train must extend the `TrainStepPage` (a special type of `Page`) and be associated with a controller (`TrainStepController`). In this case, the controller is a special type of `PageController`, and includes support for the `init(Page)` method that enables you to initialize the contents of the train page. Because the page is within a train, it might refer to either its own page model (such as `model.property`) or it might refer to data stored in the train model (such as `train.model.property`).

Finally, in either the train step controller or the train controller, the code can check for state and if the train can complete, that is, all the required information is entered, then the controller code can call `train.setMayFinish()`.

Note: For more information, see the `trainSamp` examples from the Demo Sample.

8.6.5 Defining URLs

`UrlActivityDef` support the ability to define a URL that can be accessed using the `invokeActivity` directive from a UI component click handler (for example, `InfoItem`, `ImageLink`, and `Button`). The URL can be represented as an absolute URL including all request parameters, or parameters can be supplied at runtime. To define the URL that should have URL parameters substituted at runtime, define the `UrlActivityDef` to include `inputParams` as follows:

```
<mp:UrlActivityDef id='oracle' label='myExtApp' urlBase="http://www.extapp.com" >
  <mp:inputParams>
    <mp:InputParam name='pageId' />
  </mp:inputParams>
</mp:UrlActivityDef>
```

To reference the URL the `invokeActivity` directive used specifying the id of the `UrlActivityDef` and passing a bean that includes the parameter and the appropriate value. The parameters provided will be added to the URL as request parameters.

```
<mp:InfoItem id="currentLoad" label="CPU Load"
  value="{respData.result.getString('','Load')}}"
  click="{invokeActivity('extapp', bean('pageId','Load'))}"
/>
```

In this example, the URL that is accessed is `http://www.extapp.com&pageId=Load`.

8.7 Packaging the MPCUI Implementation With the Plug-in

Include the MPCUI implementation in a plug-in by placing a metadata definition of the MPCUI in the `/mpcui` subdirectory of the plug-in stage directory. For information about the structure and packaging of plug-ins, see [Chapter 13](#).

Metadata-only Implementation

Put the MPCUI metadata file in the following directory:

```
plugin_stage/oms/metadata/mpcui/my_mpcui_metadata.xml
```

Flex Implementation

Include a single metadata file and the SWF file (the Flex application that is the custom UI) in the `/mpcui` subdirectory:

```
plugin_stage/oms/metadata/mpcui/my_mpcui_metadata.xml
```

```
plugin_stage/oms/metadata/mpcui/MyMpcui.swf
```

Note: In the previous examples, set the names of the XML (*my_mpcui_metadata.xml*) and SWF (*MyMpcui.swf*) files according to your requirements as a plug-in developer.

8.8 Converting a Metadata-based UI to a Flex-based UI

The tag language used to define pages for the metadata-based UI is a subset of the tag language supported for the Flex-based implementation. Therefore, if you have a metadata-based UI and decide that you want additional features supported in the Flex implementation only, use your metadata-based UI to create the Flex-based UI. You do not have to start again because you can reuse the metadata.

Take a page defined in a metadata file for the metadata-based implementation and turn it into a page that is part of an MPCUI application (SWF file) by following these steps:

1. From the metadata file, identify the pages that you want to convert.
2. Copy the `ActivityDefinition` block for a page and save the block as an MXML file.
3. Repeat step 2 for each page that you want to convert. For example, if you have four page definitions, then you must create four MXML files.
4. Copy the `Integration` block from the metadata file and save the `Integration` block as an MXML file.
5. At the top of each MXML file, add the appropriate MXML namespace.
6. Create an application MXML file as described in [Section 8.6, "Defining the MPCUI Application"](#).
7. Create a Flex-based metadata file as described in [Section 8.4, "Creating the MPCUI Metadata File"](#)

8.9 Defining System Home Pages

For target types identified as system targets, there are three options for which home page is rendered for the system target.

1. Display the Enterprise Manager default system home page.

This page shows a summary of the availability and incidents for the system members. This option is enabled by either of the following:

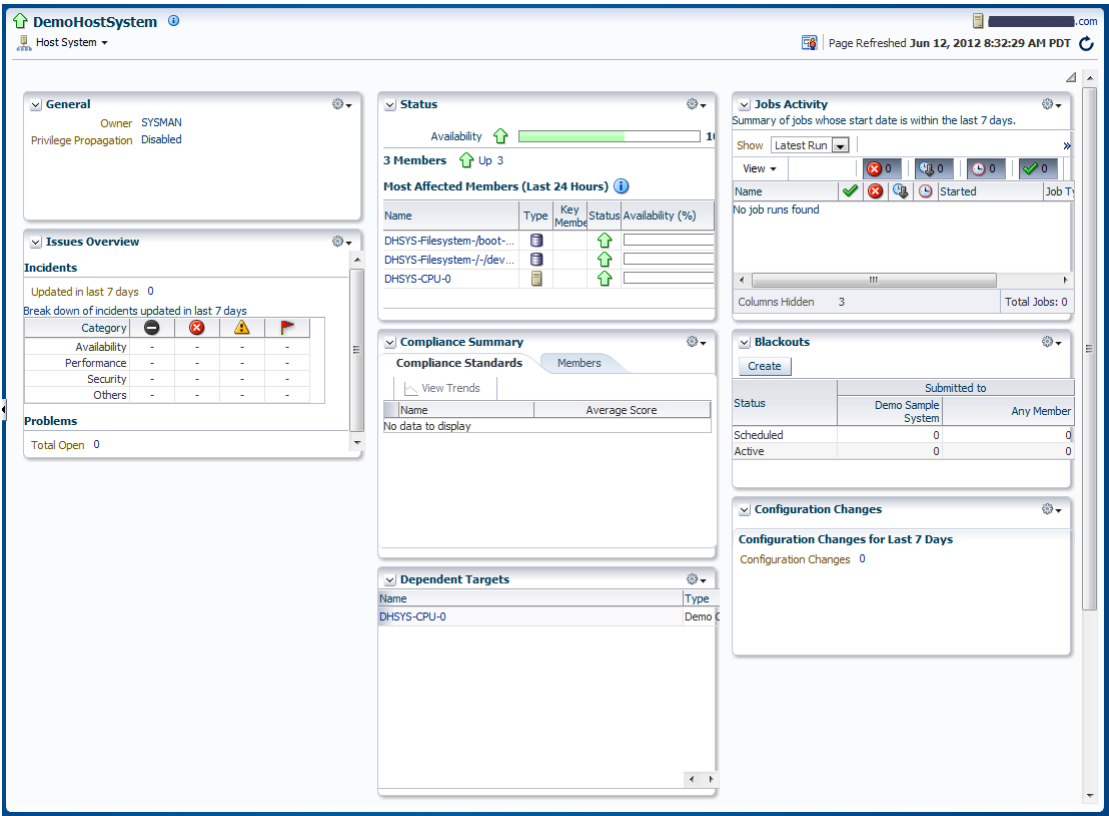
- Omitting MPCUI metadata from your plug-in
- Including MPCUI metadata in the plug-in and including the following `<EmuiConfig>` element in the MPCUI metadata file:

Example 8–8 Using the Default System Home Page

```
<CustomUI target_type="demo_
hostsystem" xmlns="http://www.oracle.com/EnterpriseGridControl/MpCui">

    <EmuiConfig>
        <use-framework-homepage>true</use-framework-homepage>
    </EmuiConfig>
</CustomUI>
```

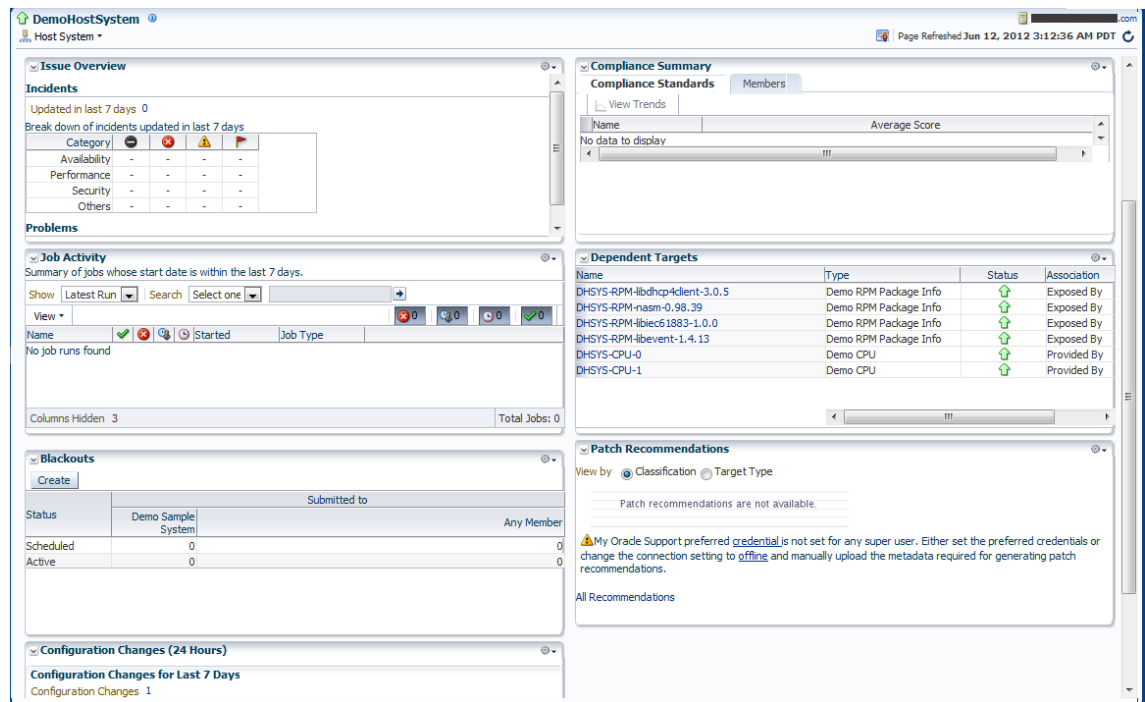
Figure 8–1 Default System Home Page



2. Display the Enterprise Manager default system home page, with some customized content.

The home page can show a number of prepackaged regions in a customized layout. The use of the default home page is controlled by metadata as illustrated in [Example 8–8](#).

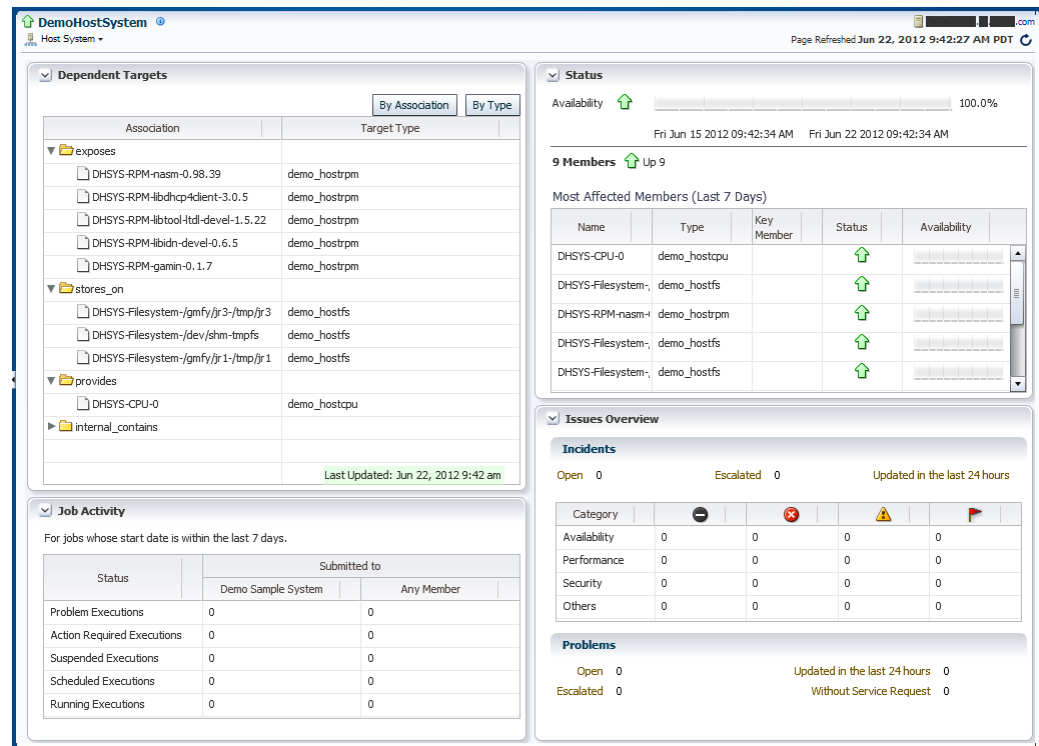
The selection of regions and their layout on the home page is specified by including `systemUiIntegration` metadata in the plug-in. For more information, see [Section 8.9.1, "Defining systemUiIntegration Metadata"](#)

Figure 8–2 System Home Page With Some Customization

- Construct a custom home page using the MPCUI capabilities included with the EDK.

The home page is constructed using either MPCUI metadata or using the MPCUI Flex libraries. There are several data services and UI components that are provided by MPCUI specific to system or composite target types. For more information, see [Section 8.9.2, "Defining System Regions"](#)

Figure 8–3 Customized System Home Page



8.9.1 Defining systemUiIntegration Metadata

To use the default system home page with some customized content:

1. Define a systemUiIntegration Metadata XML file for your target type including the following information:

- Preferred layout
- Add or remove regions (only required if you want to modify regions)

[Example 8–9](#) provides an example of a systemUiIntegration Metadata XML file.

For information about the XML Schema Definition (XSD) that governs the systemUiIntegration Metadata XML file, see *ORACLE_HOME/sysman/emSDK/core/system/xml/SystemUiIntegration.xsd*.

Example 8–9 systemUiIntegration Metadata XML

```
<systemUiIntegration xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://www.oracle.com/EnterpriseGridControl/SystemUiIntegration.xsd"
  xmlns="http://www.oracle.com/EnterpriseGridControl/SystemUiIntegration">
  <general targetType="demo_hostsystem"
    defaultLayout="twoColumnNarrowLeft"
    showOptionalRegions="false"
    topLevelTarget="true"
    allowCreateFromSystemsUi="true"/>
```

```
<region
taskFlowId="/WEB-INF/db/system/region/db-system-region-hihgavail-task-flow.xml#db-
system-region-hihgavail-task-flow"
    titleResBundle="oracle.sysman.db.rsc.inst.DBMsg"
        titleNlsId="GENERAL"
        titleDefText="General"
        regionType="add"
        displayOrder="1" />

<region
taskFlowId="/WEB-INF/sdk/core/regions/events/console/incident-overview-task-flow.x
ml#incident-overview-task-flow"
    titleResBundle="oracle.sysman.core.groups.ui.CoreGroupsUiMsg"
        titleNlsId="ISSUE_OVERVIEW"
        titleDefText="Issue Overview"
        regionType="add"
        displayOrder="4" />

<region
taskFlowId="/WEB-INF/sdk/core/regions/jobs/jobs-activity-task-flow.xml#jobs-activi
ty-task-flow"
    titleResBundle="oracle.sysman.db.rsc.inst.DBMsg"
        titleNlsId="JOB_ACTIVITY"
        titleDefText="Job Activity"
        regionType="add"
        displayOrder="7" />

<region
taskFlowId="/WEB-INF/db/system/region/db-system-region-dep-members-task-flow.xml#d
b-system-region-dep-members-task-flow"
    titleResBundle="oracle.sysman.core.groups.ui.CoreGroupsUiMsg"
        titleNlsId="DEPENDENT_TARGETS"
        titleDefText="Dependent Targets"
        regionType="add"
        displayOrder="9" />

<region
taskFlowId="/WEB-INF/sdk/core/regions/gccompliance/target/compliance-overview-task
-flow-brief.xml#compliance-overview-task-flow-brief"
    titleResBundle="oracle.sysman.core.groups.ui.CoreGroupsUiMsg"
        titleNlsId="COMPLIANCE_SUMMARY"
        titleDefText="Compliance Standard Summary"
        regionType="add"
        displayOrder="6" />

<region
taskFlowId="/WEB-INF/sdk/core/regions/mos/patch/target-patch-recommendation-task-f
low.xml#target-patch-recommendation-task-flow"
    titleResBundle="oracle.sysman.db.rsc.inst.DBMsg"
        titleNlsId="PATCH_RECOMMEND"
        titleDefText="Patch Recommendations"
        regionType="add"
        displayOrder="12" />

<region
taskFlowId="/WEB-INF/config/adfc/blackout/region/emcore-groups-blackout-task-flow.
xml#blackout_group_taskflow"
    titleResBundle="oracle.sysman.core.groups.ui.CoreGroupsUiMsg"
        titleNlsId="BLACKOUTS"
        titleDefText="Blackouts"
```



```

        regionType="add"
        displayOrder="2" />

<region
taskFlowId="/WEB-INF/sdk/core/regions/ecm/history/config-history-task-flow.xml#con
fig-history-task-flow"
        titleResBundle="oracle.sysman.db.rsc.inst.DBMsg"
        titleNlsId="CONFIG_CHANGES"
        titleDefText="Configuration Changes (24 Hours)"
        regionType="add"
        displayOrder="5" />

</systemUiIntegration>

```

2. Save the systemUiIntegration Metadata XML file to the following directory:

```
plugin_stage/stage/oms/metadata/systemUiIntegration
```

3. If your plug-in is deployed already, then you can use the `emctl register oms metadata` command to update the MPCUI part of your plug-in only. For more information about the `emctl register oms metadata` command, see [Section 13.7](#).

8.9.2 Defining System Regions

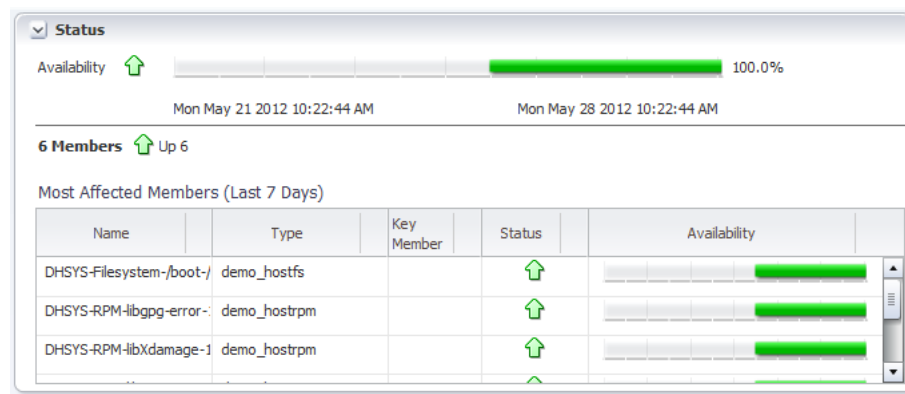
The MPCUI framework supports a number of regions that can be used as part of a home page built to display information for a system target.

8.9.2.1 Defining System Status Region

The system status region shows the recent availability of the system target and all of its members. The region is included in the system home page by using the following tag:

```
<mp:StatusOverviewRegion id="statusOverview" height="50%" />
```

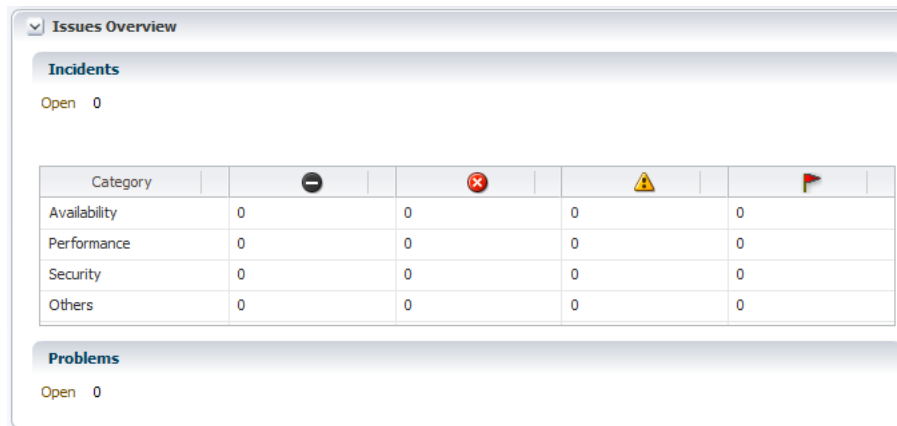
Figure 8–4 System Status Region



8.9.2.2 Defining System Issues Region

The system issues region shows the summary count of incidents for all of the targets in the system. The region is included in the system home page by using the following tag:

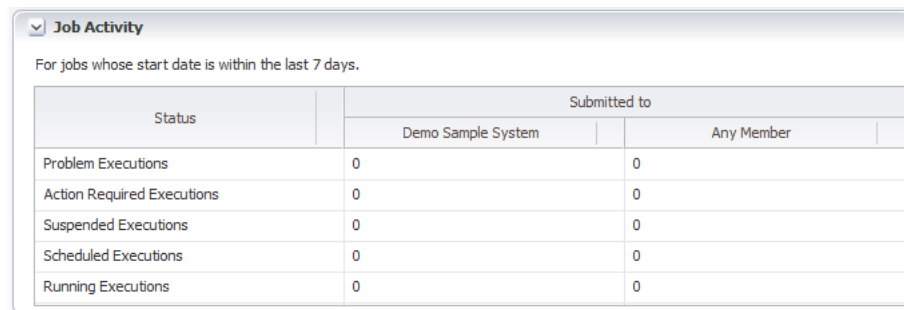
```
<mp:IssuesOverviewRegion id="issuesOverview" height="50%" />
```

Figure 8–5 Issues Overview Region

8.9.2.3 Defining the System Job Activity Region

The system job activity region displays the number of jobs in each of the primary job status for the system target and the summary for all the system members.

```
<mp:JobsActivityRegion id="jobsOverview" height="40%"/>
```

Figure 8–6 System Job Activity Region

8.10 Defining Navigation

Navigation in the MPCUI application can be either of the following:

- Between activities defined in the application. For more information, see [Section 8.10.1, "Navigation to Activities"](#).
- To other URLs, where URL refers to other Enterprise Manager pages or to external URLs. For example:

```
http://www.example.com
```

For more information, see [Section 8.10.2, "URL and Links"](#).

8.10.1 Navigation to Activities

[Section 8.5.3, "Defining Navigation"](#) describes the approach to navigating between activities from a metadata implementation. These descriptions apply to navigating to activities from the menu or from another activity defined in MXML.

This section describes how to navigate to another activity from within the controller code, that is the `ActionScript` code associated with an activity.

```
public function showProcessorHistory(event:MouseEvent):void
{
    // show an example of invoking an activity (a dialog in this case) and
    // getting information from the dialog when it returns (is closed)

    // create the context to be passed to the dialog
    var bean:Bean = new Bean('targetName',
        ApplicationContext.getTargetName(), 'targetType',
        ApplicationContext.getTargetType(),
        'metric', 'CPUProcessorPerf', 'columns', ['CPUIdle'],
        'period', 'LAST_DAY', 'title', 'Metric History');

    page.invokeActivity('metricHistory', bean, processorHistoryDone);
}
```

The preceding example shows a controller method that uses the `page.invokeActivity` method to redirect to another activity (in this case, a dialog).

The significant difference between this method and the method available from within the MXML page (described in [Section 8.5.3, "Defining Navigation"](#)) is the ability to associate a callback (`processorHistoryDone` in this example) that will be called when the called activity completes. This callback is only useful for activities that do not cause the current activity to go out of scope.

8.10.2 URL and Links

There are a number of different methods for navigating from components in the MPCUI application to other locations through a URL. Use the `Link` component to render an HTML-style link including a tool tip and location.

- **Absolute URL (external to Enterprise Manager)**

To provide a link to an absolute URL, use the `UrlAbs` class and an instance of this class can then be associated with a `Link` destination or can be accessed through the `invokeActivity` method.

In the Page Class:

```
<comp:Link id="gotoOracle" label="Oracle" destination="{model.oracleUrl}"
/>
```

In the Controller Class:

```
page.setModel("oracleUrl", new UrlAbs("http://www.oracle.com", "Oracle"));
```

Alternative method using `invokeActivity`:

In the Page Class:

```
<mx:Button label="Go To Oracle" click="{invokeActivity(model.oracleUrl)}"
/>
```

In the Controller Class:

```
page.setModel("oracleUrl", new UrlAbs("http://www.oracle.com", "Oracle"));
```

- **Link to Enterprise Manager Page Using Page Constants**

In addition to absolute URLs, the MPCUI framework supports the ability to link to well known Enterprise Manager pages by constructing a “UrlEm” object that can be referenced from the Link destination or passed to the `invokeActivity` method as part of a click handler. The reference guide includes a complete list in the `oracle.sysman.emx.Constants` class of all page constants available and the corresponding parameters that must be specified to produce a URL.

```
// setup link to availability page
var availLink:UrlEm = new UrlEm(Constants.PAGE_AVAILABILITY,
    [new InputParam(Constants.P_TARGET_NAME,
        ApplicationContext.getTargetName()),
      new InputParam(Constants.P_TARGET_TYPE,
        ApplicationContext.getTargetType()),
      new InputParam(Constants.P_PAGE_TYPE,
        Constants.BY_DAY)]);
page.setModel("availPageLink", availLink);
```

- **Link to Enterprise Manager Pages That Do Not Have Constants Defined**

Note that `UrlEm` can only be used to access pages that are supported via page constants in the `oracle.sysman.emx.util.Constants` class. For pages that do not currently have constants defined, you can access a page by creating a `UrlRel` object containing the page's ADF view ID value.

For example, to access the Bare Metal Provisioning dashboard, you would specify the page's view ID (`/faces/core-bmp-dashboard`) as follows:

```
var url:UrlRel = new UrlRel("/faces/core-bmp-dashboard", null);
```

The easiest way to find the view ID for a given ADF page is in the page URL; it is the string following `http://<server:host>/em/`.

- **Link to Enterprise Manager Target Home page**

A special case is to produce the URL to an Enterprise Manager target home page. For this situation, use the static `UrlEm.homepageUrl` method:

```
page.setModel("relatedHostLink", UrlEm.homepageUrl(host.name, host.type));
```

- **Dynamic URL Using “DIRECT_URL”**

For cases where a URL must be constructed dynamically at runtime from a data service, the following option may be used. The activity id “DIRECT_URL” is reserved for the special case and is provided by the framework. No `UrlActivityDef` is declared in this case, but instead the `invokeActivity` directive is passed a bean that specifies the “url” property. The value provided for that property will be used as the URL to direct to when the component is clicked.

In the following example, the data service “respData” is queried to obtain a URL. This would be replaced by whatever data service is used within the page to obtain the necessary URL. This may be a `MetricValuesDataService` or a `SqlDataService`.

```
<mp:InfoItem id="currentLoad" label="CPU Load"
  value="{respData.result.getString('','Load'))}"
  click="{invokeActivity('DIRECT_URL',
    bean('url',respData.result.getString('','Load')))}"
/>
```

8.10.3 Adding Links to External Applications

Providing the ability to link to other applications outside of Enterprise Manager is not currently supported.

8.11 Accessing Enterprise Manager Data

The MPCUI framework provides access to Enterprise Manager services through ActionScript interfaces to the Enterprise Manager Web services layer. You can access these client services directly when necessary. Although in many cases, the services are further abstracted through UI components that utilize them to interact with the Enterprise Manager server to obtain the appropriate data to be displayed in the management UI.

The following sections describe the various services included in the MPCUI framework and provide brief examples of how these services can be used from your code.

Note: The EDK does not support accessing arbitrary Web services from the Flex UI. The appropriate way to access Web services would through the Management Agent residing on the service host, as either metrics, jobs or remote commands invoked by a fetchlet.

8.11.1 Metric Services

The MPCUI provides a simple service for retrieving metric data from the Management server in either real-time or historical form. For real-time data, the Oracle Management Service accesses the Management Agent to retrieve the data, so use this for cases where the metric can be collected efficiently in real time.

8.11.1.1 Using the Metric Values Service Transparently

Usually the metric values service is used transparently from a chart by specifying the metric to be displayed in the chart and in the case of a line chart, the periodicity of the data.

```
<mp:LineChart id="cacheChart"
    width="100%" height="100%"
    metricName="MSSQL_MemoryStatistics"
    metricColumns="['cache_hit_ratio']"
    timePeriod="REALTIME" interval="15" >
</mp:LineChart>
```

In this case, the caller never interacts directly with the service. The MPCUI framework uses the service to retrieve the data for the chart.

In the case of the table component, you can specify the metric directly also:

```
<c:Table id="processesTable" width="100%" height="100%"
    metricName="CPUProcessesPerf"
    metricColumns="['ProcUser', 'ProcCPU', 'ProcCmd']"
    timePeriod="REALTIME"
    interval="30"
    >
    <c:columns>
        <mx:AdvancedDataGridColumn width="50" dataField="key" />
        <mx:AdvancedDataGridColumn width="100" dataField="ProcUser" />
        <mx:AdvancedDataGridColumn width="80" dataField="ProcCPU" />
        <mx:AdvancedDataGridColumn width="400" dataField="ProcCmd" />
    </c:columns>
</c:Table>
```

8.11.1.2 Using the MetricValuesDataService Tag

Use the `MetricValuesDataService` tag within a page (or dialog) to display metric data in a table component, where the `dataProvider` attribute of the table is set to the data service. Then the data from the metric service is displayed in the table or when data from the service will be shared between multiple components (for example, the table and a link or label).

Example 8–10 Using the MetricValueDataService Tag

```
<intg:services>
  <dataserv:MetricValuesDataService id="mv1" flattenData="true"
    targetName="{ApplicationContext.getTargetName()}"
    targetType="{ApplicationContext.getTargetType()}"
    metricName="Load" columns="{['cpuUtil', 'cpuUser',
'cpuKernel']}"
    timePeriod="{MetricCollectionTimePeriod.LAST_DAY}"
  />
</intg:services>
<comp:Table id="mvTable" dataProvider="{mv1}" />
```

8.11.1.3 Calling the Metric Value Service From a Controller

The metric value service can be called from within a controller. This is the most flexible means of using the service and allows the caller to manipulate the data as necessary before adding the final results to the model so that it can be displayed in the UI.

Retrieving Individual Values from the Metric Service (MXML)

You can retrieve individual values from the metric service in order to display them in a Label, InfoItem, or other such component.

```
<ds:MetricValuesDataService id="procData"
  flattenData="true"
  targetName="{appModel.target.name}"
  targetType="{appModel.target.type}"
  metricName="CPUProcessorPerf"
  columns="{['CPUIdle']}"
  timePeriod="REALTIME"
  interval="15" />
```

Then from the component that will display the value:

```
<components:InfoItem label="CPU(0) Idle %"
  value="{procData.result.getString('0','CPUIdle')}" />
```

Example 8–11 The Metric Service from a Controller

```
var cpuPerf:Metric =
    ApplicationContext.getTargetContext().getMetric("CPUPerf");
var cpuPerfSel:MetricSelector = procMetric.getSelector(
    ['system','idle','io_wait']);
cpuPerfSel.getData(cpuDataHandler, MetricCollectionTimePeriod.CURRENT,
    page.getBatchRequest());
```

Use the metric service by creating a `MetricSelector` for a particular metric, and then calling the `getData` method on that selector. When calling the `getData` method, two parameters are passed:

- the handler that will be called with the result of the request

- the periodicity of the selection

When the service request has completed, either successfully or with an error, the handler is called and passed the results of the request and a fault. The caller must check for the presence of the fault before proceeding with any processing of the data result.

Example 8–12 Metric Service Result Handler

```
public function
cpuDataHandler(cpuData:MetricResultSet,fault:ServiceFault):void
{
    if(fault != null) return;    // handle this better!

    var dataPoint:TimestampMetricData = cpuData.results[0];
    var collectionTime:Date = dataPoint.timestamp;
    var idleTime:Number = dataPoint.data[0]['idle'];
    var systemTime:Number = dataPoint.data[0]['system'];
    var ioWaitTime:Number = dataPoint.data[0]['io_wait'];
}
```

To access the data, you must have the reference to `dataService.result.getString('key', 'column')`. The key is required to identify the row in the sample to be returned in cases where the metric supports multiple keys. If the metric does not include a key column, then the key value should be passed as "" or null. The column is the data column to be retrieved from the metric definitions.

Each data point (`TimestampMetricData`) has a time stamp member that tells you when that data point was collected, and includes a data array that is effectively a table for that metric.

If the metric has multiple keys (such as process, file systems, and so on), then the data array has multiple rows, one for each key, and each row has the requested data columns. In the previous examples, the data array contains one row for each process. If your metric does not include key columns, then the data array contains a single row only.

Each row in the data array is a `KeyMetricData` object. If your metric has keys, then the `metricKey` property tells you to which key the row applies. If you have no key for your metric, then ignore this property. The `KeyMetricData` is a dynamic object into which you can index, using the column name to get the value for that column.

In the previous examples, the code walks the rows in the data array, and for each row (`KeyMetricData`) it gets the `'ProcUser'` column from the data. The original request also included the `'ProcCPU'` and `'ProcCmd'` columns, so those could be accessed in the same way, that is, `data['ProcCPU']` or `data['ProcCmd']`.

8.11.1.4 Metric Data Source Filters

When using the metric data service through the `MetricValuesDataService` in MXML or the `MetricSelector` in ActionScript, it may be useful to request that the set of data returned by the service be filtered according to some additional selection criteria. This can be accomplished within the controller by implementing a custom data source and then filtering the results of the metric service in the controller and populating the custom data source with the results.

It is also possible to define a metric filter that can be applied to the request which will cause the service itself to filter the results and return only the filtered set to the client for display.

The metric filter, referred to as `MetricPredicate`, is made up of several elements, including individual column filters, the filter operator, and the optional order by criteria. Each column filter specifies a column to filter, the operator to filter by, and a value to filter against. The column filters support typical operators for numeric data, including `GT`, `LT`, `GE`, and `LE`.

For string data, the operators include `EQ`, `NE`, and `REGEX`. The `REGEX` operator will perform a regular expression string match using each value with the filter input value as the pattern. The regular expression pattern match is done using Java regex libraries, so the pattern should conform to the requirements of Java pattern matching.

The predicate operator combines the column filters into a single expression and supports either an `AND` (all column filters must be satisfied) or an `OR` (any of the column filters being satisfied is sufficient). The order by criteria specifies a column to order the results by and a row count to limit to. This is useful in cases where a "top-N" result is desired.

When constructing a metric filter, the columns filters can be optional and an order by only specified. Alternately, the order by can be optional and the column filters only are specified. When constructing a metric filter, all columns included in column filters and in the order by must be part of the same metric, and it must be the metric that is being selected in the corresponding metric data service request. Combining columns from multiple metrics into the same filter is not supported.

The following example describes the process for defining a metric filter on a `MetricValuesDataService` tag. The data service tag includes the "predicate" property which is bound to the corresponding metric filter (`MetricPredicate`) as such:

```
<mp:MetricValuesDataService
  id="fsMetDs"
  metricName="FilesystemPerf"
  columns="['MountPoint','Utilization','FreeKB','UsedKB','TotalKB','FSType',
'FSName']"
  targetName="{appModel.targetName}"
  targetType="{appModel.targetType}"
  timePeriod="LAST_HOUR"
  predicate="{model.fsFilter}"
/>
```

In the controller associated with the page, the filter is constructed by specifying the filter columns, operator, and order by criteria. In the following example, the file systems metric request from the service above is filtered to those filesystems with a `TotalKB` size of greater than 1000kb and a regular expression match on the filesystem name (`FSName`) of `'.net'`. Finally, the results are ordered by the `FreeKB` column descending limited to the first five filesystems.

```
private function createFsFilter():MetricPredicate
{
    var filters:Array = new Array(
        new MetricFilter('TotalKB', MetricOperator.GT, 1000),
        new MetricFilter('FSName', MetricOperator.REGEX, '(.*)net(.*)') );
    var orderBy:MetricOrderBy = new MetricOrderBy('FreeKB',
        MetricOrderBy.DESC, 5);
    var predicate:MetricPredicate = new MetricPredicate(filters,
        MetricOperator.AND, orderBy);

    return predicate;
}
```


8.11.2 Custom Data Source

In addition to the metric and SQL data sources (and service tags) that can be used to obtain data for charts, tables and other components, you can construct your own custom data source for these components. This is useful in situations where you want to obtain data from other MPCUI services and manipulate it before display. For example, to combine data from two metrics, filter the data in some way, or otherwise aggregate the data.

Creating a custom data source requires the use of controller code to obtain the source data and then to manipulate it to create the data source. The custom data source provides the following important behavior:

- Set column descriptors for the data included in the data source to provide help to the UI component when displaying the data. The descriptor contains properties such as data type, and display label (for legends or column headers).
- Support multiple data points to enable the display of the data in a time-series chart.
- Support caching and modification of the data source allowing components to show updated data as information underlying the data source changes.

8.11.2.1 Creating the Custom Data Source

Typically the custom data source

(`oracle.sysman.emx.model.CustomDataSource`) is constructed and set in the page model using `Page.setModel`. When constructing the data source, you must specify the columns (or data items) that make up the data source along with a flag that can indicate the following:

- If the data should be treated as if it includes a key
Specify the key only if the data source will be displayed in a chart that honors keys such as a bar or column chart. If the data will be shown in a tabular view or a non-chart component, then you do not have to identify one of the columns as a key.
- If the data should be treated as if it includes multiple timestamp samples
Specify that the data includes timestamps only if the data will be displayed in a time-series chart (`LineChart`) and might have data samples added to the data source over time by using the MPCUI polling mechanism.

```
public function CustomDataSource(columns:Array, hasKey:Boolean=false,
isTimeSeries:Boolean=false)
```

The Array of columns specifies the data items included in the data source. This array can be either:

- an array of strings, with each string specifying the label of the data item
- an array of column descriptors (either `QueryColumnDesc` or `CustomColumnDesc`)
Specifying a column descriptor enables you to specify a label for the column and a data type (for `QueryColumnDesc`) or to specify additional properties to display the data in a tabular display such as the column width, that is, if the column is sortable, and so on (for `CustomColumnDesc`).

[Example 8–13](#) shows a result handler in the controller that is set up to handle data returned from a request to the `SqlQueryService`.

Example 8–13 Handling Data Returned From a Request to the *SQLQueryService*

```
// execute a SQL query and then massage the data for display
var query:SqlQueryService = new SqlQueryService('CPU_USAGE',
    [SqlQueryInput.createParam("TARGET_GUID",
        ApplicationContext.getTargetContext().guid)]);
query.execute(cpuQueryHandler, page.getBatchRequest());
}

public function cpuQueryHandler(result:SqlQueryResultSet,
    fault:ServiceFault):void
{
    if(fault != null || result.getError() != null) return;

    cpuSqlData = new CustomDataSource([
        new QueryColumnDesc("Processor", QueryColumnType.STRING),
        new QueryColumnDesc("Idle Percentage",
QueryColumnType.DECIMAL),
        new QueryColumnDesc("Used Percentage",
QueryColumnType.DECIMAL)
    ], true);
    page.setModel("cpuSqlData", cpuSqlData);

    if(result.rows != null)
    {
        for(var r:int=0; r<result.rows.length; r++)
        {
            var id:String = result.getString(r, 'CPU Number');
            var idle:Number = result.getNumber(r, 'Idle %');
            var used:Number = result.getNumber(r, 'Used %');
            cpuSqlData.setRow("Processor #"+id, idle, used);
        }
    }
}
```

In [Example 8–13](#), the data source is constructed with three columns and the data types are specified. The second parameter to the constructor is passed as `true`, indicating that the data should be treated as if it has a key. In this case, the first column in the list is always treated as the key. You cannot specify a different position in the data.

Finally, for each row in the `SqlQueryResultSet` (`result.rows`), the code constructs a row in the custom data source.

See Also: For a complete working example, see the `demo_hostsample, ProcessesPageController.as` in the EDK.

8.11.2.2 Binding the Data Source to a UI Component

In the page layout (for example, `ProcessesPage.mxml`), the data is bound to the UI component using the `dataProvider` property. In [Example 8–14](#), note `cpuSqlTable`. This is a table that displays the data loaded into the `cpuSqlData` custom data source.

Example 8–14 Binding the Data Source to a UI Component

```
<mp:Region id="cpuUtilRegion" width="100%" height="100%" title="CPU Utilization" >
    <mx:HBox width="100%" height="100%">
        <mp:LineChart id="cpuUtilChart" width="60%" height="100%"
            dataProvider="{model.cpuChartData}"
```

```

        legendLocation="right" showLegend="true" />
<mp:Table id="cpuSqlTable" dataProvider="{model.cpuSqlData}"
        width="40%" height="100%" />
</mx:HBox>
</mp:Region>

```

Figure 8–7 shows what Example 8–14 displays.

Figure 8–7 Table Displaying Data Loaded into the `cpuSqlData` Custom Data Source

Processor	Idle Percentage	Used Percentage
Processor #0	99.8	0.2
Processor #1	99.8	0.2

8.11.2.3 Updating the Custom Data Source

Because the data source is bound to the UI component, when you update it, the UI displays the new data automatically. You have two options to update a custom data source:

1. Call either the `CustomDataSource.setRow` or `setRows` methods.

These methods are used when you have a data source that does not include timestamped data. In this case, you are modifying the row or rows included in the data source.

2. If the data source includes timestamped data, then call the `CustomDataSource.setTimestampedRows` method.

This method adds a new sample to the time series and typically is used in the case where the data source is displayed in a line chart. Adding a new sample by calling this method causes a new time slice to appear on the line chart.

For more information about these methods, see the API Reference and the `demo_hostsampl` for examples using the Custom Data Source.

8.11.3 Computed Data Source

The Enterprise Manager metric collection framework supports the ability to compute values from counters. However, the values are only guaranteed to be correct when retrieved from the historical data collected by the agent and stored in the repository. Attempting to query these values from a real-time request (for example, `MetricValuesDataService` with `timePeriod` set to `REALTIME`) can lead to unexpected results as the value computed utilizes the last counter stored during historical collection and not a counter stored for the real-time collection. As such, if you require realtime display of computed metrics you may need to consider using the computed data source.

The computed data source provides the ability to combine data from two metrics into a single display. This is useful when the metric to be displayed is based on a compute expression using a stored counter as described previously.

To use this data source, you typically define two metrics. One metric computes the values to be collected and stored in the repository for historical purposes. This metric includes the compute expressions that consume the stored counters. The other metric would be a transient metric that only collects the counters themselves. This metric

would not be collected for historical purposes as the raw counter values are typically not useful.

You need both sets of metrics when constructing the computed data source in the UI code. The first metric, which specifies the metrics in their computed form, is called the "source" metric. The data source uses these metrics to define the display attributes for the data, including the labels for the columns, and will also retrieve any required historical data.

```
/**
 * Construct a data source that shows the CPU-System% and CPU-Idle% from historical
 * data and then appends data to it from a real-time data source that acquires
 * counter columns and derives the values from the counters. First declare the
 * columns to be shown on the chart, the labels will be based on the metric-column
 * labels and will obtain the historical data that initially populates the chart.
 */
var srcCols:Array = [
    new ComputeColumnDesc( ApplicationContext.getTargetContext(), "CPUPerf",
        "system"),
    new ComputeColumnDesc( ApplicationContext.getTargetContext(), "CPUPerf",
        "idle"),
];

/**
 * These are the counter columns; they do not need to be from the same metric as
 * the source columns, however the counter columns must be from the same metric as
 * all other counters.
 */
var ctrCols:Array = [
    new ComputeColumnDesc( ApplicationContext.getTargetContext(), "CPUPerf",
        "systemCounter"),
    new ComputeColumnDesc( ApplicationContext.getTargetContext(), "CPUPerf",
        "idleCounter"),
];

/**
 * create the data source and pass the source columns, the counter columns and a
 * pointer to the compute function. Finally pass the page the data source will be
 * consumed on and the interval to be used to populate the data. The compute
 * function will be called at each interval.
 */
var computedDataSource:ComputeDataSource = new ComputeDataSource(srcCols, ctrCols,
    computeFunction,
        page, PollingInterval.EVERY_15_SECONDS);
page.setModel("compDataSource", computedDataSource);
```

The computed data source will then send a request for the historical data, and will then begin polling for the counters at the interval specified. Each time a sample is retrieved, the compute function will be called and passed a reference to the computed data source and a data point (TimestampMetricData) that contains the latest set of values for the counter metrics.

The compute function can then compute values using the counters and must return a data point that contains the metrics with the same name as those that were identified in the source columns. In the previous example, the counter columns are "systemCounter" and "idleCounter", but the data point that is returned from the compute function must include a value for the source columns, "system" and "idle".

```
public function computeFunction(ds:ComputeDataSource,
    dp:TimestampMetricData):TimestampMetricData
{
```

```

// retrieve the counter values from the data point passed; could also retrieve
// any necessary context from the data source
var systemCounter:Number = dp.data[0]["systemCounter"];
var idleCounter:Number = dp.data[0]["idleCounter"];

// compute values; this is where you would replicate the logic in your
// computed metric
var systemValue:Number = systemCounter+Math.floor(Math.random()*(50 - 20 + 1))
+ 20;
var idleValue:Number = idleCounter+Math.floor(Math.random()*(120 - 80 + 1)) +
80;

// you must now return a TimestampMetricData object. You can use the one
passed and return
// it, but to do so you must add columns to the data point. The index
reference [0] is
// a reference to the fact that the datapoint could have multiple rows, one for
each key
// but the example does *NOT* support multiple keys. Also, if you created a
new
// data point to return you would need to set the timestamp of the datapoint
// correctly, using the timestamp of the sourced datapoint
dp.data[0]["system"] = systemValue;
dp.data[0]["idle"] = idleValue;

return dp;
}

```

8.11.4 Packaged SQL and the Query Service

While the MPCUI framework provides access to the most useful data through either UI components or simplified services (such as the metric service), inevitably you must have access to other information stored in the Management Repository in a more unstructured form. The MPCUI framework provides a SQL query service for this access.

The SQL query service enables you to package SQL statements with your plug-in and then run the statements through a Web service and then bind that data to UI elements in your custom UI. The SQL query service does *not* provide an open-ended or scriptable API to the Management Repository as this would expose a potential security risk.

The SQL query service can only run SQL statements that have been deployed to the Management repository through the Enterprise Manager Extensibility Framework. This ensures that the statements can access EDK views only. This still provides you with a lot of flexibility and the ability to access data from your own views (for example, views generated from Enterprise Manager configuration data) along with Enterprise Manager partner EDK views.

You can encapsulate the query service entirely within the page code by using the `SQLDataService` tag. This tag allows the caller to specify the SQL to be processed and the parameters to be passed. This data service object can then be bound to a table or to other UI components that support it.

Example 8–15 Using the `SQLDataService` Tag

```

<intg:services>
  <ds:SQLDataService id="dbSummaryDS" queryID="DATABASE_SUMMARY"
    properties="{model.dbSummProp}" />
</intg:services>

```

```

<t:Table id="dbSummaryTable" dataProvider="{dbSummaryDS}">
  <t:columns>
    <mx:AdvancedDataGridColumn width="100" dataField="Name"/>
    <mx:AdvancedDataGridColumn width="100" dataField="Status"/>
    <mx:AdvancedDataGridColumn width="500" dataField="Database
File Location"/>
  </t:columns>
</t:Table>

```

Retrieving Individual Values From the SQL DataService

To reference a specific cell returned from `SQLDataService` for use within a component (such as `Link` or `Label`), the following type of reference is used:

```

<ds:SQLDataService id="ids" queryID="INSTANCE_INFO"
  properties="{props('TARGET_GUID',appModel.target.guid)}" />

<components:InfoItem label="CPU Model"
  value="{ids.result.getString(0,'CPU Model')}" />

```

The reference to the data service is through `dataService.result.getString(rowIndex, 'column')`, where `rowIndex` is the row returned from the query and `column` is the name of the column as specified in the original SQL query.

The query service can also be called from within a controller, providing much more flexibility in terms of how the data is manipulated before it is displayed. There are two APIs that provide access to the query service:

- **SqlQuery interface**

The `SqlQuery` interface allows for a single SQL query to be processed, passing the bind variable and receiving a result set in return. The result set provides an interface quite similar to that of the `JDBC ResultSet`.

Example 8–16 Using the `SqlQuery` API:

```

var getInfoSvc:SqlQuery = new SqlQuery("GET_TARGET_INFO",
    [{"TARGET", name}, {"TYPE", type}]); // bind variables
getInfoSvc.execute(getTargetInfoHandler);

public function getTargetInfoHandler(resultSet:ResultSet,
fault:ServiceFault):void
{
    var target:Target;
    if(fault == null)
    {
        if(resultSet != null && resultSet.getError() == null)
        {
            target.setGuid(resultSet.getBase64Binary(0, "TARGET_GUID"));
            target.setTypeMetaVer(resultSet.getString(0, "TYPE_META_VER"));

            var props:Array = new Array();
            for(var i:int=1; i<Target.NUM_PROPERTIES+1; i++)
                props.push(resultSet.getString(0, "CATEGORY_PROP_"+i));
            target.setCatProperties(props);
        }
    }
}

```

- BulkSqlQuery interface

The bind variables are referenced by name and correspond to the variables as represented in the packaged SQL statement:

```
SELECT target_guid, type_meta_ver, category_prop_1, category_prop_2,
       category_prop_3, category_prop_4, category_prop_5
FROM mgmt_targets
WHERE target_name = ?TARGET?
AND target_type = ?TYPE?
```

When a number of queries can be processed in a single request, you can use the BulkSqlQuery interface. Each query must be added to the bulk query and when all queries to be processed have been added, the BulkSqlQuery.execute method is called and passed the result handler that will be called with the results.

When a result handler for the SqlQuery is passed a single SqlQueryResultSet for the processed query, the result handler for the BulkSqlQuery is passed a BulkResultSet. Then it must retrieve the SqlQueryResultSet for each query using the request id specified when the query was added.

A separate request id is required to support the case where the same query can be processed multiple times with different bind variables as part of the same bulk request.

Example 8-17 Using the BulkSqlQuery API

```
var guidParam:Array = ["TARGET_GUID",
ApplicationContext.getTargetContext().guid]];

var bulkQuery:BulkSqlQuery = new BulkSqlQuery();
bulkQuery.addQuery("INSTANCE_INFO", "INSTANCE_INFO", guidParam);
bulkQuery.addQuery("PROCESS_STATES", "PROCESS_STATES", guidParam);
bulkQuery.addQuery("PROCESS_INFO", "PROCESS_INFO", guidParam);

bulkQuery.execute(pageDataHandler, page.getBatchRequest());

public function pageDataHandler(bulkResult:BulkResultSet,
fault:ServiceFault):void
{
    var info:SqlQueryResultSet = bulkResult.getResultSet("INSTANCE_INFO");
```

8.11.4.1 Guidelines for Writing Packaged SQL

Adhere to the following guidelines when writing packaged SQL for the MPCUI:

- Packaged SQL can only access views that are part of the partner EDK. This includes any views that are generated as a result of configuration metric definitions.
- Any SQL that attempts to modify data (update or delete) will be filtered by the MRS during plug-in deployment.
- SQL statements that attempt data definition language (DDL) will be filtered out by the MRS and are not allowed
- Anonymous PL/SQL (for example, begin, end constructs) are not allowed as access to PL/SQL procedures is not allowed from packaged SQL

- Bind variables must be identified by a text identifier and prefixed and suffixed by a ?. For example, ?TARGET_TYPE?
- Bind variables are not case sensitive
- The query service restricts the size of result sets to 1000 rows or 100,000 bytes, so care should be taken to limit the size of the possible result set returned by a query.

8.11.4.2 Packaging SQL in the Plug-In

SQL Statements used in the MPCUI code are packaged with the MPCUI metadata using the `SqlStatements` element

For information about the location of SQL statements in the MPCUI metadata, see [Section 8.4.1, "Overview of MPCUI Metadata Elements"](#). For information about the MPCUI metadata XSD, see the EDK Metadata API reference.

8.11.4.3 Getting Target Type Information

For cases where the plug-in UI requires information about a related target type, such as its display name, but does not require the details about a specific instance of that target type, the `TargetFactory` provides a function to retrieve this summary information.

The `TargetFactory.getTargetTypeInfo` function returns a `TargetTypeInfo` object that contains the display name of the target type. When calling this function, pass a `TargetTypeInfo` object with the internal `targetType` provided (for example, "oracle_database") and a handler function. The handler will be called with the `TargetTypeInfo` and any fault that occurred during the processing of the request:

```
var typeInfo:TargetTypeInfo = new TargetTypeInfo("oracle_database");
TargetFactory.getTargetTypeInfo(typeInfo, getTypeInfoHandler);
}

private function getTypeInfoHandler(typeInfo:TargetTypeInfo,
fault:ServiceFault):void
{
    if(fault != null)
    {
        MpLog.logError(fault, "Getting Target Type Info");
        return;
    }

    MpLog.info("Target Display Label for (oracle_database):
"+typeInfo.typeDisplayName);
}
```

8.11.5 Working With Target Services

In addition to the services described previously, the MPCUI framework provides a number of other services that are an integral part of the `Target` object (`oracle.sysman.emx.model.Target`). When the MPCUI application is running, the `ApplicationContext.getTargetContext()` call returns the `Target` instance for the primary target.

You can construct other target instances for associated targets. In either case, use the following methods to obtain additional information for these targets through the MPCUI service layer.

8.11.5.1 Target Properties Service

For any instance of the Target class, you can call the `getTargetInfo()` method to retrieve the target properties associated with that target instance. The returned target information populates the properties of the Target instance including: `guid`, `catProperties`, `typeMetaVer`, `timezoneRegion`, and so on.

For information about these properties, see the Target class documentation in the EDK (`/doc/partnersdk/mpcui/emcore/doc/oracle/sysman/emx/model/Target.html`).

When calling the `getTargetInfo()` method, you must provide a handler. This handler will be called when the `targetInfo` service returns. It is passed the fully populated Target instance and a fault object that is set to include any errors that occurred during the processing of the request to retrieve target properties:

```
var target:Target = ApplicationContext.getTargetContext();
target.getTargetInfo(targetInfoHandler);

public function targetInfoHandler(target:Target, fault:Fault):void
```

Note: In the case of

`ApplicationContext.getTargetContext()`, the current target information is loaded when the application starts and it is not necessary to call `getTargetInfo()` for that target instance unless you think that target properties have changed.

8.11.5.2 Associated Targets Service

Use the `target.getAssociatedTargets()` method to retrieve the set of targets related to a target instance. This method is called and passed an array of association types and a handler that is called with the list of associated targets. Refer to the API documentation for a full description of the types of the objects returned by this method:

```
// get associated host
var target:Target = ApplicationContext.getTargetContext();
var assocTypes:Array = [ AssociationDataService.HOSTED_BY ];
target.getAssociatedTargets(assocTypes, assocHandler);

public function assocHandler(assocResult:GetAssociationsResult,
                             fault:ServiceFault):void
{
    var host:ManageableEntityComponent =
        assocResult.getAssoc(AssociationDataService.HOSTED_BY);
    if(host != null)
        page.setModel("relatedHost", host.name);
}
```

8.11.5.3 Metric Metadata Service

Use the `target.getMetricMetadata()` method to retrieve the metric definitions information for a target instance. The metric metadata information is retrieved by calling the `Target.getMetric()` method which returns a Metric object for a specified metric name. Refer to the API documentation for a full description of the types of the objects returned by this method:

```
var target:Target = ApplicationContext.getTargetContext();
```

```
target.getMetricMetadata(metadataHandler);
```

```
public function metadataHandler (target:Target,  
    fault:Fault):void
```

Note: In the case of `ApplicationContext.getTargetContext()`, the current target metric metadata is loaded when the application starts and it is not necessary to call `getMetricMetadata()` for that target instance unless you think that target metadata has changed (which is unlikely).

8.11.5.4 Availability Service

Use the `target.getAvailability()` method to retrieve current availability information for a target instance. The availability information (`AvailOverviewData`) includes the current status, the up time (%) for the last 24 hours and so on. Refer to the API documentation for a full description of the types of the objects returned by this method:

```
var target:Target = ApplicationContext.getTargetContext();  
target.getAvailability(targetAvailHandler);  
  
public function targetAvailHandler(availInfo:AvailOverviewData,  
    fault:Fault):void
```

8.11.6 Monitoring Service Request Performance

MPCUI includes a tracing service that enables you to monitor the performance of service requests made from the MPCUI code. This is useful when attempting to troubleshoot slow pages or to identify the amount of time spent processing the request in the Management server.

To enable service tracing:

1. Depending on your implementation, choose one of the following:
 - Flex-based UI
 - a. Locate the `html-template/data/mpCuiProperties.xml` file. It is included in the project directories (if you are using the Demo Sample as a template).
 - b. Add the following line to the `mpCuiProperties.xml` file:

```
<traceEnabled>true</traceEnabled>
```
 - c. Rebuild your application, and then launch it using the Flex Builder debugger or run option:
 - Metadata-only UI
 - a. When running the MPCUI page in the console, ensure that the plug-in is deployed.
 - b. Access the target home page associated with the plug-in.
 - c. In the address box of the browser window, add the following to the end of the URL:

```
&traceEnabled=true
```

- d. Press **Enter** to reload the page.

The MPCUI application loads and the home page appears with the Activity Tracing dialog similar to [Figure 8–8](#).

The Activity Tracing dialog displays the set of pages accessed in the current session, and below each page the set of requests made to the Management Server.

Figure 8–8 Activity Tracing Dialog

Page/Service	Start Time	ECID	Total Time (ms)	Service Time (ms)
LoginService.login	3:18:20:943 PM	1D41C425-9188-4605-C6	1,437	n/a
TargetInfoService.getTargetInfo	3:18:22:364 PM	D2BC9EA1-837E-B923-8C	391	167
MetricMetadataService.getTargetMetricMetadata	3:18:22:380 PM	8322373D-1937-0BAF-00	1,234	104
PageActivityDef[homePg]	3:18:22:958 PM	HomePage375	2,171	n/a
JobService.getJobActivity	3:18:23:83 PM	6321D21E-A931-ESA7-C6	1,156	38
IncidentSummaryService.getSummary	3:18:23:364 PM	5CC07725-B295-E96D-A7	843	24
BatchService.submitBatchRequest	3:18:23:395 PM	BAT4F1E0698-B1A9-060E	5,733	2,825
AssociationService.getAssociationsFromRoc	3:18:23:442 PM	35EDD4B3-739F-BD46-4E	922	39
PollService[15].BatchService.submitBatchRe	3:18:23:489 PM	BAT2D3162ED-71E2-233C	5,779	4,569
AvailOverview.getAvailOverview	3:18:23:536 PM	A6EF0501-286B-220F-EE	781	31
AssociationService.getAssociationsFromRoc	3:18:23:551 PM	8697474E-C570-BCF0-EE	891	20

Name	Start Time	End Time	Time (ms)
MarshalInputParamPhase	3:18:23:364 PM	3:18:23:364 PM	0
RemoteCallPhase	3:18:23:379 PM	3:18:24:192 PM	813
UnmarshalOutputRespPhase	3:18:24:192 PM	3:18:24:192 PM	0
DispatchResponse	3:18:24:192 PM	3:18:24:207 PM	15
Client-Overhead	n/a	n/a	15

2. Expand or collapse the dialog using the controls in the upper right-hand corner. It continues to refresh while the MPCUI application is active.

Note: When you select a service request in the top pane of the dialog, the Total Time (round trip) is shown as well as the time spent processing the request in the Management Server (Service Time).

3. In the details pane, click **Show Item Details** to view the body of the request and response messages sent between your application and the Management server.

8.11.7 Automated Polling of Service Requests

Note: An important use of the “REALTIME” data selection for any chart, table, or data service is that it initiates automated polling of the data at the specified interval.

The MPCUI framework supports a limited subset of intervals (15, 30, 60, 90 seconds) so that requests can be grouped together to avoid a large number of requests to the Management Server.

The MPCUI framework starts and stops the polling of these requests automatically as each page or dialog appears or is removed (goes out of scope).

You cannot initiate a polling request that is persistent beyond the scope of a page or dialog.

8.11.8 Batching of Service Requests

In addition to the batching of polling requests, the MPCUI framework provides the ability to explicitly batch requests made at runtime from activity (page or dialog) controllers. Batching of requests is a good practice as it avoids additional round trips to the Management Server which slows the performance of your UI pages and adds additional overhead to the Management Server.

The most common opportunity to batch requests is as part of the activity initialization.

- For data services declared in the page layout (MXML file), the MPCUI framework will batch the requests for you.
- For service requests you make from your `controller.init()` method, you can pass the page's batch request to the service methods. The MPCUI framework calls the `init()` method after your page is loaded.

Example 8–18 is extracted from the `HomeController.as` file in the Demo Sample. Note the instances of `page.getBatchRequest()` in the method. All requests made in this way will be performed over a single pass to the Management Server.

Example 8–18 Batching Requests as Part of the Activity Initialization

```
override public function init(pg:IActivity):void
{
    super.init(pg);
    page = pg as HomePage;

    var guidParam:Array = [{"TARGET_GUID",
        ApplicationContext.getTargetContext().guid}];
    var bulkQuery:BulkSqlQuery = new BulkSqlQuery();
    bulkQuery.addQuery("INSTANCE_INFO", "INSTANCE_INFO", guidParam);
    bulkQuery.addQuery("CPU_USAGE", "CPU_USAGE", guidParam);
    bulkQuery.execute(queryResultHandler, page.getBatchRequest());

    // get processes metric to get process summary information
    var procMetric:Metric = ApplicationContext.getTargetContext()
        .getMetric("CPUProcessesPerf");
    var procSelector:MetricSelector = procMetric
        .getSelector(['ProcUser', 'ProcCPU', 'ProcCmd']);
    procSelector.getData(processesHandler,
        MetricCollectionTimePeriod.CURRENT, page.getBatchRequest());

    var cpuPerf:Metric = ApplicationContext.getTargetContext()
        .getMetric("CPUPerf");
    var cpuPerfSel:MetricSelector = cpuPerf
        .getSelector(['system', 'idle', 'io_wait']);
    cpuPerfSel.getData(cpuDataHandler,
        MetricCollectionTimePeriod.REALTIME, page.getBatchRequest());

    // get associated host
    var target:Target = ApplicationContext.getTargetContext();
    var assocTypes:Array = [ AssociationDataService.HOSTED_BY ];
    target.getAssociatedTargets(assocTypes, assocHandler,
        page.getBatchRequest());
}
```

You can use batch requests elsewhere in controller code by creating a `MultiServiceRequestor` (batch request) and passing it to each request made. For example, suppose that in response to a button click in the page, two requests will be made to the Management Server to retrieve information. They could each be made separately (resulting in two trips to the server) as shown in [Example 8-19](#):

Example 8-19 Creating Individual Batch Requests

```
var procMetric:Metric = ApplicationContext.getTargetContext()
    .getMetric("CPUProcessesPerf");
var procSelector:MetricSelector = procMetric
    .getSelector(['ProcUser', 'ProcCPU', 'ProcCmd']);
procSelector.getData(processesHandler,
    MetricCollectionTimePeriod.CURRENT);    // 1st round trip

var cpuPerf:Metric = ApplicationContext.getTargetContext()
    .getMetric("CPUPerf");
var cpuPerfSel:MetricSelector = cpuPerf.
    getSelector(['system', 'idle', 'io_wait']);
cpuPerfSel.getData(cpuDataHandler,
    MetricCollectionTimePeriod.REALTIME);    // 2nd round trip
```

Alternatively, you can combine the batch requests into a single batch request avoiding the additional round trip to the Management server as shown in [Example 8-20](#):

Example 8-20 Combining Batch Requests

```
var batchRequest:MultiServiceRequestor = new MultiServiceRequestor();

    var procMetric:Metric = ApplicationContext.getTargetContext()
        .getMetric("CPUProcessesPerf");
var procSelector:MetricSelector = procMetric
    .getSelector(['ProcUser', 'ProcCPU', 'ProcCmd']);
procSelector.getData(processesHandler,
    MetricCollectionTimePeriod.CURRENT, batchRequest);

var cpuPerf:Metric = ApplicationContext.getTargetContext()
    .getMetric("CPUPerf");
var cpuPerfSel:MetricSelector = cpuPerf.
    getSelector(['system', 'idle', 'io_wait']);
cpuPerfSel.getData(cpuDataHandler,
    MetricCollectionTimePeriod.REALTIME, batchRequest);

batchRequest.sendRequest();    // 1st and ONLY round trip!
```

Note: You must call the `sendRequest()` method to commit the batch request. Otherwise, no requests will be sent. In the case of the `PageController.init` use of `page.getBatchRequest()`, this is *not* necessary because the MPCUI framework will do it for you.

8.11.9 Software Library Search Service

When the plug-in UI requires information about Software Library entities, it can search using different criteria, including name, status, maturity level, or entity attributes. The desired entities can be filtered by specifying the query criteria using the `SearchField` enumeration. A list of `EntityInfo` objects that represent an entity revision that match

the query criteria is returned. The URN in the EntityInfo object can be used as a unique identifier for the entity revision. If any error has occurred, it will be set in ListSwlibEntitiesResult.errorMessage.

```
// search the software library
var swSearch:ListSwlibEntities = new ListSwlibEntities();
swSearch.addSearchInput(new SearchInput(SearchField.NAME, "oracle"));
var swlib:Swlib = Swlib.getSwLib();
swlib.listEntities(swSearch, swSearchHandler);
}

private function swSearchHandler(result:ListSwlibEntitiesResult,
fault:ServiceFault):void
{
    var r:ListSwlibEntitiesResult; var e:EntityInfo;
    if(fault != null)
    {
        MpLog.logError(fault, "Search Software Library");
        return;
    }

    for(var i:int=0; i<result.swlibEntitiesList.length; i++)
    {
        var entity:EntityInfo = result.swlibEntitiesList[i];
        MpLog.info("Swlib Entity: "+entity.toString());
    }

    page.setModel("swlibContents", result.swlibEntitiesList);
}
```

8.12 Performing Task Automation

The following sections describes how to perform task automation with examples.

It includes the following:

- [Automation Services](#)
- [Working With Credentials](#)

8.12.1 Automation Services

One of the more powerful aspects of the MPCUI framework is the ability to provide access to administrative features through a UI customized to that purpose. The framework supports the processing of administrative tasks through the Enterprise Manager job system and Web services that provide access to the job system.

The MPCUI provides the following job services:

- Job.submit
- Job.runSynchronous
- JobExecution.getStatus
- JobExecution.getDetails
- JobExecution.stopJob
- JobExecution.deleteJob
- JobExecution.getJobDetailsURL

- RemoteOp.performOperation

8.12.1.1 Submitting or Running a Job

The job service allows any job that is registered with the plug-in target type to be submitted for processing. The service does not support the ability to submit system job types at this time.

Scheduling of jobs through the job service supports a limited set of the scheduling options supported by the job system. The job schedule supports the following options:

- Immediate, once, hourly, daily, weekly, monthly, yearly
- Start and end time for repeat submissions
- Repeat count and frequency
- Starting period grace time
- Management Repository or target time zone

Supported job parameter types include Vector, Scalar, Large and ValueOf.

As with other services, the Adobe Flex framework issues requests asynchronously. This requires that a handler is provided that will be called when the request has completed (or failed). When submitting a job, the result handler is called and passed a JobExecution object. This object contains the processing context for the job that was submitted, and can be used to retrieve the status of the job and operate on the job (stop or delete it).

Example 8-21 Submitting a Job

```
var job:Job = new Job("backup", "MyBackup", null,
    ApplicationContext.getTargetContext(),
    [Job.jobParam("dsn", "AdminDS"), Job.jobParam("sql_cmd",stmt)],
    JobSchedule.IMMEDIATE);
job.submit(jobSubmitHandler);
}

private function jobSubmitHandler(exec:JobExecution, fault:Fault):void
{
    // using exec (JobExecution) can now get current status of job,
    // get step details, and start or stop the job
    var execId:JobExecutionId = exec.getExecutionId();
}
```

When a job is run in this way (using the submit method), the job is submitted for processing and the service returns immediately. Therefore, the status of the job may change from submitted to running, and then to complete and the client must check the status periodically.

The job service also provides a way to submit a job for immediate processing and will wait (synchronously) until the job execution completes, fails or reaches a timeout. The client handler will not be called until this state is reached.

Example 8-22 Running a Synchronous Job

```
var job:Job = new Job("backup", "MyBackup", null,
    ApplicationContext.getTargetContext(),
    [Job.jobParam("dsn", "AdminDS"), Job.jobParam("sql_cmd",stmt)],
    JobSchedule.IMMEDIATE);
job.runSynchronous((jobRunHandler, 30); // 2nd param is timeout
}
```

```
private function jobRunHandler(exec:SynchronousJobExecution, fault:Fault):void
{
    // using exec (SynchronousJobExecution) can get details about job execution;
    // this handler will not be called until the job completes, fails
    // or the timeout is reached
    var execId:JobExecutionId = exec.getExecutionId();
}
```

The Task interface is a simplified way of submitting a job for immediate processing, without requiring all of the additional settings associated with the `Job.submitJob` API.

Example 8–23 Using the Task API

```
var task:Task = new Task("TTisql", null, [Job.jobParam("dsn",
"AdminDS"),
                                Job.jobParam("sql_cmd",stmt)]);
task.execute(createTableHandler, 10); // timeout is 10s
}

private function createTableHandler(result:SynchronousJobExecution,
fault:Fault):void
{
    var status:JobStatus = result.getRunStatus();
    if(status == JobStatus.RUNNING)
    {
        // timed out while waiting for job... still running
    }
}
```

In the case of a synchronous job, the status of the job is available immediately from the result passed to the handler; however, it should be checked to see if it equals `JobStatus.RUNNING`. If so, then the request reached the specified timeout and the caller must treat the job execution as if it were submitted asynchronously.

8.12.1.2 Getting Job Status and Step Details

After a job has been submitted, there are several APIs available to get the status of the job and the details of each step including job output. To use these APIs, the caller must have a valid `JobExecution` object, which is passed to the result handlers of `submit` and `runSynchronous` APIs. Currently, there is no service provided that allows a client to search for a job execution.

Example 8–24 Getting Job Status

```
private function submitHandler(exec:JobExecution, fault:Fault):void
{
    exec.getStatus(statusHandler);
}

private function statusHandler(status:JobStatus, fault:Fault):void
{
    if(status.getStatus() == JobStatus.FINISHED)
```

Getting the job status for a submitted job requires service request, and therefore requires a handler to be called with the result (and possibly a fault if the request processing fails). In addition to the status of the job, the job step details can be retrieved.

Getting Job Details:

Use the `JobExecution` object that is passed to the submit handler, to retrieve step output details as well as the job status. If the job has failed or if the step has not completed, then no data will be returned.

In the case of a synchronous job execution, the handler for the `Job.runSynchronous` or `Task.execute` call can check the job status and if complete, retrieve the job details from the result directly:

```
private function createTableHandler(result:SynchronousJobExecution,
fault:Fault):void
{
    var status:JobStatus = result.getRunStatus();
    if(result != null && result.Status() == JobStatus.COMPLETED)
    {
        var steps:Array = result.getStepDetail();
        for(var i:int=0; i<steps.length; i++)
        {
            var detail:JobStepDetail = steps[i];
            proc.addDetailText(detail.output);
        }
    }
}
```

In the case of an asynchronous job execution, the handler for the `Job.submit` handler must call `JobExecution.getStatus`, and then `JobExecution.getDetails`. Each call requires a request to the server:

```
private function submitJob():void
{
    var params:Array = new Array();
    params.push(Job.jobParam("p0", "p0value"));
    var job:Job = new Job(type, name, desc, ApplicationContext.getTarget(),
params, Job.immediateSchedule());
    job.submit(submitHandler);
}

private function submitHandler(result:JobExecution, fault:Fault):void
{
    if(fault == null && result != null)
    {
        // get the job status; calls the server
        result.getStatus(statusHandler);
    }
}

private function statusHandler(result:JobStatus, fault:Fault):void
{
    if(fault == null && result != null)
    {
        if(result.getStatus() == JobStatus.COMPLETED)
        {
            // now can get job output details
            result.getJobExecution().getDetails(detailsHandler);
        }
    }
}

private function detailsHandler(result:JobExecutionDetails, fault:Fault):void
{
    if(fault == null && result != null)
    {
        var steps:Array = result.getStepDetail();
        for(var i:int=0; i<steps.length; i++)
        {
            // ...
        }
    }
}
```

```
        {  
            var detail:JobStepDetail = steps[i];  
            proc.addDetailText(detail.output);  
        }  
    }  
}
```

8.12.1.3 Using a Timer to Periodically Check Job Status

If a job is submitted asynchronously (`Job.submit`) or runs synchronously but the request reaches a timeout and returns a job status of `JobStatus.RUNNING`. This indicates that the job is still running, and you might want to check the status of the job again at a later point.

The easiest thing from a code perspective is to expose a UI element that the user interacts with to cause the application to check the status of the job. For example, the UI might show a *Running* indicator with a button or link labeled **Check Status Now**. When the user clicks the button or link, it calls the `JobExecution.getStatus` method to retrieve the updated status.

If the required interaction pattern is that the UI remains active while the job is running in the background, and periodically updating the UI with information about the status of the job, then the MPCUI provides a job API to perform period checking of the job status. Each update calls a handler to provide the caller with the opportunity to process the current status and update the UI with that information.

8.12.1.4 Stopping and Deleting a Job

Jobs submitted through the MPCUI APIs can be stopped or deleted using the following APIs:

- [Stopping a Job](#)
- [Deleting a Job](#)

Example 8–25 Stopping a Job

```
private function stopJob(exec:JobExecution):void  
{  
    // NOTE: the JobExecution must be a valid job context obtained from submitted  
    a job  
    exec.stopJob(stopJobHandler);  
}  
  
private function stopJobHandler(exec:JobExecution, fault:Fault):void  
{  
    if(fault == null && exec != null)  
    {  
        // job was successfully stopped  
    }  
}
```

Example 8–26 Deleting a Job

```
private function deleteJob(exec:JobExecution):void  
{  
    // NOTE: the JobExecution must be a valid job context obtained from submitted  
    a job
```

```

        exec.deleteJob(deleteJobHandler);
    }

    private function deleteJobHandler(exec:JobExecution, fault:Fault):void
    {
        if(fault == null && exec != null)
        {
            // job was successfully deleted
        }
    }
}

```

For jobs that are submitted using the `Job.runSynchronous` API, the job can be deleted when completed by passing the `deleteWhenFinished` parameter as `true`. It is the third parameter and it defaults to `false`:

```

var job:Job = new Job("backup", "MyBackup", null,
    ApplicationContext.getTargetContext(),
    [Job.jobParam("dsn", "AdminDS"), Job.jobParam("sql_cmd", stmt)],
    JobSchedule.IMMEDIATE);
job.runSynchronous((jobRunHandler, 30, true);

```

8.12.1.5 Remote Operations

Using a job to perform administrative tasks is the most flexible approach in terms of scheduling and control (start, interrupt, or stop), but does come with the additional overhead of managing the task being processed. For simple tasks that do not require control over schedule and that are expected to be performed quickly, use the `RemoteOp` service.

This service allows the execution of a script packaged with the plug-in to be executed directly through the Management Agent.

Note: The script must be packaged with the plug-in in the agent/scripts directory (as described in the following section), and might require credentials or parameters to be processed.

Packaging Scripts for Remote Operation

Scripts included in a plug-in for remote operations must be included in the staging area:

```
stage/agent/scripts
```

You can create additional subdirectories under `/scripts` if required. Scripts placed in this location can be referenced using the `RemoteOp` service by referencing the `%scriptsDir%` variable. For example:

Plug-in Stage Area

```
./stage/agent/scripts/process/kill_process.pl
```

MPCUI Code (ActionScript)

```

var params:Array = [
    RemoteOp.param("%scriptsDir%/process/kill_process.pl"),
    RemoteOp.param(pid) ];
var remoteOp:RemoteOp = new RemoteOp("perl", params);
remoteOp.performOperation(killProcessHandler);

```

In this example, a `RemoteOp` object is constructed using the shell / command to run and the parameters to pass into that shell. The first parameter must always be the location of the script to be run, referencing its location relative to the `%scriptsDir%` variable. Subsequent parameters are included as required for the script being run.

To run the remote operation, the `RemoteOp.performOperation` method is called and passed a function that will be called when the remote operation completes processing. This handler has the following signature:

```
private function killProcessHandler(remoteOp:RemoteOp, fault:Fault):void
```

If the remote operation failed to be communicated to the Management Agent, then the `fault` parameter will include the details of that error. If the remote operation was processed, then the fault will be null and the `remoteOp` parameter supplied.

Check the `remoteOp` parameter status because it can indicate an error status returned during command execution on the Management Agent. [Example 8-27](#) shows this check being performed.

Example 8-27 Checking the remoteOp Parameter Status

```
/**
 * Check status; could be any number of problems some of which may result
 * in step output, some of which (like missing creds) result in a non-successful
 * run status but no step details.
 *
 * result.getRunStatus() - the status of the job, refer to Constants.JOB_*_STATUS
 * result.getStepDetail().stepName/detail - name and output from each step in the
job
 * result.getJob() - the complete job Object, to reference parameters:
 *   result.getJob().parameter[0].paramName/paramValue[0]
 *
 */
if(remoteOp.result.status != Constants.JOB_COMPLETED_STATUS)
{
    // job did *NOT* complete successfully
    var pid:String = remoteOp.getParameter(2).paramValue[0] as String;
    var msg:String = "Unable to successfully kill process ["+pid+
        "]. The status of the command was: "
        +Util.getCatString(Util.JOB_STATUS, remoteOp.result.status)
        +"\nReturn Code: "+remoteOp.result.returnValue
        +"\nCommand Output: "+remoteOp.result.commandOutput;
    MessageAlert.show(msg, "Failed to Kill Process", Alert.OK);
}
else
{
    // successful job execution; process was killed; can look at the
    // step details to get possible output from the job
    MpLog.debug("Command was successful: "
        +"\nReturn Code: "+remoteOp.result.returnValue
        +"\nCommand Output: "+remoteOp.result.commandOutput);
    page.processesTable.refreshImmediate();
}
```

8.12.2 Working With Credentials

The Enterprise Manager credentials model supports three different modes for performing operations that require credentials:

- Preferred Credentials

Specific credentials are set for a target or all targets of a particular type. In this mode, the user does not select a set of credentials or provide credential values.

- **Named Credential Set**

Sets of credentials are created for a target or all targets of a particular type, and each set is assigned a name. In this mode, the user is presented with a list of named sets and can select the set that they would use to perform the operation.

- **Override Credentials**

In this mode, the user can supply credentials at runtime that are used to perform the operation.

8.12.2.1 Retrieving Credential Information

MPCUI provides the facilities for retrieving credential information about a particular target. The services can return information only that the current user is privileged to see. This ensures that there is no unauthorized access to secure information. It also requires that you must handle a situation where credential information might not be available to the user accessing the MPCUI code.

8.12.2.1.1 Check for Preferred Credentials To check if a target has preferred credentials set for it, call the `CheckPreferredCredsService.getPreferredCredsInfo` method as follows:

```
var ccSvc:CheckPreferredCredsService = new CheckPreferredCredsService();
ccSvc.getPreferredCredsInfo(ApplicationContext.getTargetName(),
    ApplicationContext.getTargetType(),
    'HostCreds', checkPrefCredsHandler);
```

The service returns a `CheckPreferredCredsResult` object, which indicates whether global (applying to all target instances of the type) or instance (applying only to the single target instance) credentials are available:

```
public function checkPrefCredsHandler(result:CheckPreferredCredsResult,
    fault:ServiceFault):void
{
    if(fault != null)
        MessageAlert.show(fault.faultDetail, "Error Checking Preferred Creds");
    else
    {
        var msg:String = "Checked for Preferred Credentials for
            target (" + ApplicationContext.getTargetName() + ", " +
            ApplicationContext.getTargetType() + " for set (HostCreds)
            user (" + ApplicationContext.getEmUser() + ") \n" +
            "Global Set = " + result.globalSet + " Instance Set = "
            + result.instanceSet;
        MessageAlert.show(msg, "Check Preferred Creds Result");
    }
}
```

Note: If preferred credentials are set, you can submit a job or perform a remote operation without passing any credentials information. In this case, the preferred set will be used.

8.12.2.1.2 Retrieve Named Credentials Sets You can retrieve the named credentials sets available for a particular target and:

- display the named credentials set in a choice (list or combo box)
- select from the named credentials set based on a convention required by your plug-in

The following code requests all named sets for two different credentials types for the current target, and calls the `credSetResultHandler` handler with the result:

```
var target:Target = ApplicationContext.getTargetContext();
target.getCredentialSets(["HostCreds", "HostSampleCreds"],
    credSetResultHandler);
```

The results handler can then consume the named sets return as appropriate (in this example, constructing a data source for display in a table):

```
var credTableData:ArrayCollection;
if(creds.credSet != null)
{
    credTableData = new ArrayCollection(creds.credSet);

    // check to see if there are sets for both types
    var hostFound:Boolean = false;
    var sampFound:Boolean = false;
    for(var c:int=0; c<creds.credSet.length; c++)
    {
        if(creds.credSet[c].credentialType == "HostCreds")
            hostFound = true;
        else if(creds.credSet[c].credentialType == "HostSampleCreds")
            sampFound = true;
    }

    var missingType:String = ( !hostFound ? "HostCreds" :
                                "HostSampleCreds" );
    var empty:CredentialSet = new CredentialSet();
    empty.credentialType = missingType;
    empty.name = "<No Sets Found>";
    empty.guid = "";
    credTableData.addItem(empty);
}
else
{
    empty = new CredentialSet();
    empty.credentialType = "<No Credential Sets Defined>";
    empty.name = "";
    empty.guid = "";
    credTableData = new ArrayCollection([empty]);
}
```

8.12.2.2 Passing Credentials to Jobs and Remote Operations

This section discusses passing preferred credentials and named set credentials to jobs and remote operations.

Preferred Credentials

If the task (job or operation) to be performed attempts to use preferred credentials, then the credentials parameter passed to the task is omitted. Both the job and remote operation services will attempt to perform the task using preferred credentials. If no preferred credentials are set, then an error will be returned

Named Set

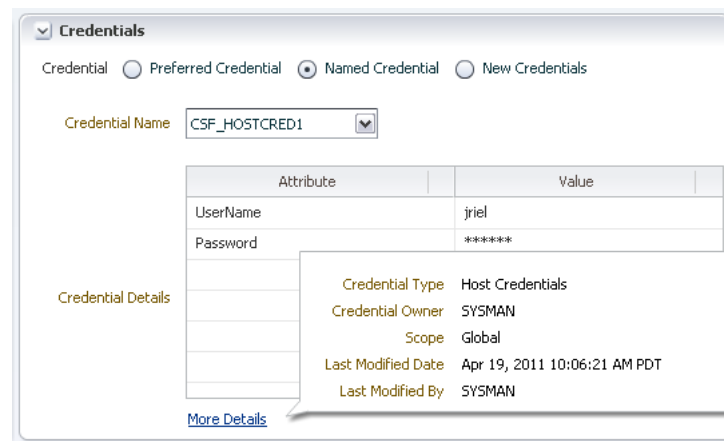
To specify that a named set be used to perform a task, the credentials are passed in a `JobCredential` (for jobs) and an `OpCredential` (for remote operations). In both cases, the credentials object includes the following four properties that must be set:

- `targetName`: the target the credentials apply to, usually `ApplicationContext.getTargetName()`
- `targetType`: the type of the target the credentials apply to, usually `ApplicationContext.getTargetType()`
- `usage`: the credentials usage as defined for the operation (see the job type definition). This usage specifies which credentials types are required and where they are applied during job execution
- `credGuid`: the identifier of the named set to be used. This is one of the properties of the `CredentialSet` class, which holds named credential sets. For more information, [Section 8.12.2.1, "Retrieving Credential Information"](#).

8.12.2.3 Reusable Credentials UI Components

MPCUI provides a credentials region that may be included in a page to allow the end user to interact with the Enterprise Manager credentials subsystem to view the set of credentials available and to select preferred, named, or override credentials when performing a task (job or remote operation).

Figure 8–9 Credentials Region



To include this region in a page, add the following MXML:

Example 8–28 Adding a Credentials Region

```
<credentials:CredentialsRegion id="credRegion" width="40%" height="100%"
  title="Credentials" target="{ApplicationContext.getTargetContext()}"
  credentialType="HostCreds" />
```

From the page controller associated with the page, retrieve the settings applied by the end user to this region:

Example 8–29 Retrieving Selected Credential Information

```
public function getCredsEntered(event:MouseEvent):void
{
    var mode:String = page.credRegion.getMode();
```

```

var msg:String = "Credential Option Selected: "+mode+"\n";

var namedSet:String;
var creds:Array;
if(mode == CredentialsRegion.NAMED_MODE)
{
    namedSet = page.credRegion.getNamedSet();
    msg += "Named Set Selected: "+namedSet;
}
else if(mode == CredentialsRegion.OVERRIDE_MODE)
{
    try
    {
        creds = page.credRegion.getOverrideCredentials();
        for(var c:int=0; c<creds.length; c++)
            msg += "Field:"+creds[c].label+", "+creds[c].value+"\n";
    }
    catch(e:Error)
    {
        msg += "Error Entering Credentials:\n";
        msg += e.message;
    }
}
else
{
    // preferred selected...
}
MessageAlert.show(msg, "Credentials Entered");
}

```

In [Example 8–29](#), note that the mode determines if the user selected preferred, named, or override credentials. Depending on the mode, the named set can be retrieved (`CredentialsRegion.getNamedSet()`) or the override credentials can be retrieved (`CredentialsRegion.getOverrideCredentials()`).

8.12.2.4 Managing Monitoring Credentials

The Target class provides the ability to retrieve and set the monitoring credentials for the current target. To retrieve the monitoring credentials, an instance of the Target class is required and the `getMonitoringCredentials` function is called, passing the results handler that will receive the credential meta data including the monitoring credentials set:

```

// get monitoring credentials
target.getMonitoringCredentials(getMonCredResultHandler, page.getBatchRequest());

```

The handler would appear as follows:

```

private function getMonCredResultHandler(cred:CredentialTypeMetadata,
fault:ServiceFault):void
{
    if(fault != null)
    {
        if(cred != null && cred.isMissingCredentials())
        {
            // no monitoring credentials are set
            MpLog.info("Monitoring Credentials have not been set:
"+fault.faultDetail);
        }
    }
    else

```



```

        MpLog.logError(fault, "Get Monitoring Credentials");
    return;
}

/**
 * The CredentialTypeMetadata returned includes the meta-data for the
 * credentials as well as the actual values. NOTE: credentials never
 * return the actual values for any field identified as a password it only
 * returns the masked "****" value. You should never have any need to
 * access the actual values for a password field as any time credentials
 * are passed you are passing a credential set and don't need the actual
 * values of a pre-existing credentials set
 */

    var credFieldValues:Array = (cred.credentialSets[0] as
CredentialSet).columnValues;
    for(var i:int=0; i<credFieldValues.length; i++)
    {
        var credField:CredentialColumnValue = credFieldValues[i];
        MpLog.debug("Monitoring Credentials["+credField.label+"] =
"+credField.value);
    }
}

```

To set the monitoring credentials, the credentials fields according to the credential type specified for monitoring credentials. This is defined in your target metadata.

```

/**
 * the CredentialSet passed contains the monitoring credentials to be set.
 * Note that only the columnValues property of the credentials needs
 * to be set when updating monitoring credentials as the framework
 * derives the values for the credential set and type. It is CRITICAL
 * that the label set for each columnValue is the column NAME and not
 * the display label for that column. The name is the identifier assigned
 * to the credential column in the target meta-data.
 *
 * In the demo_hostsampl, for example, the credentials fields are:
 *   name: SampleCredUser      label: User Id
 *   name: SampleCredPassword  label: Password
 *   name: SampleCredRole      label: Role
 */

var monitoringCreds:CredentialSet = new CredentialSet();
monitoringCreds.columnValues = [
    new CredentialColumnValue("SampleCredUser", "myMonitoringUser"),
    new CredentialColumnValue("SampleCredPassword", "myMonitoringPassword"),
    new CredentialColumnValue("SampleCredRole", "myMonitoringRole")
];

var target:Target = ApplicationContext.getTargetContext();
target.setMonitoringCredentials(monitoringCreds, setMonCredResultHandler);

```

Then the handler would appear as:

```

private function setMonCredResultHandler(cred:CredentialSet,
fault:ServiceFault):void
{
    if(fault != null)
    {
        MpLog.logError(fault, "Set Monitoring Credentials");
        return;
    }
}

```

```
    }

    /**
     * if the set monitoring credentials was successful then the handler is
     * called with no fault and the set of credentials that were passed in
     */
}
```

8.13 Storing Session State

The session state service provides the ability to store global state associated with the Flex application. This is useful for cases where state should be maintained for the current user, even as they move between pages outside of the Flex application that defines the custom UI. For example, if the user modifies the state of the home page and then navigates to the "All Metrics" page, and then upon returning to the home page you wish to restore the state of the page as the user left it. Because the user has left the pages that make up the Flex application, it is necessary to store the state required on the server-side session established for this user session and not within the Flex application itself.

To set the session state, call the `EmUser.setSessionData` function, passing a `SessionAttributes` object. The session attributes contain an array of `SessionAttribute` objects, each of which has a corresponding name-value pair for the attributes stored.

```
public function setSessionAttributes():void
{
    var sessionData:SessionAttributes = new SessionAttributes();
    var item1Value:String = page.item1Input.text;
    var item2Value:String = page.item2Input.text;

    sessionData.attributes.push(new SessionAttribute("attr1",
        item1Value));
    sessionData.attributes.push(new SessionAttribute("attr2",
        item2Value));

    EmUser.setSessionData(sessionData, setSessAttrResultHandler);
}

public function setSessAttrResultHandler(attr:SessionAttributes,
    fault:ServiceFault):void
{
    {
        if(fault != null)
        {
            MpLog.logError(fault, "Set Session Data");
            return;
        }
    }
}
```

To retrieve the session state, use the corresponding `EmUser.getSessionData` service function. This function will retrieve the session state requested by passing a list of `SessionAttributes` and a handler that will be called with the result. This will be the same `SessionAttributes`, populated with the data retrieved from the session:

```
public function getSessionAttributes():void
{
    var sessionData:SessionAttributes = new SessionAttributes();
    sessionData.attributes.push(new SessionAttribute("attr1"));
    sessionData.attributes.push(new SessionAttribute("attr2"));
    EmUser.getSessionData(sessionData, getSessAttrResultHandler);
}
```

```

public function getSessAttrResultHandler(attr:SessionAttributes,
fault:ServiceFault):void
{
    if(fault != null)
    {
        MpLog.logError(fault, "Get Session Data");
        return;
    }

    for(var i:int=0; i<attr.attributes.length; i++)
    {
        var item:SessionAttribute = attr.attributes[i];
        if(item.name == "attr1")
            page.setModel("lastItem1Value", item.value);
        else if(item.name == "attr2")
            page.setModel("lastItem2Value", item.value);
    }
}

```

8.14 Defining Page Layout Components

To ensure that the MPCUI page resizes correctly when the browser window resizes, Oracle recommends the following guidelines for page layout of an MPCUI-based page:

- Use the HBox (horizontal box) and VBox (vertical box) containers
- Set the height and width of the containers using percentage sizes and not absolute pixel sizes

For example, to create a layout of three rows, each occupying one third of the height of the page, then enter the following in the MXML file:

Example 8–30 Defining a Page Layout of Three Rows

```

<mx:VBox height="100%" width="100%">
    <!-- 1st row -->
    <mx:HBox height="33%" width="100%">
    </mx:HBox>

    <!-- 2nd row -->
    <mx:HBox height="33%" width="100%">
    </mx:HBox>

    <!-- 3rd row -->
    <mx:HBox height="33%" width="100%">
    </mx:HBox>
</mx:VBox>

```

Then enter the following to split each row horizontally into two separate or equal sections:

Example 8–31 Splitting Rows into Two Equal Sections

```

<mx:VBox height="100%" width="100%">
    <!-- 1st row -->
    <mx:HBox height="33%" width="100%">
        <mx:VBox height="100%" width="50%" >
        </mx:VBox>
    </mx:HBox>

```

```

        <mx:VBox height="100%" width="50%" >
        </mx:VBox>
    </mx:HBox>

    <!-- 2nd row -->
    <mx:HBox height="33%" width="100%">
        <mx:VBox height="100%" width="50%" >
        </mx:VBox>
        <mx:VBox height="100%" width="50%" >
        </mx:VBox>
    </mx:HBox>

    <!-- 3rd row -->
    <mx:HBox height="33%" width="100%">
        <mx:VBox height="100%" width="50%" >
        </mx:VBox>
        <mx:VBox height="100%" width="50%" >
        </mx:VBox>
    </mx:HBox>
</mx:VBox>

```

Within each section, include individual components to fill out the layout of the page.

8.14.1 Defining Regions

The Enterprise Manager UI style guides recommends that you organize the information on the page into *regions*. A region is a visual box with a title that can be expanded and collapsed. For example, in [Example 8–30](#), each of the rows could be split up into separate regions rather than more vertical containers:

Example 8–32 Defining Regions

```

<mx:VBox height="100%" width="100%">
    <!-- 1st row -->
    <mx:HBox height="33%" width="100%">
        <components:Region height="100%" width="50%" title="Row 1 Region 1" >
        </components:Region>
        <components:Region height="100%" width="50%" title="Row 1 Region 2" >
        </components:Region>
    </mx:HBox>

    <!-- 2nd row -->
    <mx:HBox height="33%" width="100%">
        <components:Region height="100%" width="50%" title="Row 2 Region 1" >
        </components:Region>
        <components:Region height="100%" width="50%" title="Row 2 Region 2" >
        </components:Region>
    </mx:HBox>

    <!-- 3rd row -->
    <mx:HBox height="33%" width="100%">
        <components:Region height="100%" width="50%" title="Row 3 Region 1" >
        </components:Region>
        <components:Region height="100%" width="50%" title="Row 3 Region 2" >
        </components:Region>
    </mx:HBox>
</mx:VBox>

```

[Example 8–32](#) results in a display similar to [Figure 8–10](#). You can use each of the regions to contain other UI components (such as tables and charts) to display

meaningful information. For more detailed examples, see the examples in the Demo Sample included in the Extensibility Development Kit.

Figure 8–10 Regions



8.15 Including Packaged Regions

The Region component is an empty container within which you can display any number of components to construct your UI. MPCUI supplies several packaged regions that can be included in your page with a single simple tag

8.15.1 Availability Region

The availability region displays the availability of the target for the period specified in the AvailabilityRegion tag daySpan property. It shows a segmented bar that shows details of the target availability (such as outages) over that same period:

```
<avail:AvailabilityRegion width="25%" height="100%" daySpan="1"/>
```

Figure 8–11 Availability Region

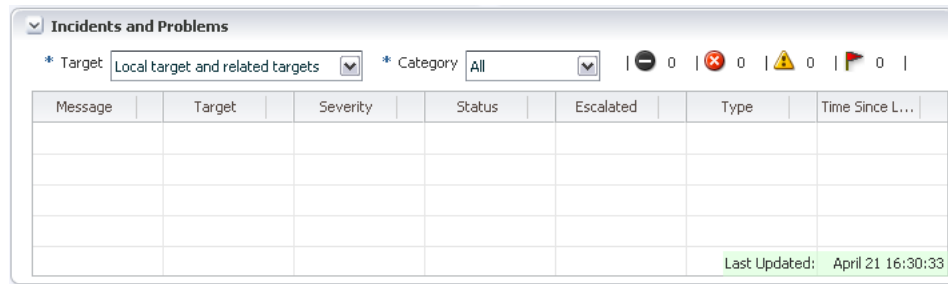


8.15.2 Incidents and Problems Region

The incidents region shows the set of open incidents for the current target and all related targets. It provides the option to filter the list of displayed incidents. The only necessary settings for the region are the size (width/height):

```
<events:IncidentRegion width="50%" height="100%"/>
```

Figure 8–12 Incidents and Problems

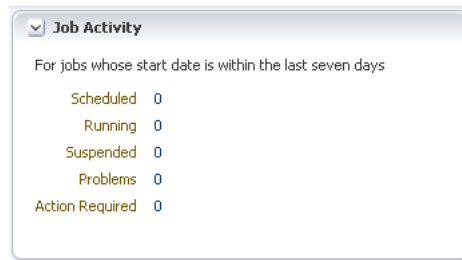


8.15.3 Job Summary Region

The jobs summary region displays the count of jobs by status.

```
<jobs:JobSummaryRegion width="20%" height="100%" />
```

Figure 8–13 Job Summary



8.15.4 Credentials Region

For information about reusable credentials UI components, see [Section 8.12.2.3, "Reusable Credentials UI Components"](#).

8.16 Defining Charts

MPCUI supports three chart components. All chart components have integral support for displaying metric information by specifying the metric properties. Additionally, you can construct your own data for the chart using information obtained from other services including `SQLDataService` and map it to the charts using the `dataProvider` property.

The following examples and documentation for each chart type are a brief summary of the various options available for each chart. For a complete description of each chart's properties, refer to the API documentation. For examples of how these charts work at runtime, see the Demo Sample included in the Extensibility Development Kit.

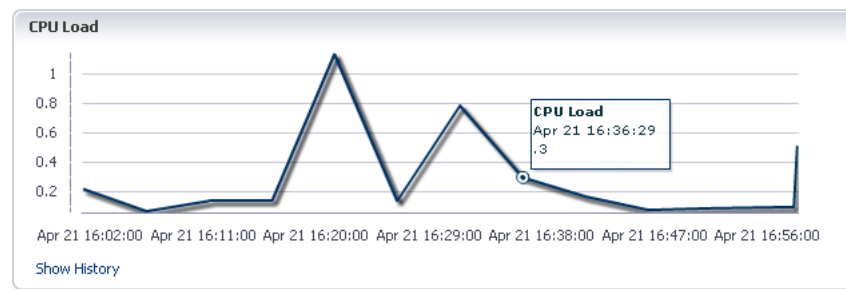
8.16.1 Line Chart

Typically, the line chart displays information over time (often referred to as a time-series chart). Therefore, it lends itself to the display of metric information either historically or in real-time. The chart includes properties for specifying the

`metricName` and `metricColumns` (an array) that should be shown in the chart and a `timePeriod` property that can be set to show historical data or real-time sampled data. When `timePeriod` is set to "REALTIME", the chart manages an automatic polling request for you and updates the chart data as new samples arrive.

```
<components:Region title="CPU Load" width="75%" height="100%" >
  <charts:LineChart id="cpuload" width="100%" height="100%"
    metricName="Response"
    metricColumns=["Load"]
    timePeriod="REALTIME" interval="15" />
  <components:Link label="Show History"
    click="{invokeActivity('metricHistory')}}" />
</components:Region>
```

Figure 8–14 Example of a Line Chart



8.16.1.1 Providing Line Chart Data Source

In addition to specifying metrics to be plotted using the line chart, you can create your own data source that is used by the chart to display data. For example, data obtained through the SQL data service or some other means such as by using the polling service and then creating the data samples to be added to the chart.

In the following example, the page includes a chart with the `chartDataSource` mapped to an item in the page model that is constructed in the page controller.

■ **ProcessesPage.mxml**

```
<ch:LineChart id="cpuUtilChart" width="100%" height="100%"
  chartDataSource="{model.cpuChartData}" />
```

■ **ProcessesPageController.as** (init method)

```
// setup a data provider for the CPU line chart; it will be
// updated each time a new data sample comes back for this metric

// first get the polling context for a 15 second interval
var pollingCtx:PollingContext =
    page.pollingContext.getContext(PollingInterval.EVERY_15_SECONDS);

// now get the metric to be selected and initiate the request (won't start
until
// "startPolling" is called)
var cpuPerf:Metric =
    ApplicationContext.getTargetContext().getMetric("CPUPerf");
var cpuPerfSel:MetricSelector = cpuPerf.getSelector(['system', 'idle', 'io_
wait']);
cpuPerfSel.getData(cpuDataHandler, MetricCollectionTimePeriod.REALTIME,
pollingCtx);
```

```
// start polling; this will automatically stop when the user moves to another
page
pollingCtx.startPolling();
```

■ **ProcessesPageController.as** (cpuDataHandler method)

```
public function cpuDataHandler(cpuData:MetricResultSet,
fault:ServiceFault):void
{
    if(fault != null) return;    // handle this better!

    // get the current data point and derive a new one to
    // add to the charts data source
    var dataPoint:TimestampMetricData = cpuData.results[0];
    var systemTime:Number = dataPoint.data[0]['system'];
    var ioWaitTime:Number = dataPoint.data[0]['io_wait'];

    // create a new data point; this is added to the chart
    // data source (ChartData) below
    var dataSample:ChartDataSample = new ChartDataSample();
    dataSample.timestamp = dataPoint.timestamp;
    dataSample["cpuTime"] = systemTime + ioWaitTime;

    // check if the chart data source is there yet and if
    // not create it and add it to the page model
    var cpuChartData:ChartData = page.model["cpuChartData"] as ChartData;
    if(cpuChartData == null)
    {
        cpuChartData = new ChartData();
        page.setModel("cpuChartData", cpuChartData);

        // define the series "cpuTime" in the chart including a label
        page.cpuUtilChart.setLineSeries(["cpuTime"], ["CPU Time %"]);
    }
    cpuChartData.addDatapoint(dataSample);
}
```

8.16.1.2 Controlling the Legend

All charts can include a legend that displays the items shown in the chart. Use the following example to position the legend in one of four locations (top, bottom, left, right).

```
<c:LineChart id="hchart" timePeriod="REALTIME"
            showLegend="true" legendLocation="top"
```

8.16.2 Area Chart

The area chart is similar to the line chart and has the same attributes. It displays data in the same way as LineChart. The showCumulativeLine property controls the display of an area chart. For most area charts, this property should be included in set to "true" to show a stacked or cumulative area chart. Otherwise, the area chart overlays the fill areas for each series included in the chart.

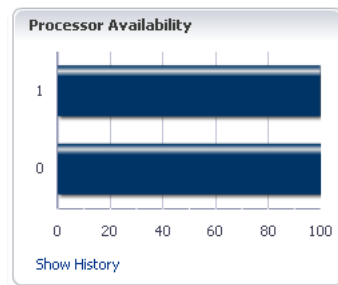
```
<charts:AreaChart id="cpuutil" width="100%" height="100%"
    metricName="CPUProcessorPerf"
    metricColumns="{['CPUIdle']}"
    timePeriod="LAST_DAY" />
```


8.16.3 Bar (Horizontal) Chart

The bar chart exposes the same properties as the column chart both for visible attributes and for specifying control over the data source:

```
<charts:BarChart id="spaceChart" timePeriod="CURRENT"
    width="100%" height="100%"
    groupBy="byKey"
    metricName="MSSQL_Database"
    metricColumns="{['spaceavailable']}" />
```

Figure 8–15 Bar Chart



8.16.3.1 Grouping Bars

The `groupBy` property (available for bar and column charts) enables you to organize data by key or by column. The default (by column) applies when the data set does not include keys.

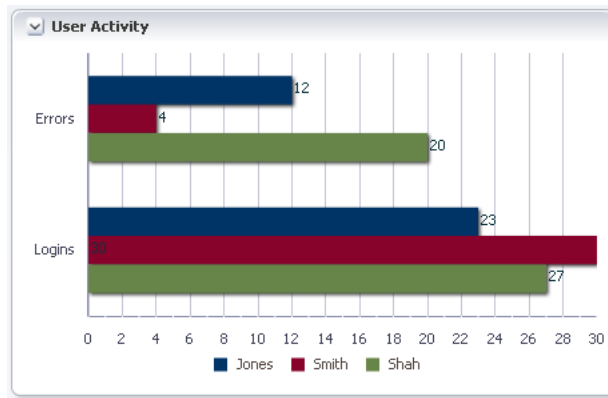
For example, assume you have the following data set where the `User` column is treated as the key to the data:

User	Logins	Errors
Jones	23	12
Smith	30	4
Shah	27	20

In [Example 8–33](#), the `groupBy` property is set to **byColumn**. This creates two groups of columns, one for each data column, with all three keys appearing in each group as displayed in [Figure 8–16](#).

Example 8–33 Group By Column

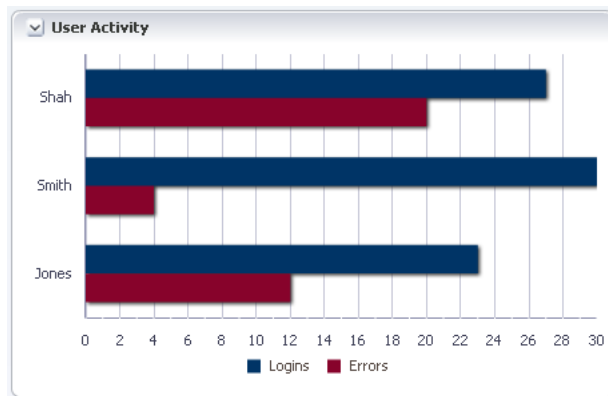
```
<mp:BarChart id="userBarChart"
    dataProvider="{model.userData}"
    showLegend="true"
    groupBy="byColumn"
/>
```

Figure 8–16 Group By Column

In [Example 8–34](#), the `groupBy` property is set to **byKey**. This creates three groups, one for each key, with both columns (the data items) appearing in each group as displayed in [Figure 8–17](#).

Example 8–34 Group by Key

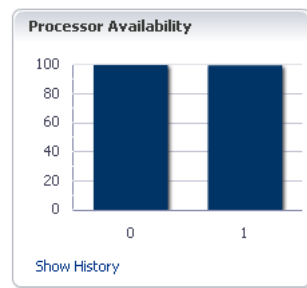
```
<mp:BarChart id="userBarChart"
              dataProvider="{model.userData}"
              showLegend="true"
              groupBy="byKey"
            />
```

Figure 8–17 Group By Key

8.16.4 Column (Vertical Bar) Chart

The column chart is a vertical bar chart and exposes the same properties as the bar chart both for visible attributes and for specifying control over the data source:

```
<charts:ColumnChart id="bchart" timePeriod="LAST_DAY"
                     width="100%" groupBy="byKey"
                     metricName="CPUProcessorPerf" metricColumns="{['CPUIdle']}" />
```

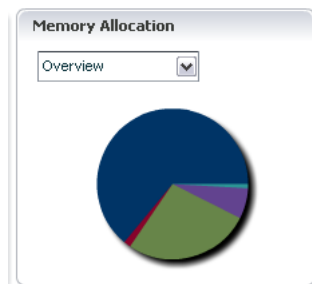
Figure 8–18 Column Chart

8.16.5 Pie Chart

In the following example, the code constructs a pie chart by specifying the metric name and metric columns. The MPCUI framework performs the necessary requests to obtain information from the Management Server and populates the values in the chart.

Note: For the `metricColumns` attribute, the value is set in the controller (see the `HomeController.as` example) in response to the user changing the value of the combo box above the chart.

```
<charts:PieChart id="memChart"
  targetName="{appModel.target.name}"
  targetType="{appModel.target.type}"
  metricName="MemoryPerf"
  metricColumns="{model.memoryColumns}"
  timePeriod="LAST_DAY"
  labelPosition="none" />
```

Figure 8–19 Pie Chart

8.17 Defining Tables

The following sections describe the different methods of defining tables, providing examples of each method.

8.17.1 Data Service

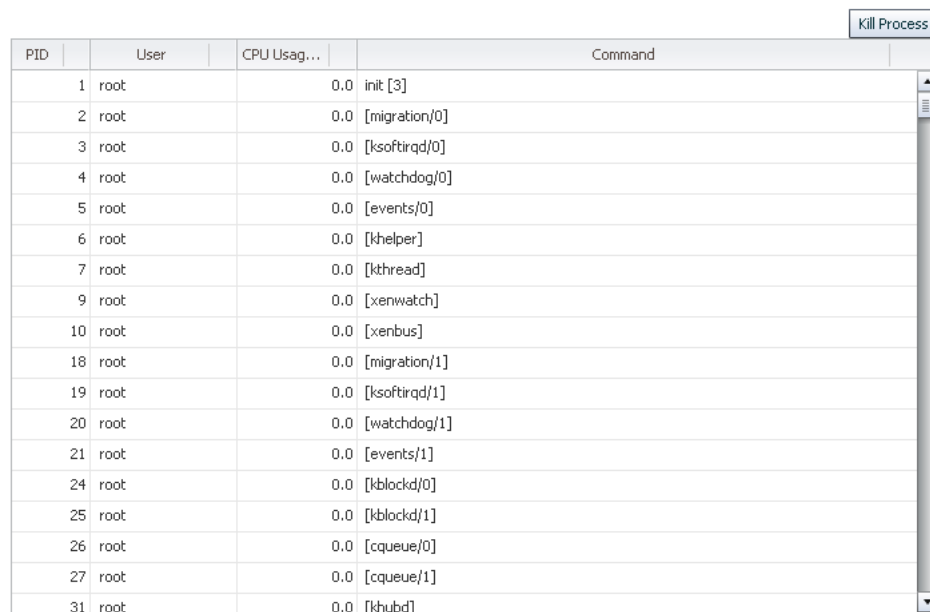
Example 8–35 maps the table to the `MetricDataService` by specifying the `metricName` and `metricColumns`. You do not have to specify the `headerText` attributes for the columns because it will be filled with the metric column labels. You can override these labels if required.

Example 8–35 Mapping a Table to the MetricDataService

```

<c:Table id="processesTable" width="100%" height="100%"
    metricName="CPUProcessesPerf"
    metricColumns="['ProcUser', 'ProcCPU', 'ProcCmd']"
    timePeriod="REALTIME"
    interval="30"
    >
    <c:adminElements>
        <mx:Button id="killProcessButton" label="Kill Process"
            click="controller.killProcess(event)"/>
    </c:adminElements>
    <c:columns>
        <mx:AdvancedDataGridColumn width="50" dataField="key" />
        <mx:AdvancedDataGridColumn width="100" dataField="ProcUser" />
        <mx:AdvancedDataGridColumn width="80" dataField="ProcCPU" />
        <mx:AdvancedDataGridColumn width="400" dataField="ProcCmd" />
    </c:columns>
</c:Table>

```

Figure 8–20 Data Service


PID	User	CPU Usag...	Command
1	root	0.0	init [3]
2	root	0.0	[migration/0]
3	root	0.0	[ksoftirqd/0]
4	root	0.0	[watchdog/0]
5	root	0.0	[events/0]
6	root	0.0	[khelper]
7	root	0.0	[kthread]
9	root	0.0	[xenwatch]
10	root	0.0	[xenbus]
18	root	0.0	[migration/1]
19	root	0.0	[ksoftirqd/1]
20	root	0.0	[watchdog/1]
21	root	0.0	[events/1]
24	root	0.0	[kblockd/0]
25	root	0.0	[kblockd/1]
26	root	0.0	[cqueue/0]
27	root	0.0	[cqueue/1]
31	root	0.0	[khubd]

8.17.2 Custom Data Provider

In [Example 8–36](#), the data for the table is loaded in the controller, and mapped to the page model `processInfoData` item. The `processInfoData` is an array of objects (of any type). The `dataField` property specified for each column identifies the public property that will be displayed in each column. In this case, the `dataField` name will also be used as the `headerText`. You can supply the `headerText` property to override this label.

Example 8–36 Mapping a Table to the processInfoData Item

```

<tbl:Table id="processInfoTable" dataProvider="{model.processInfoData}">
    <tbl:columns>
        <mx:AdvancedDataGridColumn width="100" dataField="Process ID"/>
    </tbl:columns>
</tbl:Table>

```

```

        <mx:AdvancedDataGridColumn width="250" dataField="User"/>
        <mx:AdvancedDataGridColumn width="100" dataField="Database"/>
        <mx:AdvancedDataGridColumn width="100" dataField="Status"/>
        <mx:AdvancedDataGridColumn width="250" dataField="Command"/>
        <mx:AdvancedDataGridColumn width="100" dataField="CPU Time"/>
        <mx:AdvancedDataGridColumn width="100" dataField="Memory Usage"/>
    </tbl:columns>
</tbl:Table>

```

8.17.3 Getting Selected Rows

The rows currently selected in the table can be obtained by looking at the `selectedRows` property of the table. This property is an array of selected rows, where each row is a Dictionary object that contains data in the row keyed by the column name. If the row is based on a custom data source, then the row will be whatever object was mapped into the table data source.

```
var process:Dictionary = page.processesTable.selectedRows[0];
```

If the table is set to allow single selection, then the `selectedRows` array includes a single entry only (or none if no row is selected).

8.18 Defining Dialogs

When you construct a dialog, typically you require an MXML class only, extending the `oracle.sysman.emx.intg.Dialog` class.

8.18.1 Dialog Registration

To make a dialog available to be displayed using the `invokeActivity` method, you must register it as an activity as part of the Integration class. In [Example 8–37](#), note the following:

- `id` attribute: The ID is used to reference this dialog from other activities within the application. It must be unique across all activities included in the application.
- `dialogClass` attribute: The `dialogClass` attribute is a reference to the MXML class that extends `Dialog` and that is the implementation for this dialog.

`inputParams` are optional, but they enable the dialog to be reused in situations where input parameters are required and you want to pass an object as context directly from the MXML using the bean directive. The MPCUI framework maps the input object parameters to the dialog parameters.

If you do not define `inputParams` as part of the dialog definition, then any input data required by the dialog (such as any custom properties) would have to be set in ActionScript and the dialog shown using the `Dialog.show` method.

Example 8–37 *Registering a Dialog*

```

<intg:DialogActivityDef id='metricHistory' label='Metric History'
dialogClass='{MetricHistoryDialog}' >
    <intg:inputParams>
        <intg:InputParam name='targetName' />
        <intg:InputParam name='targetType' />
        <intg:InputParam name='metric' />
        <intg:InputParam name='columns' />
        <intg:InputParam name='period' />
    </intg:inputParams>
</intg:DialogActivityDef>

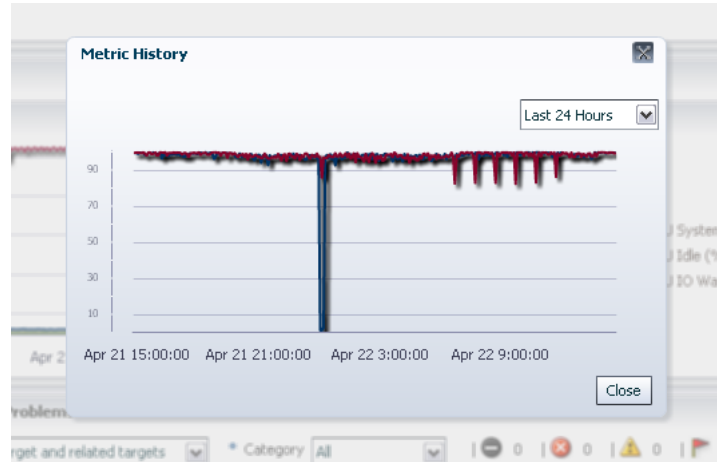
```

```

        <intg:InputParam name='title' />
    </intg:inputParams>
</intg:DialogActivityDef>

```

Figure 8–21 Metric History Dialog



8.18.2 Displaying a Dialog and Waiting for Close Events

If the dialog includes some state that is required when the dialog closes, then a close handler can be passed to the `invokeActivity` method. This handler is called with the `CloseEvent`. This handler identifies which button was pressed to close the dialog and retrieves the dialog object itself to retrieve information from it.

Example 8–38 Waiting for a Close Event

```

public function showCpuMetricDetails(event:MouseEvent):void
{
    var bean:Bean = new Bean("targetName", page.appModel.target.name,
        "targetType", page.appModel.target.type,
        "metric", page.cpuutil.metricName,
        "columns",page.cpuutil.metricColumns,
        "period", "REALTIME", "title", "Metric Details (Current)");
    page.invokeActivity("metricDetails", bean, metricDialogClosed);
}

public function metricDialogClosed(event:CloseEvent):void
{
    MpLog.debug("Metric Details Dialog Closed");
    var button:int = event.detail; // Alert.YES, Alert.NO, Alert.OK etc...
    var metDetailsDialog:MetricDetailsDialog = event.currentTarget
        as MetricDetailsDialog;
}

```

In [Example 8–38](#), the `metricDialogClosed` function is passed to `invokeActivity`. When the dialog closes, the method is called and passed a `CloseEvent`. From this event, the `currentTarget` property contains the dialog itself, and the `detail` property indicates which button was pressed to close the dialog.

8.19 Defining Trains

The train allows the definition of a multi-step UI, with next and previous buttons to navigate between each step. The train is typically used in cases where the user is going to create or modify an entity that has a large number of attributes that can be organized into categories.

The train must be registered with the integration class, and includes a controller that extends `TrainController` and a page class for each step in the train and extends `TrainStepPage`. Each step (train step page) can have its own controller class. Because each step is a page with a controller, the layout, management of data and response to events within the step is exactly the same as any other page in the application. For more information about the Page class, see [Section 8.6.2.1, "Page Class"](#).

The train step controller can access the train itself by referencing the `TrainStepPage.train` property. Use this to access other information maintained within the train object or its model.

8.19.1 Train Definition Example

[Example 8–39](#) provides a definition of train and [Figure 8–22](#) shows the train.

Example 8–39 Defining a Train

```
<intg:TrainActivityDef id='addNewUserEmbeddedTrain' label='Add New User'>
  <intg:stepActivities>
    <mx:Array>
      <intg:TrainStepActivityDef id='anuStep1' label='User Info'
pageClass='{trainSamp.S1_UserInfo}'
pageControllerClass='{trainSamp.AddNewUserTrainStepController}'/>
      <intg:TrainStepActivityDef id='anuStep2' label='Expiry'
pageClass='{trainSamp.S2_Expiry}'
pageControllerClass='{trainSamp.AddNewUserTrainStepController}'/>
      <intg:TrainStepActivityDef id='anuStep3' label='Credentials'
pageClass='{trainSamp.S3_Credentials}'
pageControllerClass='{trainSamp.AddNewUserTrainStepController}'/>
      <intg:TrainStepActivityDef id='anuStep4' label='Schedule'
pageClass='{trainSamp.S4_Schedule}'
pageControllerClass='{trainSamp.AddNewUserTrainStepController}'/>
      <intg:TrainStepActivityDef id='anuStep5' label='Notifications'
pageClass='{trainSamp.S5_Notifications}'
pageControllerClass='{trainSamp.NotificationsTrainStepController}'/>
      <intg:TrainStepActivityDef id='anuStep6' label='Confirmation'
pageClass='{trainSamp.S6_Confirm}'
pageControllerClass='{trainSamp.AddNewUserTrainStepController}'/>
    </mx:Array>
  </intg:stepActivities>
</intg:TrainActivityDef>
```

Figure 8–22 Train Example

The screenshot shows a web interface for defining a 'Train' process. At the top, a progress bar indicates six steps: User Info, Expiry, Credentials, Schedule, Notifications, and Confirmation. The 'Expiry' step is currently selected and highlighted. Below the progress bar, the 'Expiry' section contains the text 'When do you want the account to expire:'. To the right of this text is a calendar widget for April 2011. The calendar shows the days of the week (S, M, T, W, T, F, S) and the dates. The 22nd of April is highlighted. Below the calendar, there is a label '# of days after expiration that the account is disabled:' followed by a slider control. The slider has a range from 0 to 7, with the current value set to 0.

8.19.2 Train Controller

The train controller is used to managed state kept across all pages in the train and respond to changes in the train (movement between steps) and respond to the train completing when the user clicks either the Finish or Cancel button.

8.19.3 Train State

State may be maintained in the Train model using the `Train.model` property. This property is a dynamic property that can be used to hold any information appropriate to the train. Individual pages can store their own state in their own model properties and may also access information stored in the train model.

8.19.4 Train Events

Each train step controller can implement the `init` and `destroy` methods that are called when the step starts or stops. The step can do either of the following:

- Perform a step-specific processing step
- Access the train and allow it to process higher level logic

The train controller can also be called when the train ends (Finish or Cancel) by adding a listener function for the train done event:

Example 8–40 Adding a Listener Function

```
// register a listener for the train complete event, this may be a cancel or
finish.
train.addEventListener(TrainEvent.TRAIN_EVENT, trainDone);
```

The listener (`trainDone` in [Example 8–40](#)) can inspect the train state and determine if processing should continue or not. It can choose to direct control to some other activity (page) or can set the train back to another step:

Example 8–41 Defining Actions at the End of the Train

```
public function trainDone(event:TrainEvent):void
```



```

{
  // train cancel/finish button was pressed, so caller can now validate
  // the train (look at the model). The caller has the various options indicated
  below.
  var train:Train = event.target as Train;

  if(train.model["isComplete"])
  {
    // want to end the train, but go somewhere else (otherPage is a page id)
    train.endTrain("otherPage");
  }
  else
  {
    // go back to train at a certain step
    train.controller.setStepById("step2");
  }
}

```

8.20 Defining Information Item and Information Displays (Label-Value Pairs)

The `InfoDisplay` and `InfoItem` classes allow you to display a set of label-value pairs in a group with the labels right-aligned and the values left-aligned. Each entry (`InfoItem`) in the display specifies a label, value, optional icon, destination, or click property.

The destination or click properties cause the value to appear as a link. You can set destination to either of the following:

- String that is the identifier for some other activity (page or dialog)
- URL object constructed in the controller (see `HomePage.mxml` and `HomePageController` for examples)

You can specify the click handler instead of the destination and set it to a function within the controller that will be called when the item is clicked by the user.

Example 8-42 Defining Label Value Pairs

```

<components:InfoDisplay width="100%" height="100%" >
  <!-- ref to SQLDataService -->
  <components:InfoItem label="CPU Model"
    value="{ids.result.getString(0,'CPU Model')}}" />
  <!-- ref to MetricDataService for metric with a key -->
  <components:InfoItem label="CPU(0) Idle %"
    value="{procData.result.getString('0','CPUIdle')}}" />
  <!-- ref to MetricDataService for metric with no key -->
  <components:InfoItem label="Current Load"
    value="{respData.result.getString('', 'Load')}}" />
  <!-- ref to page model; model set in controller in SQL svc handler -->
  <components:InfoItem label="{model.osVersLabel}" value="{model.osVersion}" />
  <components:InfoItem label="Hosted By" value="{model.relatedHost}"
    destination="{model.relatedHostLink}" />
</components:InfoDisplay>

```

Figure 8–23 Label Value Pairs

Status	
Current Status	Up
Up Since	Apr 19 13:06:19
Availability %	100.0%

8.21 Using Built-in Renderers

In addition to the ability to define custom renderers for table columns, headers, and other UI elements using the capabilities provided by the Flex framework, the MPCUI framework also provides several built-in renderers that can be used to display custom icons in a table or for an `InfoItem`. These renderers are also useful for meta-data only implementations where it is not possible to create controller code to implement a renderer.

These built-in renderers show an icon in place of a text value, either in a `Table` or `InfoItem` component. The renderer is specified by using the `"appModel.renderer"` directive in the MXML and specifies a renderer id to select the renderer and then a set of input parameters for the renderer depending on the renderer type.

For example, the following code would result in an icon being displayed next to the value on the `InfoItem` and will show a check mark, warning, or error icon depending on the value displayed in the `InfoItem`:

```
<mp:InfoItem id="currentLoad" label="CPU Load"
  value="{respData.result.getString('','Load'))}"
  imageRenderer="{appModel.renderer('CHECK_MARK',
    bean('type','number','warning','0.1','critical','0.4'))}" />
```

The first parameter, `"CHECK_MARK"` indicates which renderer to be used, see the complete list below. The second parameter, `"bean"` specifies the input parameter for the check mark renderer. This parameter will be different and in some cases optional depending on the renderer selected. Refer to the API documentation for details regarding what each renderer requires for input parameters.

The built-in renderers supported in 12.1.0.3.0 of the EDK include the following:

- **CHECK_MARK**
Displays a check mark, warning icon or error icon depending on the value provided. The renderer can either be used to display a check mark or error icon in the case where a Boolean value is shown. The Boolean may be true/false, t/f or 0/1. If the `'type'` parameter is specified as `'number'`, then the value will be compared to thresholds provided in the input parameters to also show a different icon if the following is beyond the specified threshold.
- **TARGET_TYPE**
Displays the icon associated with a target type value. This is the internal target type id, e.g. `'oracle_database'`, not the actual displayed string representation of the target type.
- **TARGET_STATUS**

Displays the icon associated with a status value. The status value will be up/down, true/false, or 0/1 and will display and up or down arrow according to that value.

When the renderer is associated with an `InfoItem`, by default the value shown in the `InfoItem` will be passed to the renderer to determine which icon should be displayed. In cases where an alternative value should be used to control the renderer, the `InfoItem` provides the "imageDataSource" property. This property can be bound to a data item that is different than the displayed value.

In the following example, the `TARGET_TYPE` renderer is used to show the target type icon next to the name of the target. In order to do this the `imageDataSource` property is set to an item that will contain the internal target type needed by the renderer.

```
<mp:InfoItem id="relatedHost" label="Hosted By" value="{model.relatedHost}"
  imageRenderer="{appModel.renderer('TARGET_TYPE')}"
  imageDataSource="{model.relatedHostType}"
  destination="{model.relatedHostLink}" />
```

8.22 Defining Links

Use the link component to display what appears to the user to be a link to a URL. The link specifies a label property and also either a destination or click handler property. The destination can be an activity id or a URL object constructed in the controller. For information about the `InfoItem` class, see [Section 8.20, "Defining Information Item and Information Displays \(Label-Value Pairs\)"](#).

Figure 8–24 Link Example

Hosted By .com

8.23 Including Enterprise Manager Images

To reference one of the images shipped with the Enterprise Manager product from the MXML in both Flex and metadata implementations, use the `appModel.emImage` function to refer to the desired image. Note that the list of images and their filenames is not currently part of the Enterprise Manager EDK and therefore is subject to change.

You should verify and test any use of this information with each new release of Enterprise Manager. In a typical deployment, the images are located under the `emcore_war/images` directory. For example:

```
<mx:Image source="{appModel.emImage('yellowFlag.gif')}" />
```

8.24 Displaying a Processing Cursor

The UI displays the processing or busy cursor automatically when:

- Any new activity is accessed (page, train or dialog)
- Any request is made to the Management Server for data

Typically, you do not have to show the busy cursor. However, if you feel that you must show the busy cursor, take care that the cursor is ended cleanly. Ensure that if exceptions are thrown while the busy cursor is shown, that they are caught and the cursor is removed.

To show the cursor, call the `MpCursor.setBusyCursor` method.

To remove the cursor, call `MpCursor.removeBusyCursor`. Both methods accept an optional `owner` parameter. This parameter allows you to nest multiple cursors calls.

8.25 Defining a Processing Window

The processing window displays a dialog that is displayed during long running tasks. It can be updated periodically with status information as the task runs. You can show the dialog with:

- an infinite completion, which shows a spinning status icon
- a finite completion which shows a percentage bar that can be updated as the task executes (from 0-100% complete)

Example 8–43 *Defining a Processing Window*

```
private var proc:ProcessingWindow;
private var indeterminate:Boolean = false;

public function showDialog():void
{
    proc = new ProcessingWindow();
    proc.title="Processing";
    proc.operationText = "Get Metric Info Many Times...";
    proc.showDetailText = true;
    indeterminate = !indeterminate;
    proc.indeterminate = indeterminate;
    proc.show();

    taskCount = 0;
    longRunningTask();
}

private var taskCount:int = 0;
private function longRunningTask():void
{
    proc.addDetailText("Starting next item of work...");
    var querySvc:SqlQuery = new SqlQuery(allMetricsHandler);
    var queryRequest:RunQueryRequest = new RunQueryRequest();
    var tgt:Target = ApplicationContext.getTargetContext();
    queryRequest.addSqlQueryInput(new SqlQueryInput("GET_ALL_METRIC_INFO", "GET_
METRIC_INFO", [
        ["TARGET_TYPE", tgt.type],
        ["TYPE_META_VER", tgt.typeMetaVer],
        ["CAT_PROP_1", tgt.catProperties[0]],
        ["CAT_PROP_2", tgt.catProperties[1]],
        ["CAT_PROP_3", tgt.catProperties[2]],
        ["CAT_PROP_4", tgt.catProperties[3]],
        ["CAT_PROP_5", tgt.catProperties[4]],
        ["TARGET_GUID", tgt.guid]
    ], SqlQueryInput.GENERIC_QUERY_TYPE));
    querySvc.runQuery(queryRequest);
}

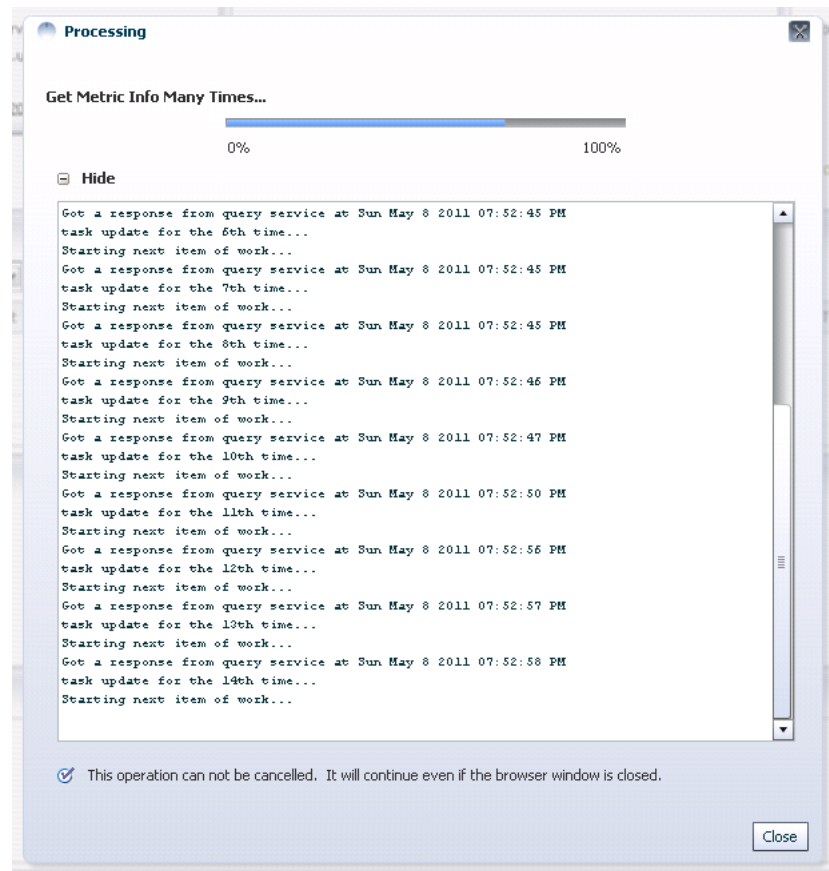
public function allMetricsHandler(event:RunQueryResultEvent):void
{
    proc.addDetailText("Got a response from query service at "+new
Date().toLocaleString());
    taskCount++;
    proc.addDetailText("task update for the "+taskCount+"th time...");
}
```

```

if(taskCount < 20)
{
    if(!indeterminate) proc.setPercentComplete(taskCount*5);
    longRunningTask();
}
else
{
    proc.addDetailText("All work complete.");
    proc.setComplete();
}
}
...

```

Figure 8–25 Processing Window Example



8.26 Defining Icons for Target Types

You can specify icons to associate with a target type to be displayed in the Cloud Control console wherever a target type icon is shown (such as next to the target menu).

MPCUI supports the following graphic formats for icons:

- PNG
- JPG
- GIF

Oracle recommends the following sizing for icons:

- Small icon: 16x16
- Large icon: 24x24

Example 8–44 Defining Icons

```
<EmuiConfig>
  <large-icon>demo_hs_large_icon.png</large-icon>
  <small-icon>demo_hs_small_icon.png</small-icon>
</EmuiConfig>
```

Figure 8–26 and Figure 8–27 provide examples of a small and a large icon.

Figure 8–26 Small Icon



Figure 8–27 Large Icon



8.27 Displaying the Target Navigator

The target navigator can be displayed on the left side of the home page of any composite target or any of its members. The target navigator displays the composite target at its root and then shows all members of the composite by searching for any *contains* associations below it. Targets that are associated with the composite target can have other *non-contains* associations with the composite target or with other targets. However, only those targets with *contains* associations with the composite target are shown in the target navigator. You can add these containment associations through any of the supported mechanisms for discovering or deriving associations. For more information, see [Chapter 10](#).

To enable the target navigator, the MPCUI metadata must include the `<EmuiConfig>` element with the `context-pane-visible` property set to true. This must be set for the composite target type as well as any of its member targets. If it is not set for member targets, then the navigator will not appear showing the other members of the composite when the home page is displayed for those targets.

By default, the `context-pane-visible` property is set to false and the target navigator is not displayed.

Note: If there are no *contains* associations, then the target navigator will not appear, even if the `context-pane-visible` property is set to true.

Example 8–45 Enabling the Target Navigator

```
<?xml version = '1.0' encoding = 'UTF-8'?>

<CustomUI target_type="demo_
hostsystem"xmlns="http://www.oracle.com/EnterpriseGridControl/MpCui">

  <EmuiConfig>
    <context-pane-visible>true</context-pane-visible>
  </EmuiConfig>
```

</CustomUI>

8.28 Defining a UI for Guided Discovery

The MPCUI framework supports the ability to define a custom user interface that can be registered as part of a guided discovery flow. After registration, this discovery flow is available from the **Add Targets Manually** page. For more information about adding targets manually, see [Section 11.6](#).

8.28.1 About Guided Discovery

The guided discovery flow provides you with the ability to add targets and associations to Enterprise Manager by running discovery scripts on selected Management Agents and calling service APIs to add the appropriate entities. This process is driven from a user interface wizard (train) and can use information supplied by the end user to guide the process. It is up to you to determine (based on your specific requirements) the information required from the end user during this process. For examples, see the plug-in samples in the EDK (demo_hostsample and demo_hostsystm). For more information about discovery scripts, see [Chapter 11](#).

The services typically used during guided discovery include the following:

- TargetInfo services to retrieve Management Agents and targets, for example, for target existence or target properties
- AssociationInfo services to retrieve existing associations
- Discovery service to run discovery scripts on selected Management Agents
- TargetManagement services to create or delete targets
- AssociationManagement services to create or delete associations

For more information about these services, see [Section 8.28.5, "Using Discovery Service"](#), [Section 8.28.6, "Using Target Information Services"](#), and [Section 8.28.7, "Using Target Management Services"](#).

8.28.2 Supporting Guided Discovery

To add guided discovery to a plug-in, ensure that the following directories contain the required files:

- *plugin_stage/discovery*
 - Scripts that will be executed from the guided discovery flow. These scripts might include multiple targets and associations. For more information about discovery scripts, see [Section 11.3](#).
 - customdiscover.lst file. This file must include one line for each discovery script to be provided with the plug-in. Each entry must reference a discovery category, which is a unique identifier that will be used to identify the script to be executed when calling the discovery service. The following entry shows a discovery category (DHS_DISC) that is used to refer to the demo_hostsample_discovery.pl script during the guided discovery flow.

```
DHS_DISC|demo_hostsample_discovery.pl
```

- *plugin_stage/oms/metadata/discovery*

Discovery metadata file (*plugin_discovery.xml*). For more information about the discovery metadata file and an example of the discovery XML, see [Section 11.2](#). For guided discovery, there are a number of attributes that must be specified correctly to allow your guided discovery to be registered correctly.

- `<DiscoveryModule name="DemoHostSample">`

This is the unique name for the discovery module and must match the module name used to register the discovery SWF in the MPCUI metadata file.

- `<NlsValue>Discover Demo Host Sample Targets</NlsValue>`

This is the label that appears in the Target Types list on the **Add Targets Manually** page of the Cloud Control console.

```
<CustomDiscoveryUI>
  <LaunchADF>
    <DestOutcome>goto_core-mpcustomdiscovery-page</DestOutcome>
  </LaunchADF>
</CustomDiscoveryUI>
```

This must be exactly the same in your discovery metadata file. It ensures that the guided discovery UI that you built and included in your plug-in will be launched.

- *plugin_stage/oms/metadata/mpcui*

- *discovery.swf*

Similar to the MPCUI application created for a target home page, the guided discovery UI is constructed as a Flex application.

- *MyMpcui.swf*

In addition to the discovery SWF, the MPCUI metadata file must include an entry to register the discovery SWF with the appropriate discovery module.

Example 8–46 SwfFiles Tag From MPCUI Metadata File for Flex-based UI

```
<SwfFiles>
  <Swf is_homepage="true">HostSample.swf</Swf>
  <Swf discovery_module="DemoHostSample">HostSampleDiscovery.swf</Swf>
</SwfFiles>
```

The `<Swf discovery_module="DemoHostSample">HostSampleDiscovery.swf</Swf>` entry shows a discovery SWF being registered with the discovery module that was registered in the discovery metadata. This UI is launched when this discovery module is selected by the user in the **Add Targets Manually** page in the Cloud Control console.

8.28.3 Constructing the Guided Discovery User Interface

The guided discovery UI is built using the MPCUI features for constructing a Flex UI. The UI components, such as regions, buttons, tables, dialogs, and so on are used to construct a user interface to guide the user through the process of adding new targets to Enterprise Manager. For information about adding these UI components, see the relevant sections of this chapter, such as [Section 8.17, "Defining Tables"](#) or [Section 8.18, "Defining Dialogs"](#).

8.28.3.1 Discovery Application and Integration Class

The discovery application must extend the `MpDiscoveryApplication` and *not* the `MpApplication` class. The discovery application is an MXML file that specifies the name of the application (taken from the name of the file) and the integration class that defines the activities included in the discovery application:

Example 8–47 Application MXML

```
<?xml version="1.0" encoding="utf-8"?>
<mp:MpDiscoveryApplication xmlns:mx="http://www.adobe.com/2006/mxml"
                           xmlns:mp="http://www.oracle.com/mpcui"
    backgroundColor="#EFF3F7" preloader="oracle.sysman.emx.MpPreloader" >
  <mx:Script>
    <![CDATA[
      import discovery.DiscoveryInteg;
      override public function getIntegrationClass():Class
      {
        return discovery.DiscoveryInteg;
      }
    ]]>
  </mx:Script>
</mp:MpDiscoveryApplication>
```

The integration class for the discovery application defines the set of activities used by the discovery UI. The discovery application must include at least one `PageActivity` that is defined with the `isDefaultPage=true` property indicating that this is the page that will be loaded when the guided discovery starts. In [Example 8–48](#) (extracted from the `demo_hostsystem` sample plug-in), take note of the `discoHomePg` activity.

Example 8–48 Integration Class

```
<mp:PageActivityDef id='discoHomePg' label='Discovery Console'
  pageClass='{discovery.DiscoveryTrainPage}'
  pageControllerClass='{discovery.DiscoveryTrainPageController}'
  isDefaultPage="true" />
<mp:TrainActivityDef id='discoTrain' label='Discover New Targets'>
  <mp:stepActivities>
    <mx:Array>
      <mp:TrainStepActivityDef id='selAgentsStep' shortLabel="Select Agents"
        label='Add Demo Host System Target: Select Agents'
        pageClass='{discovery.train.SelectAgentsStep}'
        pageControllerClass='{discovery.train.DiscoveryTrainStepController}' />
      <mp:TrainStepActivityDef id='selHostSysStep' shortLabel="Select
        System" label='Add Demo Host System Target: Select Demo Host System'
        pageClass='{discovery.train.SelectHostSystemStep}'
        pageControllerClass='{discovery.train.DiscoveryTrainStepController}' />
      <mp:TrainStepActivityDef id='selTargetsStep' shortLabel="Configure
        Targets" label='Add Demo Host System Target: Configure Targets'
        pageClass='{discovery.train.SelectTargetsStep}'
        pageControllerClass='{discovery.train.DiscoveryTrainStepController}' />
      <mp:TrainStepActivityDef id='confTargetsStep' shortLabel="Confirm
        Changes" label='Add Demo Host System Target: Confirm Changes'
        pageClass='{discovery.train.ConfirmChangesStep}'
        pageControllerClass='{discovery.train.DiscoveryTrainStepController}' />
      <mp:TrainStepActivityDef id='finTopo' shortLabel="Summary" label='Add
        Demo Host System Targets: Apply Changes'
        pageClass='{discovery.train.FinalizeTopoStep}'
        pageControllerClass='{discovery.train.DiscoveryTrainStepController}' />
    </mx:Array>
  </mp:stepActivities>
</mp:TrainActivityDef>
```

8.28.4 Structure of the Discovery Application

The discovery application is often a single page that either has a train embedded in it, or that displays dialogs to obtain information from the end user to guide the discovery process. The steps of the guided discovery flow depends on the requirements, but often involve the following:

1. Determine on which Management Agents to run a discovery script
2. Run the discovery script
3. Process the results of the discovery script, adding additional information provided by the end user
4. Call APIs to add or delete targets

One important consideration about guided discovery is that it can be used to update the topology of existing composite targets as well as discover new targets. In the case of the sample plug-in (`demo_hostsystem`), the guided discovery UI allows the user to add new system targets, but also allows the user to add or remove members from an existing system.

This use case also illustrates the requirement to use Enterprise Manager APIs to query for the set of existing targets known to Enterprise Manager to compare the set with information returned from the discovery script to identify which targets are already managed by Enterprise Manager and which are not. For example, the result might be a list of new targets that should be added and a list of other targets that no longer exist in the managed configuration and must be removed from Enterprise Manager.

This scenario also illustrates that the discovery application might also be integrated with the custom UI built for the target home page. This provides the user with the ability to update the configuration of an existing composite target directly from the composite target home page.

See the **HostSystemConfiguration** page in the `demo_hostsystem` sample plug-in for an example of using discovery UI from within a target home page.

8.28.5 Using Discovery Service

The Discovery service is used to run a discovery script included with your plug-in. For a description of your plug-in requirements to support discovery, see [Section 8.28.2, "Supporting Guided Discovery"](#).

Example 8–49 (included in the `demo_hostsystem` sample plug-in in the `DiscoveryTrainStepController`) shows calling the discovery service (`TargetFactory.discoverTargets`). This service includes an `addRequest` method that can be called multiple times to process discovery on multiple Management Agents if required.

Each call to the `addRequest` method is passed the following along with the handler that will be called with the results of the discovery script:

- Request ID
A unique identifier (assigned by you) associated with that particular request which will enable you to retrieve the specific results associated with that request.
- Agent
the Management Agent Target that the discovery should be run against

- Plug-in ID

The plug-in ID associated with the discovery to be run. A plug-in can include multiple discovery modules and categories.

- Discovery category

The discovery category. This must map to a discovery script through an entry in the `discover.lst` file included in the agent part of the plug-in.

Example 8–49 Discovery Service

```
/**
 * when doing discovery, the service will accept multiple requests to be
 * processed at the same time. this would typically be the case if multiple
 * agents were involved in the process, but could also be if different discovery
 * categories (scripts) were to be processed.
 *
 * the discovery request includes the following elements:
 *   requestId  a unique identifier associated with that particular request
 *               that will allow you to retrieve the specific results associated
 *               with that request.
 *   agent      the agent Target that the discovery should be run against
 *   pluginId   the pluginId associated with the discovery to be run; a plug-in
 *               can include multiple discovery modules and categories
 *   discCat    the discovery category; this must map to a discovery script via
 *               an entry in the discover.lst file included in the agent part of
 *               the plugin
 *   params     parameters that are to be passed to the discovery script
 *
 * Note on discoveryModule - in addition to the pluginId, the discovery UI is
 * passed the discovery module associated with this discovery pass. If you've
 * chosen to implement multiple types of discovery operations from a single
 * discovery UI you may retrieve the discoveryModule to determine in what context
 * the UI was launched.
 */

var requestId:String = "DiscReq1";
var pluginId:String = ApplicationContext.getPluginId();
var discoveryCategory:String = "DHSYSTEM_DISC";
discSvc.addRequest(requestId, agent, pluginId, discoveryCategory, params);

TargetFactory.discoverTargets(discSvc, discoverResultsHandler);
```

The discovery results handler, declared as follows, is passed a fault object and the discovery results.

```
public function discoverResultsHandler(disc:Discovery, fault:ServiceFault):void
```

If a fault did not occur during processing of the discovery script, then the fault object will be null. The discovery object passed to the handler includes an Array of the `DiscoveryRequest` objects constructed by calling the `addRequest` method. Each request includes the properties specified (such as agent, category, and so on) and also includes a `DiscoveredTargets` object. The `DiscoveredTargets` object includes the list of targets returned from the discovery script that was run on the target Management Agent for the specified request. For more information about discovery scripts, see [Section 11.3](#) and for information about the discovery objects returned by the `DiscoveryService`, see the API documentation in the EDK.

8.28.6 Using Target Information Services

During the discovery process it is often necessary to obtain target information such as a list of agents or the set of targets of a particular type. The target information service provides a number of APIs that can be used for such purposes. This section provides an overview of these services. For additional information, see the API documentation in the EDK and for examples of their use, see the `demo_hostsystem` sample plug-in.

- `TargetFactory.getAgents`

The `getAgents` API enables you to retrieve a set of Management Agents that can be used to perform discovery. You can filter the list by specifying selection properties (Array of `TargetProperty`) such as selecting all the Management Agents running on a Windows host.

- `TargetFactory.getTargets`

Use the `getTargets` API to retrieve a list of targets specifying any number of selection criteria including hosts, target types, managed status, or metadata version. Each item is specified as a list of possible values and the request can include one or more selection criteria.

- `Target.getSystemMembers`

Use the `getSystemMembers` API to retrieve the list of system member targets. These are targets that are included in the system through the `systemStencil`. For information about the system targets, see the *Oracle Enterprise Manager Cloud Control Extensibility Programmer's Guide*:

<http://www.oracle.com/pls/em121/homepage>

- `Target.getCompositeMembers`

Use the `getCompositeMembers` API to retrieve the list of composite member targets. Composite members are those included in a composite (or system target) through containment associations. For information about composite targets, see the *Oracle Enterprise Manager Cloud Control Extensibility Programmer's Guide*:

<http://www.oracle.com/pls/em121/homepage>

8.28.7 Using Target Management Services

The target management services provide you with the ability to create or delete targets or associations. In the case of target management, associations can also be passed as part of the target definition and the associations are added as part of the process of adding the target itself. This section provides an overview of these services. For additional information, see the API documentation in the EDK and for examples of their use, see the `demo_hostsystem` sample plug-in.

- `TargetFactory.createTargets`

Use the `createTargets` API add targets to Enterprise Manager. The process of adding targets to Enterprise Manager forces the deployment of the necessary plug-in to the Management Agents associated with each target. The request to this API is a list of `Target` objects, each of which must, at a minimum, specify a name, type, and agent. Typically, target instance properties (if used for this target type) can also be specified.

- `TargetFactory.deleteTargets`

Use the `deleteTargets` API to remove targets from Enterprise Manager. The request to this API is a list of `Target` objects. Removing a target from Enterprise Manager

should be done with care as deleting the target is not reversible and it removes all target, metric, and configuration history.

- **Target.createAssociations**

Use the createAssociations API to add associations between the specified target and another target. Associations can be created in this way when creating them by using derived associations or by using the system stencil. In all cases, the association must be associated with a corresponding allowed pairs definition. For more information, see [Chapter 10](#).

- **Target.deleteAssociations**

Use the deleteAssociations API to delete associations between the specified target and other targets.

8.29 About Logging

The following sections discuss the logging options for MPCUI.

8.29.1 Add Logging to your Code

Use the logging facility (MpLog) to log messages from your code.

Note: Do not use the Flash trace global method as this only provides output to the Flash log where the browser is running a debug version of the Adobe Flash Player.

While logging can be useful in situations where diagnostics are necessary, it has a cost in terms of code size and overhead. Therefore, use logging with care.

Perform logging by calling one of several MpLog methods (such as debug, info, error, or fatal). The methods accept a message string and an optional list of parameters that must be substituted.

To substitute parameters, indicate the parameter location using {#} format:

```
MpLog.debug("The metric {1} was not found for the target {2}.", metricName,
targetName);
```

The message generated for this log statement appears in the following log output:

```
2011-04-22 11:10:17 [MpCui] DEBUG The metric CPU was not found for the target
MYHOST
```

The level (info, debug, error, fatal) allows the user to enable log output for different classes of messages.

- By default, all error and fatal messages are sent to the log output location.
- The info and debug level messages are only sent if these levels are explicitly enabled.

Furthermore, you can direct the messages for each level to different output locations. There are three possible log locations:

- **FLASHLOG:** messages are sent to the Flash log (requires a debug Adobe Flash Player)
- **EMLOG:** messages are sent to the Enterprise Manager application logs

- **CONSOLE:** messages are sent to a small window displayed at the bottom of the MPCUI display

8.29.2 Options for Capturing Log Output

The options for capturing log output depend on your implementation:

- [Running MPCUI From Adobe Flash Builder](#)
- [Running MPCUI from Enterprise Manager Console](#)

8.29.2.1 Running MPCUI From Adobe Flash Builder

When you are developing MPCUI using Adobe Flash or Flex Builder, the log messages sent to FLASHLOG appear in the console window at the bottom of the Adobe Flash Builder integrated development environment (IDE) by default.

To change these logging settings:

1. Open the `html-template/data/mpCuiProperties.xml` file.
2. Locate the `loglevel` element:

```
<!-- Logging
      level: DEBUG, INFO, ERROR, FATAL, WARN (or ALL)
      output location: FLASHLOG, CONSOLE, EMLOG
      format: level,output;level,output (e.g. DEBUG,FLASHLOG;ERROR,EMLOG)
-->
<loglevel>ALL,FLASHLOG</loglevel>
```

3. Modify the `loglevel` element as required.
4. Rebuild your Adobe Flex application before running the application.

8.29.2.2 Running MPCUI from Enterprise Manager Console

After the plug-in is deployed, the settings for logging are detected from the HTTP request. The default setting is `FATAL,FLASHLOG;ERROR,FLASHLOG`.

The end user can modify the settings as follows:

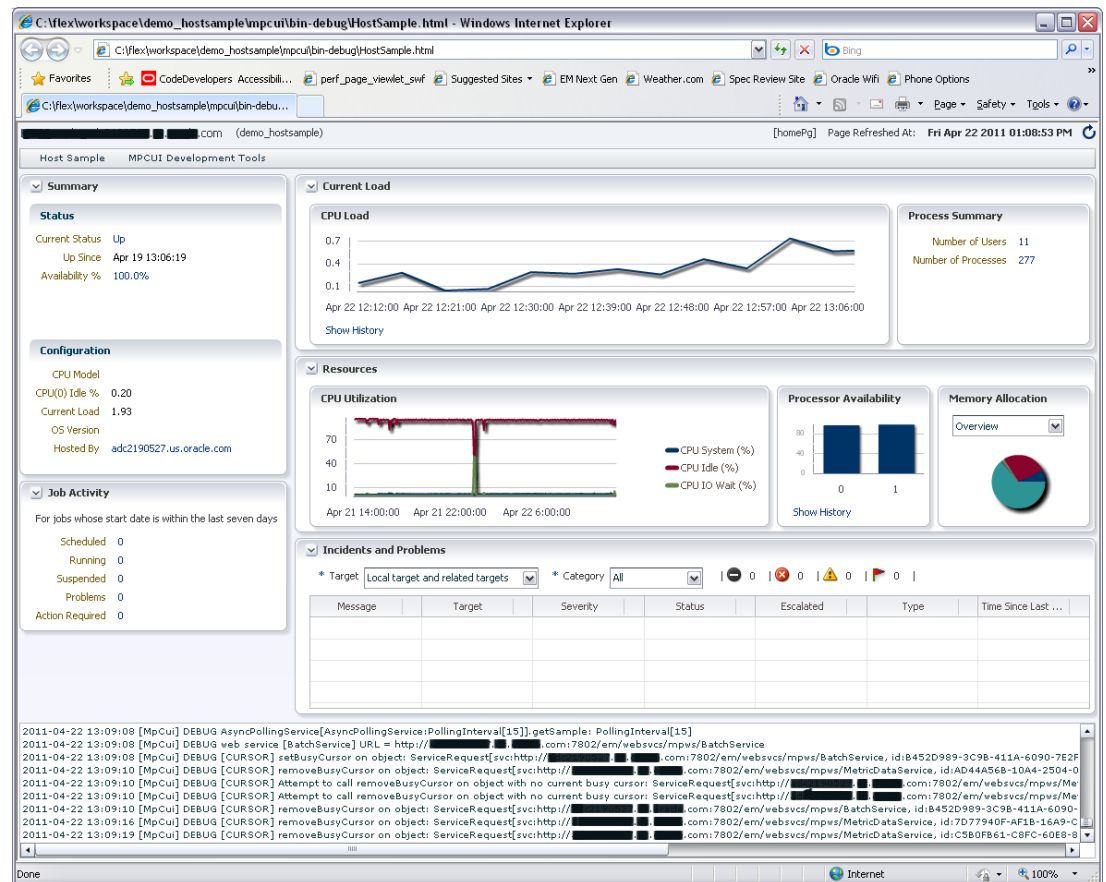
1. Append the following to the URL in the address bar of the browser:

```
&loglevel=ALL,FLASHLOG
```

2. You can substitute `ALL,FLASHLOG` with any valid logging settings, such as `ERROR,CONSOLE` and so on.
3. For diagnostic situations, add the following to the end of the URL:

```
&loglevel=ALL,CONSOLE
```

4. Refresh the page to see the page and all log messages in the console window similar to [Figure 8-28](#)

Figure 8–28 Viewing Log Messages

8.30 Development Environment Options

When building a custom UI for your plug-in, you have the following development environment options:

- **Metadata Only**

Use your preferred text editor to edit your MPCUI metadata file. For more information about the MPCUI metadata file, see [Section 8.4.1, "Overview of MPCUI Metadata Elements"](#).

- **Flex SDK and Apache Ant**

Download the Flex SDK and utilize the Apache Ant build files supplied with the EDK to build your SWF code. This option does not require a license for Flex Builder but it requires the use of the command line debugger (fdb) as opposed to the graphical debugger that is part of Flex Builder. For more information about using Flex SDK and Apache Ant, see [Section 8.30.1, "Developing MPCUI with Flex SDK and Apache Ant"](#).

- **Flex Builder**

You must acquire a license for Flex Builder to build and debug your Flex code included with your plug-in. For more information about using Flex Builder, see [Section 8.30.2, "Developing MPCUI in Adobe Flash or Flex Builder"](#).

Note: The project settings used in Adobe Flash Builder for your MPCUI project are critical to ensure that your application is built correctly and operates correctly when deployed as part of your plug-in. Oracle recommends that you use the settings of the demo_hostsample sample project as a guide and match these settings exactly.

In Adobe Flash Builder, select the project properties (from the **File** menu, select **Properties**), and ensure that the settings match those of the sample project. Ensure that the settings on the Flex Build Path properties on the Library path tab are the same as those on the Flex Compiler properties.

8.30.1 Developing MPCUI with Flex SDK and Apache Ant

If Flex Builder is not available, then install Adobe Flex SDK and use Ant to build the SWF file for your plug-in as follows:

Note: These steps build HostSample.swf using Ant on Windows.

Two Apache Ant build files are included for reference in case you want to build the Demo Sample from the command line or without using the Flex Builder IDE:

- demo_hostsample\mpcui\build.xml
 - demo_hostsample\mpcui\plugin-build-config.xml
-

1. Download and install Apache Ant from the following website:

<http://ant.apache.org>

- a. Set the ANT_HOME environment variable to the location where Apache Ant is installed.
 - b. Include \$ANT_HOME\bin in the PATH environment variable of your command shell.
 - c. Set the ANT_OPTS environment variable to **-Xmx512m**.

2. Download and install Flex SDK version 3.5:

- Download the Flex SDK 3.5 ZIP file from the following website:

http://download.macromedia.com/pub/flex/sdk/flex_sdk_3.5.zip

- Download the Flex 3.5 Data Virtualization Components for Flex Builder ZIP file from the following website:

http://download.macromedia.com/pub/flex/sdk/datavisualization_sdk3.5.zip

- Download the Flex 3.5 Automation Libraries for Flex Builder ZIP file from the following website:

http://download.macromedia.com/pub/flex/sdk/automation_sdk3.5.zip

Note: When you are adding the automation libraries to the SDK directory, they must be expanded into the FLEX_HOME/frameworks directory, and *not* in the FLEX_HOME directory.

3. Edit the demo_hostsample\mpcui\build.xml file.

Update the FLEX_HOME property so location points to the location of the Flex SDK installation.

4. Build the HostSample.swf file:

- a. `cd demo_hostsample\mpcui`

- b. **ant**

Entering this command builds the demo_hostsample\mpcui\bin-debug\HostSample.swf file.

8.30.2 Developing MPCUI in Adobe Flash or Flex Builder

Note: This section assumes the use of Adobe Flash Builder 4. However, these steps should be the same if you are using Adobe Flex Builder 3.

This section describes the process to follow when building a Flex application using the MPCUI libraries and Adobe Flash Builder. These steps assume the use of the sample application provided with the EDK referred to as the Demo Host Sample (or demo_hostsample). As with many development activities it is often easiest to start with a working example to understand how the provided APIs work and how to use them to accomplish higher-level use cases.

To simplify the process for developing your custom UI, build and run the Flex application from Adobe Flash Builder without having to redeploy the plug-in to Enterprise Manager after each change. When running from Adobe Flash Builder, your UI will not have access to the other Enterprise Manager features and pages available to the console, but you will be able to exercise your UI to ensure that it is functioning correctly before deploying it as part of your plug-in.

8.30.2.1 Configuring Adobe Flash Builder

The MPCUI libraries provided with the EDK are built with and require the use of the Flex SDK 3.5. Therefore when building your MPCUI project in Adobe Flash Builder you must use the 3.5 SDK.

To check the version:

1. From within Adobe Flash Builder, from the **Window** menu, select **Preferences**.
2. In the Preferences dialog, expand **Flash Builder**, then select **Installed Flex SDKs**.

Note: In Adobe Flex Builder 3, from the Preferences dialog, expand **Flex**.

If you do not have Flex SDK 3.5, then you must install Flex SDK 3.5:

1. Download the following Flex SDK 3.5 ZIP files:

- Download the Flex SDK 3.5 ZIP file from the following website:
http://download.macromedia.com/pub/flex/sdk/flex_sdk_3.5.zip
 - Download the Flex 3.5 Data Virtualization Components for Flex Builder ZIP file from the following website:
http://download.macromedia.com/pub/flex/sdk/datavisualization_sdk3.5.zip
 - Download the Flex 3.5 Automation Libraries for Flex Builder ZIP file from the following website:
http://download.macromedia.com/pub/flex/sdk/automation_sdk3.5.zip
2. Expand these ZIP files into a new directory on your system (for example, c:\flex3.5\sdk\).

Note: When you are adding the automation libraries to the SDK directory, they must be expanded into the FLEX_HOME/frameworks directory, and *not* in the FLEX_HOME directory.

For additional information about adding SDK versions to Adobe Flash Builder (or Flex Builder), visit the Adobe website:

<http://www.adobe.com/support/flashbuilder/>

3. From within Adobe Flash Builder, from the **Window** menu, select **Preferences**.
4. In the Preferences dialog, expand **Flash Builder**, then select **Installed Flex SDKs**.

Note: In Adobe Flex Builder 3, from the Preferences dialog, expand **Flex**.

5. Click **Add** and browse to the location where you unpackaged the ZIP files in step 2.

8.30.2.2 Setting up the Demo Sample Project

To set up the demo sample project:

1. Locate the demo_hostsample.zip file in the EDK distribution.
2. Copy this file to a location on your Windows system where you installed and configured Adobe Flash Builder.
3. From Adobe Flash Builder, from the **File** menu, select **Import Flex Project**.
4. Browse to the location where the demo_hostsample.zip file is located, select the file, then click **Finish**.

A warning dialog (Project Will Be Upgraded) appears, which indicates that "Project 'demo_hostsample' was created with a previous version of Flash Builder."

A second dialog (Choose Flex SDK Version) appears with the same warning

5. In the second dialog, select **Use a specific SDK** and from the list, select **Flex 3.5**, then click **OK**.

6. In the first dialog, click **OK** to upgrade the project.

The demo_hostsample project appears in the Adobe Flash Builder navigator.

7. Navigate to the mpcui/src directory to view the source files that comprise the plug-in UI.

When these files are compiled, the mpcui/bin-debug/HostSample.swf file is created. This is the Flex application that is the custom UI for this target and it is included in the plug-in.

For information about the SWF file, see [Section 8.7, "Packaging the MPCUI Implementation With the Plug-in"](#) and [Section 8.6, "Defining the MPCUI Application"](#).

8.30.2.3 Running Demo Sample MPCUI from Adobe Flash Builder

Note: One of the advantages of MPCUI is that it allows you to test your Flex UI as part of the deployed plug-in or by running it from within Adobe Flash Builder directly. This latter option makes iterative development much simpler, however it requires that at least one version of the plug-in is deployed to Enterprise Manager and that a target instance has already been discovered before attempting to run the UI (Flex application) from Adobe Flash Builder.

After the Demo Sample plug-in has been deployed and the demo_hostsample.zip project has been imported into Adobe Flash Builder and builds successfully, then you can run the Flex application from Adobe Flash Builder.

1. From the Navigator, select **demo_hostsample**.
2. From the **Run** menu, select **Run HostSample** or **Debug HostSample**.

A browser window appears with a Management Server Connection login dialog.

Note: If the Management Server Connection dialog does not appear or if any other error appears, then verify that the project was imported correctly and verify that no errors appear in the **Problems** or **Console** tabs that appear in the bottom of the Flash Builder window.

During normal operation, when the user accesses your UI through the Enterprise Manager console by going to a target home page, this dialog does not appear because the UI is running as an integral part of the Enterprise Manager console and is embedded in a console session.

However, when running from Adobe Flash Builder, your UI requires information to connect to the Enterprise Manager site where your plug-in has been deployed and where the target instance that you will manage is located. Enter the same information for host, port and credentials that you would use to connect to your running Enterprise Manager console. Use either http or https depending on your configuration; however you must ensure that the ports you supply are correct for the protocol supplied.

3. Below **Target to Monitor**, enter the target name and type (the internal type and not the displayed label) of a target instance associated with this plug-in. It must be a target that exists in Enterprise Manager already. If you are using the Demo

Sample, then the target type is **demo_hostsample**, and the name is the target name you provided when creating the target instance.

4. Click **OK**.

The Demo Sample home page appears and is populated with data.

Note: In this mode, the Enterprise Manager page decorations and the target context area do not appear at the top of the page but they will appear when you access the target home page from the Enterprise Manager console. A menu appears that allows you to exercise the multiple pages included in your custom UI.

8.30.2.4 Elements of the Demo Sample Flex UI

The following is a brief list of the components that make up the Demo Sample Flex UI. For more comments that describe the purpose of each file and the items demonstrated in each file, see the source code.

```
demo_hostsample
  html-template/data/demo_hostsample_menu.xml
  mpcui/src
    HostSample.mxml           The application definition, must extend
                              MpApplication and implement the single
                              method "getIntegrationClass"
    HostSampleInteg.mxml      The integration class, defines the set of
                              pages, dialogs, trains, etc. that make up
                              the application
    HomePage.mxml             The target homepage, contains the layout of
                              the UI for the homepage
    HomePageController.as      The controller for the HomePage, includes
                              the methods that load data for the page
                              and respond to events from the page
    ProcessesPage.mxml
    ProcessesPageController.as
    CredentialsPage.mxml
    CredentialsPageController.as
    AvailabilityDialog.mxml
    MetricHistoryDialog.mxml
```

8.30.2.5 Updating the Demo Sample

As you modify and rebuild your UI in Adobe Flash Builder, you can run or debug the Flex application directly from Adobe Flash Builder as you make changes.

Ensure that there are no compiler errors shown in the **Problems** tab before attempting to proceed. Adobe Flash Builder displays a warning if you try to run the application with unresolved errors.

Note: Although you can see your updated UI in Adobe Flash Builder, you have not updated the version that is deployed to Enterprise Manager. The updated version will not appear in the Enterprise Manager console until you update the plug-in.

8.30.2.6 Modifying the Deployed Plug-in

After you make changes to your UI in Adobe Flash Builder, you can apply the changes to your plug-in so that you can also view the updates from a target home page within the Enterprise Manager console.

To do this, you must either create and deploy a new version of your plug-in or use the metadata registration service (MRS).

MRS allows you to apply incremental updates to your plug-in without creating and deploying a entirely new version. For more information about MRS, see [Section 13.7](#).

After you have modified your custom UI by modifying your Flex code in Adobe Flash Builder:

1. Rebuild the SWF file (/bin-debug)
2. FTP or copy the SWF file to the server where your Enterprise Manager site is installed and where you deployed the Demo Sample plug-in (HostSample.swf) originally.
3. Copy this file to the location where you created the plug-in staging directory:

```
stage/oms/metadata/mpcui
```

Note: There is an existing version of this file (HostSample.swf) in the directory along with an MPCUI metadata XML file.

4. Update the plug-in using the following command:

```
emctl register oms metadata -sysman_pwd sysman -pluginId oracle.sysman.ohs
-service mpcui -file demo_hostsample_uimd_swf.xml
```

For information about the `emctl` command, see [Section 13.7](#).

8.30.3 Setting Up an Adobe Flash Builder Project for MPCUI

This section assumes you are using Adobe Flash Builder 4.5 but includes notes for Adobe Flash Builder 3 users.

To set up an Adobe Flash Builder project, you can create an empty project or import the `demo_hostsample` project in to Adobe Flash Builder to use as a template. If you are importing the `demo_hostsample` project, then complete the steps in [Section 8.30.3.1, "Before You Begin"](#). Otherwise, proceed to [Section 8.30.3.2, "Creating an Adobe Flash Builder Project"](#).

Note: You must have configured Adobe Flash Builder to include the Flex 3.5 SDK files. For more information about including the Flex 3.5 SDK files, see [Section 8.30.2.1, "Configuring Adobe Flash Builder"](#).

8.30.3.1 Before You Begin

If you are using the `demo_hostsample` project as a template, you must complete the following steps before you set up the Adobe Flash Builder project for MPCUI.

1. Delete the contents of the following directories:
 - /stage
 - /scripts

- /rsc

These directories provide support for deploying the sample plug-in, but are not appropriate to your new project.

2. From the mpcui/metadata directory, rename the demo_hostsample_uimd_swf.xml file to *targettype_mpcui.xml*, where *targettype* is the name of your target type.

Delete all the other contents of the mpcui/metadata directory except *targettype_mpcui.xml*.

3. From the mpcui/data directory, edit the mpCuiProperties.xml file as follows:

- Replace the OMS connection with the information for connecting to your Enterprise Manager server.

```
<!-- Default OMS Connection -->
<hostname>myhost.us.example.com</hostname>
<port>7777</port>
<emUser>sysman</emUser>
<password>sysman_password</password>
```

- Replace the <metadata> tag with the file name as created in step 2 (/metadata/*targettype_mpcui.xml*)

```
<!-- the filename that includes the mpcui meta-data that will be included
in the plug-in. This is used to populate the menus for testing of the
UI in standalone (FlashBuilder) mode
If not specified then a default filename of <targetType>_menu.xml will be
used.-->
<metadata>../metadata/targettype_mpcui.xml</metadata>
```

4. From the /src directory, rename the HostSample.mxml (application definition) and HostSampleInteg.mxml (integration class) files to file names relating to your target, such as MyTargetUi.mxml and MyTargetUiInteg.mxml. For more information about these files, see [Section 8.6, "Defining the MPCUI Application"](#).

Delete all the other contents of the /src directory except for these two files, MyTargetUi.mxml and MyTargetUiInteg.mxml files.

5. Ensure that the following two files are located in the /mpcui/libs directory:

- mpcui_12.1.0.2.0.swc
- mprslldr_12.1.0.2.0.swc

Note: These files are version specific. If you are moving from an earlier version of the EDK to a later version, then you must replace these files with the latest version and ensure the project properties are updated to reflect this change.

6. Ensure that the stylesFusionFX.swf file is located in the /html-template directory. If this file is missing, then locate the file in the demo_hostsample.zip file and add it to the /html-template directory.

7. From Adobe Flash Builder, select the project properties (from the **File** menu, select **Properties**), and select **Flex Compiler**.

Ensure that the Flex SDK version is set to Flex 3.5. If there is an error, then the location of your Flex 3.5 SDK files is not the same as the location used for the sample.

Click **Configure Flex SDKs...** and then edit the settings for Flex 3.5 to point to the location where you installed the Flex 3.5 SDK files.

8.30.3.2 Creating an Adobe Flash Builder Project

From Adobe Flash Builder, create a new project and complete the following steps to set up the project:

1. Name the project a meaningful name related to your target (such as MyTargetUI). This name will be assigned to the default application file after the project is created. For example, /src/MyTargetUi.mxml.
2. Ensure that the Flex SDK version is set to Flex 3.5. For more information, see [Section 8.30.2.1, "Configuring Adobe Flash Builder"](#). Click **Finish** to create the application.
3. In the demo_hostsample.zip file, locate the mpcui/html-template/stylesFusionFX.swf file. Make a copy of this file into your project under the /html-template directory.
4. In the demo_hostsample.zip file, locate the mpcui/libs directory, and copy the following files into your project under the /libs directory:
 - mpcui_12.1.0.2.0.swc
 - mprslldr_12.1.0.2.0.swc
5. Create a directory called metadata and create a file called *targettype_mpcui.xml*. This file contains the MPCUI metadata for your plug-in. For more information about this file, see the demo_hostsample and [Section 8.4, "Creating the MPCUI Metadata File"](#).
6. In the demo_hostsample.zip file, locate the mpcui/data/mpCuiProperties.xml file and copy it into your project in a new directory called data.

Edit this file as follows:

- Replace the OMS connection with the information for connecting to your Enterprise Manager server.

```
<!-- Default OMS Connection -->
<hostname>myhost.us.example.com</hostname>
<port>7777</port>
<emUser>sysman</emUser>
<password>sysman_password</password>
```

- Replace the <metadata> tag with the file name as created in step 5 (/metadata/targettype_mpcui.xml)

```
<!-- the filename that includes the mpcui meta-data that will be included
in the plug-in. This is used to populate the menus for testing of the
UI in standalone (FlashBuilder) mode
If not specified then a default filename of <targetType>_menu.xml will be
used.-->
```

```
<metadata>../metadata/targettype_mpcui.xml</metadata>
```

7. In the /src directory, you should have a *projectname.mxml* file. You must modify this file to correspond to the required settings of an MPCUI application and add an integration class. For more information, see [Section 8.6, "Defining the MPCUI Application"](#).
8. Ensure that the Project properties are set correctly as described in [Section 8.30.3.3, "Setting MPCUI Project Properties"](#).

8.30.3.3 Setting MPCUI Project Properties

It is critical to ensure that the project properties for an MPCUI application are set correctly to ensure that the project operates successfully.

1. From Adobe Flash Builder, select the project properties (from the **File** menu, select **Properties**), and select **Flex Compiler**.
 - Check that **Flex SDK version** is set to Flex 3.5
 - Ensure that **Enable strict type checking** is not selected
2. From Adobe Flash Builder, select **Project**, then select **Properties**, then select **Flex Build Path** and click the **Library Path** tab.
 - If you are starting from a new project, then you must edit the library path by accessing the project's build path properties page:
 - Remove the **libs** entry in the Build Path libraries.
 - Click **Add SWC** and select the `mprslldr.swc` file.
 - Repeat the previous step to select the `mpcui.swc` file.

These steps are necessary so that each file can have different linkage settings.
 - Ensure that **Verify RSL Digests** is not checked.
 - Expand the `mprslldr.swc` entry and ensure that **Link Type** is set to **Merged into code**. If not, then select **Link Type** and change the setting to **Merged into Code**.
 - If you are using Flex Builder 3, check the following:
 - Expand the MPCUI swc entry, select **Link Type** and then click **Edit**.
 - Change **Link Type** to **Runtime shared library (RSL)**.
 - Set **Verification** to **None (trusted environments only)**
 - Set **Deployment Path/URL**: to `mpcui.swf` and select **Automatically extract swf to deployment path**.
 - If you are using Adobe Flash Builder 4.5:
 - Expand the MPCUI swc entry, select **Link Type** and then click **Edit**.
 - Change **Link Type** to **Runtime shared library (RSL)**.
 - Click **Add** next to **RSL deployment paths**: and enter `mpcui.swf` for **Deployment Path/URL** and ensure that **Automatically extract SWF to deployment path** is selected.
 - If an entry already exists under **RSL Deployment paths** and it is *not* `mpcui.swf`, then edit it to ensure the correct path as noted in the previous step.

8.30.3.4 Verifying Correct MPCUI Library Linkage

When you have set up your project and are building a SWF, you must verify that it operates correctly when deployed to any of the various Enterprise Manager platform versions that support the version of the EDK you are using.

Initially, verify the size of your SWF file. If it is larger than 1 MB, then you might have linked your SWF incorrectly. A typical MPCUI SWF file is 300-600 KB. This is not a definitive test, but is a good indication whether your SWF is linked correctly.

To verify that your SWF is linked correctly:

Note: These steps are specific to deploying and running the plug-in in a 12.1.0.2.0 environment.

1. Update the SWF file in your plug-in and deploy the new version of your plug-in to the Enterprise Manager server.

Note: If your plug-in is deployed already, then you can use the `emctl register oms metadata` command to update the MPCUI part of your plug-in only. For more information, see [Section 13.7](#).

2. Access the home page for a target that is associated with this plug-in and the SWF file included in your plug-in.
3. If the home page load fails with an error in the `TargetInfoService`, you might have linked with an older swc file (pre-12.1.0.2.0) but not linked it as an RSL. If this happens, then do the following:

- Modify the URL in the browser address by appending the following text and then reload the page:

```
&traceEnabled=true
```

- When the page loads and the error appears, dismiss the error dialogs and in the request trace window, locate and select the `TargetInfoService.getTargetInfo` entry, then select the `Shot Item Details`.
- Search the response message that appears in the right-hand side of the window at the bottom for "rslVersionError". If you find this text, it indicates that the server rejected the request coming from a client that is incorrectly linked. Review the project properties to ensure that you have correctly linked the SWF as described in [Section 8.30.3.3](#), "Setting MPCUI Project Properties".

If the error is included in the response message, typically it appears as follows:

```
<rslVersionError
SOAP-ENV:actor="http://schemas.xmlsoap.org/soap/actor/next"
xmlns="http://em.oracle.com/">12.1.0.2.0
</rslVersionError>
```

4. If the home page loads correctly, then modify the URL in the browser address bar by appending the following text and then reload the page

```
&loglevel=ALL,CONSOLE
```

5. When the page has loaded, check the log output displayed in the logging window at the bottom of the page and search for the following text:

```
[MpCui] INFO Application Linked with Version
```

If the application is correctly linked, then the version of the linked swc files should appear on this line.

The next line should include the following text and indicate the version of the Enterprise Manager server to which you are connected. If not, then your SWF is linked incorrectly:

```
[MpCui] INFO RSL Library Loaded Version
```

Example 8–50 Log Output

```
2012-05-04 13:37:44.269 [MpCui] INFO Application Linked with Version: 12.1.0.1.1
2012-05-04 13:37:44.270 [MpCui] INFO RSL Library Loaded Version      : 12.1.0.2.0
2012-05-04 13:37:44.271 [MpCui] INFO RSL Library Min Supported Vers : 12.1.0.1.0
2012-05-04 13:37:44.273 [MpCui] INFO RSL Build Date                 : 2012.04.02
2012-05-04 13:37:44.274 [MpCui] INFO Application & RSL Versions are COMPATIBLE
```

8.30.3.5 Using Earlier SWC Files With a Later Enterprise Manager OMS

If you are building your MPCUI SWF against an earlier version of the EDK (such as version 12.1.0.1.0) but then want to run it against an existing OMS that is part of a later version (such as 12.1.0.2.0), you must deploy the updated SWF to Enterprise Manager as part of your plug-in.

To do this, you can use Adobe Flash Builder, however your Adobe Flash Builder project will contain and use the older version of the `mpcui.swf` file and will use the local copy within Adobe Flash Builder and not the later version from the OMS. This results in unpredictable results, therefore follow the steps below:

To run in this mode:

1. Copy the `mpcui.swf` file from the Enterprise Manager server (such as 12.1.0.2.0) add the file to the `bin-debug` directory under your project.

Note: Ensure that **Build automatically** is disabled so that your project does *not* rebuild and replace this file with an updated copy based on the earlier MPCUI swc file.

2. Replace the `mpcui.swf` file in Adobe Flash Builder each time you update and rebuild the project as this build step replaces the `mpcui.swf` file with the file associated with the MPCUI swc file that you are compiling against.

8.30.3.6 Understanding MPCUI RSL Caching

An advantage of using the MPCUI RSL is that it allows the code included in this library to be cached in the browser of Adobe Flash Player. Therefore, the code must not be reloaded each time an Enterprise Manager page that includes an MPCUI SWF file is accessed.

However, when you upgrade your Enterprise Manager site, any users that accessed pages from the previous version, will still have the older version of the MPCUI RSL cached in their browser. This release includes additional checks to the server to validate this, and the user might see an error indicating that the page could not be loaded due to a version mismatch in the MPCUI RSL. If this error appears, you must clear the Adobe Flash Player cache before continuing. For more information, see [Section 8.30.3.7, "Clearing Adobe Flash Player Cache"](#).

8.30.3.7 Clearing Adobe Flash Player Cache

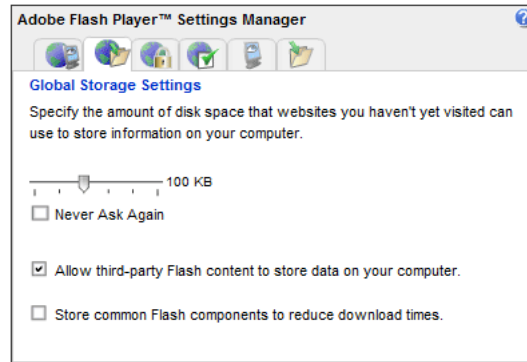
To clear the Adobe Flash Player cache without clearing the browser cache:

1. To access the Adobe Flash Player Settings Manager: Global Storage Settings page, open the following URL:

http://www.macromedia.com/support/documentation/en/flashplayer/help/settings_manager03.html

The following window appears.

Figure 8–29 Adobe Flash Player Settings Manager: Global Storage Settings



2. Deselect **Store common Flash components to reduce download times**.
3. Click **Confirm** on the confirmation window to clear any cached RSLs.
4. Select **Store common Flash components to reduce download times**.

This enables caching again and any previously cached RSLs will be downloaded.

8.31 Migrating Home Page Customizations

Earlier versions of the Enterprise Manager extensibility framework supported the ability to customize the default Enterprise Manager home page by:

- setting a set of charts to display on the home page
- defining a series of related links to display on the home page

For Enterprise Manager release 11.2, you could choose to continue to use these customizations as a basis for the UI for your target. This avoids implementing a custom UI for your target but maintains limited control over the home page.

Use the `empdk` utility to read the home page customizations from Enterprise Manager and generate the required 11.2 files to define this metadata. For information about the `empdk` utility, see [Chapter 16](#).

Note: The ability to read home page customizations from an Enterprise Manager release earlier than 11.2 and then generate 11.2 metadata is not supported for this release.

8.32 Accessibility Guidelines

The MPCUI framework is designed to support a user interface that complies with the Oracle Global HTML Accessibility Guidelines (OGHAG). This section provides information about accessibility standards for your UI implementation.

Also, Adobe provides guidelines and information to help with the implementation of Flex applications (on which MPCUI is based) to meet accessibility standards. For more information, see the following websites:

- Adobe Flex Accessibility Page
<http://www.adobe.com/accessibility/products/flex/>
- Adobe Flex Accessibility Best Practices
<http://blogs.adobe.com/accessibility/files/2011/03/Flex-4-Accessibility-Best-Practices.pdf>
- Adobe Flex Accessibility Blog
<http://blogs.adobe.com/accessibility/category/flex>

8.32.1 Accessibility Options in Enterprise Manager

Enterprise Manager provides the end user with the ability to set options for accessibility including a screenreader option. The MPCUI framework is aware of these settings and makes them available to you in your Flex code (see the `oracle.sysman.emx.util.AdaSettings` in the API reference).

Typically you do not have to check for these settings because MPCUI automatically renders accessible components when the end user sets their account to require an accessible user interface. Among other things, this replaces charts with an accessible view of the same data.

8.32.2 Summary of Critical Issues

When constructing an accessible MPCUI Flex application, consider the following items:

- Enable accessibility
Add the accessible flag to the compiler settings.
For more information, see the Adobe Best Practices and the settings in the `demo_hostsample` example plug-in
- Use MPCUI Pages, dialog and components
These components include accessibility support.
For more information, see [Section 8.32.3, "Using MPCUI Components"](#).
- Set Tab Order of Components
Sets the reading and tab order for all components.
For more information, see [Section 8.32.4, "Controlling Reading and Tabbing Order"](#).
- Set Name and Description
For components that require additional text description (such as images).
For more information, see Adobe Best Practices.
- Avoid the use of or provide accessible equivalents for drag-and-drop, audio, and decorative content.
For more information, see Adobe Best Practices.
- Avoid conveying information using color
For more information, see Adobe Best Practices.

- Do not add menus to the Flex Application.
Use the Enterprise Manager Target menu.
For more information see this guide.
- Avoid using non-accessible Flex components
For a list of accessible components, see Adobe Best Practices.

8.32.3 Using MPCUI Components

Because the MPCUI framework provides components that include accessibility support beyond the base Flex components, use the MPCUI version of those components. In addition to the specific components listed in [Table 8–2](#), the application should include only MPCUI top-level activities such as Page, Dialog, or Train. Implement your application based on the `MpApplication` base class. [Table 8–2](#) provides a list of important components included in the MPCUI framework:

Table 8–2 Important MPCUI Components

Flex Component	MPCUI Alternative
<code>mx.controls.Panel</code>	<code>oracle.sysman.emx.components.Region</code>
<code>mx.controls.Label</code>	<code>oracle.sysman.emx.components.AccLabel</code>
<code>mx.controls.AdvancedDataGrid</code>	<code>oracle.sysman.emx.table.Table</code>
<code>mx.charts.PieChart</code>	<code>oracle.sysman.emx.charts.PieChart</code>
<code>mx.charts.BarChart</code>	<code>oracle.sysman.emx.charts.BarChart</code>
<code>mx.charts.ColumnChart</code>	<code>oracle.sysman.emx.charts.ColumnChart</code>
<code>mx.charts.AreaChart</code>	<code>oracle.sysman.emx.charts.AreaChart</code>
<code>mx.charts.LineChart</code>	<code>oracle.sysman.emx.charts.LineChart</code>

8.32.4 Controlling Reading and Tabbing Order

While Flex provides support for determining the reading and tabbing order of components included in the application, it does not work well for complex layouts that include multiple regions and pages. To ensure that the application provides a tab order that make sense, set the order within each page or dialog.

The MPCUI framework provides a means of setting the tab order. Use this method instead of setting the `tabIndex` of each property. Do not set the `tabIndex` of the components included in your pages, but use the `tabOrder` property of the page or dialog.

To use the `tabOrder` property, you must assign a unique ID to every component included in the Page. All Flex and MPCUI components support the `id` property. [Example 8–51](#) shows a Page declaration that includes the `tabOrder` property.

Example 8–51 Page Declaration That Includes the `tabOrder` Property

```
<mp:Page label="Home Page"
  xmlns:mx="http://www.adobe.com/2006/mxml"
  xmlns:mp="http://www.oracle.com/mpcui"
  tabOrder="{[ summaryRegion, statusRegion, currentStatus, statusSince,
availability,
configurationRegion, cpuModel, cpuPer, currentLoad, osVersion,
relatedHost,
```

```

        installedSoftware, jobSummary,
        reportsRegion, allReports, hostPerfReport,
        currentLoadRegion, cpuLoadRegion, cpuload, showCpuHistory,
        processSummaryRegion, numUsers, numProcesses,
        resourcesRegion, cpuUtilRegion, cpuutil,
        processorRegion, processorChart, showProcessorHistory,
        memoryRegion, selMemChart, memChart,
        eventsRegion
    }}"
    defaultComponent="{currentStatus}"
>

<!-- First Column, 25% width of page, two regions stacked vertically -->
<mx:HBox width="100%" height="100%">
    <mx:VBox width="25%" height="100%" >
        <mp:Region id="summaryRegion" title="Summary" height="50%"
width="100%" >
            <mp:InnerRegion id="statusRegion" title="Status" height="40%"
width="100%" >
                <mp:InfoDisplay width="100%" height="100%">
                    <!-- reference to the AvailDataService -->
                    <mp:InfoItem id="currentStatus" label="Current Status"
value="{ads.currentStatus}" image="{ads.currentStatusIcon}"
                        click="invokeActivity(Constants.PAGE_AVAILABILITY,
bean(Constants.P_TARGET_NAME, appModel.target.name,
                            Constants.P_TARGET_TYPE, appModel.target.type));" />
                    <mp:InfoItem id="statusSince"
source="{ads.statusSinceItem}" />
                    <mp:InfoItem id="availability" label="Availability %"
value="{NumberFormat.formatNumber(ads.availPercent,1)}%" destination="availDialog"
/>

                </mp:InfoDisplay>
            </mp:InnerRegion>

```

The `tabOrder` property is an array of the component ids included in the page that must be included in the tabbing or reading order.

For more information, see the `demo_hostsample` sample application and the API reference document.

Note: In addition to setting the `tabOrder` property for components appearing to the screen reader, set the `defaultComponent` property also. This instructs the screen reader to set focus to a particular component when the page is rendered initially.

8.33 Localization Support

MPCUI does not provide support for localized text resources.

8.34 Providing Online Help

If you want to include online help for your customized UI pages, package the help JAR files in the following directory:

```
plugin_stage/oms/online_help
```

For an example of a help JAR file, see the `plugin_sample_help.jar` in the `/oms/online_help` directory of the `demo_hostsample` example in the EDK.

Customizing Incident Manager

This chapter describes how to customize the event details page to provide more diagnostic information about the event and to facilitate quicker resolution of the underlying issue. Details pages on the Incident Manager UI allow users to view the details of an event. The content of such pages helps the user understand the basic nature of the underlying issue and provides additional contextual details (such as text, links to diagnostic or resolution pages) to resolve the issue quickly.

For incidents that have only one event, the customizations applied to the event details page are automatically applied to the Incident Details page.

Note: For information about Incident Management, see the "Using Incident Management" chapter of the *Oracle Enterprise Manager Cloud Control Administrator's Guide*.

<http://www.oracle.com/pls/em121/homepage>

This chapter contains the following sections:

- [Introduction to Customizing Incident Manager](#)
- [Understanding Supported Customizations](#)
- [Creating Event-Specific Customization XML](#)
- [Adding Customized Details About the Event](#)
- [Providing Content in the Guided Resolution Region](#)
- [Defining a Search String for My Oracle Support Knowledge](#)
- [Defining Conditions for Customization](#)
- [Registering Customizations](#)
- [Testing Incident Manager After Customization](#)

9.1 Introduction to Customizing Incident Manager

As a plug-in developer, you are responsible for the following steps within customizing Incident Manager:

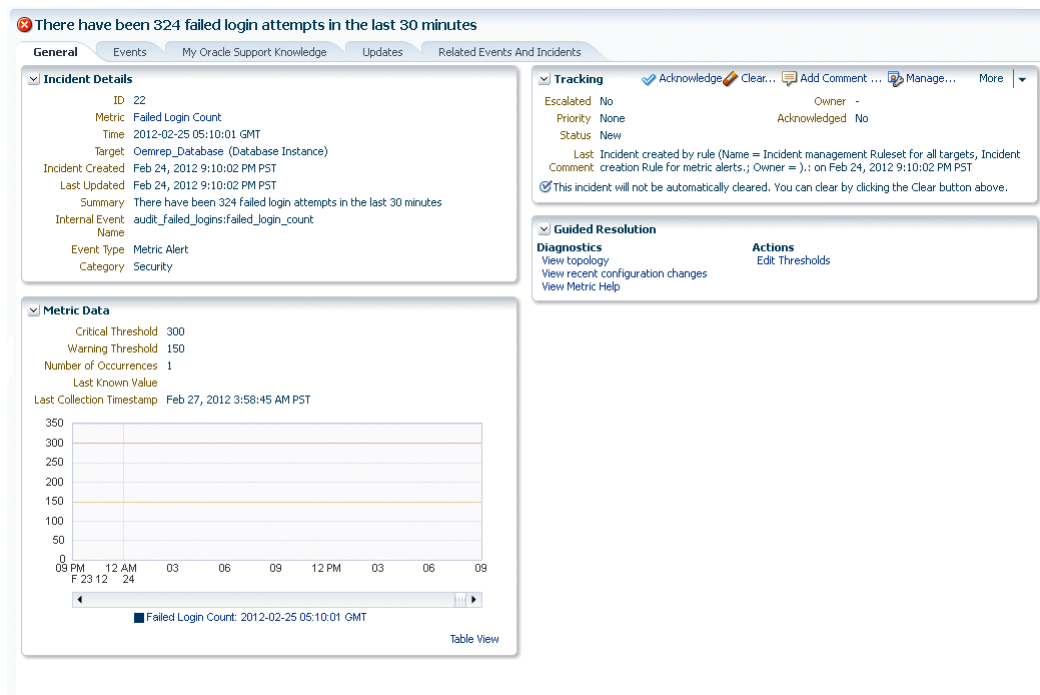
1. Determine what customizations you require for your Incident Manager UI. For fine-grained access, use conditions. For more information, see [Section 9.2, "Understanding Supported Customizations"](#).

2. Construct the customization XML according to the XSD. For more information, see [Section 9.3, "Creating Event-Specific Customization XML"](#).
3. Register your customization. For more information, see [Section 9.8, "Registering Customizations"](#).
4. Test the UI by publishing an event that matches the condition. For more information, see [Section 9.9, "Testing Incident Manager After Customization"](#).

9.2 Understanding Supported Customizations

[Figure 9–1](#) displays the General tab for a selected incident from the Incident Manager page.

Figure 9–1 Incident Manager



The following customizations are supported for this page:

- Adding name-value pairs to the Incident Details region.
For more information, see [Section 9.4, "Adding Customized Details About the Event"](#).
- Customizing Action and Diagnostic links in the Guided Resolution region.
For more information, see [Section 9.5, "Providing Content in the Guided Resolution Region"](#).
- Adding recommendations to the Guided Resolution region.
For more information, see [Section 9.5.1, "Adding Recommendations using XML"](#).
- Specifying the default search phrase for My Oracle Support Knowledge.
For more information, see [Section 9.6, "Defining a Search String for My Oracle Support Knowledge"](#).

Each customization specification has two parts:

1. Condition

This is the criteria used to identify an event for which the customized content will be rendered. For example, consider a scenario where you want to show a diagnostic link for metric alerts on a database. The condition would be "event class is metric_alert and target type is oracle_database". Another example is where you to show the region containing a metric chart. This condition would be "event class is metric_alert and metric_type is numeric".

Note: Any target type name is supported. While matching an event, you match the target type in the condition with the target type of the event

2. Action

The actions specify the customized content. For example, the specification of the diagnostic link (that is, the label and the URL to be shown under it).

9.3 Creating Event-Specific Customization XML

Oracle provides an event-specific customization XSD so that you can write XML to describe customizations for a specific event for display on the Incident Manager UI.

Note: For a complete event-specific customization XML Schema Definitions (XSD), see the Extensibility Development Kit (EDK).

The event-specific customization XSD defines how the Incident Manager UI supports UI customization.

You can define fine-grained conditions to customize the Event Details or Incident Details pages.

Example 9–1 Sample Metadata File

```
<evt:CustomUI AppliesTo="EVENT" EventClass ="metric_alert" TargetType =host">
  <evt:ConditionDetails>
    <evt:Condition>
      <evt:Attrib Name="metric_name" Value="Load"/>
      <evt:Attrib Name="metricColumn" Value="cpuUtil"/> </evt:Condition>
    </evt:ConditionDetails>
  </evt:CustomUI>
```

Oracle recommends the following naming conventions for your metadata XML:

- *event_class_description.xml*

In the preceding file name:

- *event_class* represents the name of the event class

Event customization supports the following event classes:

- * metric_alert
- * target_availability
- * job_status_change
- * cs_rule_violation

- * `cs_score`
 - * `sla_alert`
 - * `metric_error`
 - *description* represents a short description of the event customization
- For example, `job_status_change_recommendation.xml`
- *event_class_target_type_description.xml*
- In the preceding file name:
- *event_class* represents the name of the event class
 - *target_type* represents the name of the target type for which this event is generated
 - *description* represents a short description of the event customization
- For example, `host_metric_alert_diaglinks.xml`

Note: The maximum length of the file name is 255 characters.

For information about the directory location for the metadata XML, see [Section 9.8, "Registering Customizations"](#).

Note: Use the `empdk validate_plugin` command to validate the XML metadata file. For more information about the `empdk validate_plugin` command, see [Chapter 13, "Validating, Packaging, and Deploying the Plug-in"](#).

9.3.1 Overview of Event-Specific Customization Metadata Elements

[Table 9–1](#) describes the key elements that define the event-specific customization XML.

Table 9–1 Key Elements in Event-Specific Customization XML

Element	Description
<code>evt:CustomUI</code>	<p>This is the root element of the XML. It defines the customization.</p> <p>It includes the following attributes:</p> <ul style="list-style-type: none"> ■ AppliesTo: Applicable to event customizations. The only valid value is <code>EVENT</code>. ■ EventClass: Specifies the internal event class name and is applicable only when the customization applies to an event. ■ TargetType: Internal name of the target type. The customization applies to events from all targets of this target type.
<code>evt:ConditionDetails</code>	Specifies the criteria on which the customizations are to be applied.
<code>evt:Condition</code>	<p>Specifies a condition for the customization.</p> <p>Note: Oracle supports one condition only within the <code>evt:ConditionDetails</code> tag.</p>

Table 9–1 (Cont.) Key Elements in Event-Specific Customization XML

Element	Description
<code>evt:DetailUI</code>	Specifies that you are customizing the Details region of the Incident Manager UI page. For more information, see Section 9.4, "Adding Customized Details About the Event" .
<code>evt:GuidedResolutionDetails</code>	Specifies that you are customizing the Guided Resolution region of the Incident Manager UI page. Using this element, you can add action links, diagnostic links, and recommendations. For more information, see Section 9.5.1, "Adding Recommendations using XML" .

9.3.2 About Events

This section provides common event attributes and the definition of the two most commonly-used event types:

- [Common Event Attributes](#)
- [Target Availability Event](#)
- [Metric Alert Event](#)

9.3.2.1 Common Event Attributes

All events have the following common attributes:

Table 9–2 Common Event Attributes

Attribute	Description
<code>sys_event_class</code>	Event type Possible values: <ul style="list-style-type: none"> ■ <code>target_availability</code>: Target Availability events ■ <code>metric_alert</code>: Metric Alert events
<code>sys_event_name</code>	Event name to identify the nature of the event uniquely
<code>sys_event_key</code>	Name of a subcomponent within the event source object to which this event is related. This is optional. Examples include a disk name on a host, name of a tablespace, and so on
<code>sys_event_msg</code>	Event message
<code>sys_action_msg</code>	Action message
<code>sys_source_obj_type</code>	Source object type. For example, JOBS for job-based events.
<code>sys_source_obj_id</code>	Unique internal identifier of a Source object
<code>sys_target_guid</code>	Unique internal identifier of a target
<code>sys_target_name</code>	Target name
<code>sys_target_owner</code>	Target owner
<code>sys_target_version</code>	Target version
<code>sys_target_lifecycle_status</code>	Lifecycle status
<code>sys_incident_id</code>	Incident ID

Table 9–2 (Cont.) Common Event Attributes

Attribute	Description
sys_severity	Severity of the event Possible values: <ul style="list-style-type: none"> ■ 32: Fatal ■ 16: Critical ■ 8: Warning ■ 4: Minor Warning
sys_category	Event category. Possible values: <ul style="list-style-type: none"> ■ Availability: 1 ■ Configuration: 2 ■ Capacity: 4 ■ Fault: 8 ■ Load: 16 ■ Performance: 32 ■ Security: 64 ■ Jobs: 128 ■ Diagnostics: 256 ■ Error: 512 ■ Business: 1024

9.3.2.2 Target Availability Event

The Target Availability Event represents a target's availability status.

[Example 9–2](#) shows the event attributes defined by the target availability XML file.

[Table 9–3](#) provides a list of all the event attributes for target availability.

Example 9–2 target_availability.xml

```

<evt:EventClass Name="target_availability"
  NLSID="TARGET_AVAILABILITY"

  ResourceBundle="oracle.sysman.core.common.events.classes.rsc.availability.AvailabilityEventsMsg"
  TargetAware="ALWAYS"
  SourceObjectType="TARGET"
  Version="1.0"
  xmlns:evt="http://www.oracle.com/EnterpriseGridControl/eventclass_model"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"

  xsi:schemaLocation="http://www.oracle.com/EnterpriseGridControl/eventclass_model
  EventClass.xsd">

  <evt:DescriptionNLSID>TARGET_AVAILABILITY_DESC</evt:DescriptionNLSID>

  <evt:AttributeDef>

    <!--This attribute is used to store the availability status of a target-->
    <evt:Attrib Name="target_status"
      DataType="STRING"
      isReferenced="false"

```

```

        NLSID="TARGET_STATUS"
        isValueTranslatable="true">
        <evt:DescriptionNLSID>TARGET_STATUS_DESC</evt:DescriptionNLSID>
    </evt:Attrib>

    <!--The guid of the severity record associated with this availability
record-->
    <evt:Attrib Name="severity_guid"
        DataType="RAW"
        isReferenced="false"
        NLSID="SEVERITY_GUID"
        isValueTranslatable="false">
        <evt:DescriptionNLSID>SEVERITY_GUID_DESC</evt:DescriptionNLSID>
    </evt:Attrib>

    <!--The cycle guid of the severity record associated with this
availability record-->
    <evt:Attrib Name="cycle_guid"
        DataType="RAW"
        isReferenced="true"
        NLSID="CYCLE_GUID"
        isValueTranslatable="false">
        <evt:DescriptionNLSID>CYCLE_GUID_DESC</evt:DescriptionNLSID>
    </evt:Attrib>

    <!--The below attributes specifies the metric guid of response metric -->
    <evt:Attrib Name="metric_guid"
        DataType="RAW"
        isReferenced="true"
        NLSID="METRIC_GUID"
        isValueTranslatable="false">
        <evt:DescriptionNLSID>METRIC_GUID_DESC</evt:DescriptionNLSID>
    </evt:Attrib>

    <!--The below attribute represents a sub-state for availability states
like Status pending, Agent Unreachable and Blackout.
    TARGET STATUS          CODE      SUB_STATE

    Any                    0          None (Default)
    Agent unreachable     1          Normal
    Agent unreachable     2          Host Down
    Agent unreachable     3          Disk Full
    Status Pending        10         Normal
    Status Pending        11         Stuck
-->
    <evt:Attrib Name="avail_sub_state"
        DataType="NUMBER"
        isReferenced="false"
        NLSID="AVAILABILITY_SUB_STATE"
        isValueTranslatable="false">
        <evt:DescriptionNLSID>AVAILABILITY_SUB_STATE_
DESC</evt:DescriptionNLSID>
    </evt:Attrib>

    <!--The below attributes specifies the availability transition severity
that resulted in the target availability status that is specified by
target_status attribute -->
    <evt:Attrib Name="avail_severity"
        DataType="NUMBER"
        isReferenced="false"

```

```

        NLSID="AVAILABILITY_SEVERITY"
        isValueTranslatable="false">
        <evt:DescriptionNLSID>AVAILABILITY_SEVERITY_
DESC</evt:DescriptionNLSID>
    </evt:Attrib>
</evt:AttributeDef>

    <evt:RefAttribSource><![CDATA[ mgmt_avail.get_target_avail_ref_
attrs]]></evt:RefAttribSource>

    <!-- For availability we don't have any identifier attribute list. -->
    <!-- So system will use target_guid, event_class name to generate the
identifier attribute. -->

    <evt:RuleAttribs>
        <evt:RuleAttrib Name="target_status" /></evt:RuleAttrib>
        <evt:RuleAttrib Name="avail_sub_state" /></evt:RuleAttrib>
        <evt:RuleAttrib Name="avail_severity" /></evt:RuleAttrib>
    </evt:RuleAttribs>

    <evt:NotifAttribs>
        <evt:NotifAttrib Name="target_status" />
        <evt:NotifAttrib Name="severity_guid" />
        <evt:NotifAttrib Name="avail_sub_state" />
        <evt:NotifAttrib Name="avail_severity" />
        <evt:NotifAttrib Name="metric_guid" />
        <evt:NotifAttrib Name="cycle_guid" />
    </evt:NotifAttribs>

    <evt:Severities>
        <evt:Severity>FATAL</evt:Severity>
        <evt:Severity>CRITICAL</evt:Severity>
        <evt:Severity>WARNING</evt:Severity>
        <evt:Severity>MINOR_WARNING</evt:Severity>
        <evt:Severity>INFORMATIONAL</evt:Severity>
    </evt:Severities>
</evt:EventClass>

```

Table 9–3 Event Attributes for Target Availability

Attribute	Description
TARGET_STATUS	Availability status
AVAILABILITY_SUB_STATE	Availability substatus
AVAILABILITY_SEVERITY	Transition severity

9.3.2.3 Metric Alert Event

A metric alert event is generated when an alert occurs for a metric on a specific target (for example, CPU utilization for a host target) or metric on a target and object combination (for example, space usage on a specific tablespace of a database target)

[Example 9–3](#) shows the event attributes defined by the metric alert XML file. [Table 9–4](#) provides a list of all the event attributes for the metric alert event.

Example 9–3 metric_alert.xml

```
<evt:EventClass Name="metric_alert"
```

```

        NLSID="METRIC_ALERT_EVENT"
        TargetAware="ALWAYS"
        SourceObjectType="TARGET"

ResourceBundle="oracle.sysman.core.common.events.classes.rsc.metrics.MetricEventsMSG"

        Version="1.1"
        xmlns:evt="http://www.oracle.com/EnterpriseGridControl/eventclass_model"
        xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"

xsi:schemaLocation="http://www.oracle.com/EnterpriseGridControl/eventclass_model
EventClass.xsd">
    <evt:DescriptionNLSID>METRIC_ALERT_DESC</evt:DescriptionNLSID>
    <evt:AttributeDef>
        <evt:Attrib Name="metric_guid" DataType="RAW" isReferenced="false"
            NLSID="METRIC_GUID_NLSID" isValueTranslatable="false">
            <evt:DescriptionNLSID>METRIC_GUID_DESC</evt:DescriptionNLSID>
        </evt:Attrib>
        <evt:Attrib Name="key_value" DataType="STRING" isReferenced="false"
            NLSID="KEY_VALUE_NLSID" isValueTranslatable="false">
            <evt:DescriptionNLSID>KEY_VALUE_DESC</evt:DescriptionNLSID>
        </evt:Attrib>
        <evt:Attrib Name="severity_guid" DataType="RAW" isReferenced="false"
            NLSID="SEVERITY_GUID_NLSID" isValueTranslatable="false">
            <evt:DescriptionNLSID>SEVERITY_GUID_DESC</evt:DescriptionNLSID>
        </evt:Attrib>
        <evt:Attrib Name="cycle_guid" DataType="RAW" isReferenced="true"
            NLSID="CYCLE_GUID_NLSID" isValueTranslatable="false">
            <evt:DescriptionNLSID>CYCLE_GUID_DESC</evt:DescriptionNLSID>
        </evt:Attrib>
        <evt:Attrib Name="value" DataType="STRING" isReferenced="true"
            NLSID="VALUE_NLSID" isValueTranslatable="false">
            <evt:DescriptionNLSID>VALUE_DESC</evt:DescriptionNLSID>
        </evt:Attrib>
        <evt:Attrib Name="metric_group" DataType="STRING" isReferenced="true"
            NLSID="METRIC_GROUP_NLSID" isValueTranslatable="true">
            <evt:DescriptionNLSID>METRIC_GROUP_DESC</evt:DescriptionNLSID>
        </evt:Attrib>
        <evt:Attrib Name="metric_column" DataType="STRING" isReferenced="true"
            NLSID="METRIC_COLUMN_NLSID" isValueTranslatable="true">
            <evt:DescriptionNLSID>METRIC_COLUMN_DESC</evt:DescriptionNLSID>
        </evt:Attrib>
        <evt:Attrib Name="metric_description" DataType="STRING"
isReferenced="true"
            NLSID="METRIC_DESCRIPTION_NLSID" isValueTranslatable="true">
            <evt:DescriptionNLSID>METRIC_DESCRIPTION_DESC_
NLID</evt:DescriptionNLSID>
        </evt:Attrib>
        <evt:Attrib Name="coll_name" DataType="STRING" isReferenced="true"
            NLSID="COLL_NAME_NLSID" isValueTranslatable="true">
            <evt:DescriptionNLSID>COLL_NAME_DESC</evt:DescriptionNLSID>
        </evt:Attrib>
        <evt:Attrib Name="key_column_1" DataType="STRING" isReferenced="true"
            NLSID="ALERT_KEY_COL_NLSID_1" isValueTranslatable="true">
            <evt:DescriptionNLSID>ALERT_KEY_COL_DESC_1</evt:DescriptionNLSID>
        </evt:Attrib>
        <evt:Attrib Name="key_column_2" DataType="STRING" isReferenced="true"
            NLSID="ALERT_KEY_COL_NLSID_2" isValueTranslatable="true">
            <evt:DescriptionNLSID>ALERT_KEY_COL_DESC_2</evt:DescriptionNLSID>
        </evt:Attrib>

```

```

<evt:Attrib Name="key_column_3" DataType="STRING" isReferenced="true"
  NLSID="ALERT_KEY_COL_NLSID_3" isValueTranslatable="true">
  <evt:DescriptionNLSID>ALERT_KEY_COL_DESC_3</evt:DescriptionNLSID>
</evt:Attrib>
<evt:Attrib Name="key_column_4" DataType="STRING" isReferenced="true"
  NLSID="ALERT_KEY_COL_NLSID_4" isValueTranslatable="true">
  <evt:DescriptionNLSID>ALERT_KEY_COL_DESC_4</evt:DescriptionNLSID>
</evt:Attrib>
<evt:Attrib Name="key_column_5" DataType="STRING" isReferenced="true"
  NLSID="ALERT_KEY_COL_NLSID_5" isValueTranslatable="true">
  <evt:DescriptionNLSID>ALERT_KEY_COL_DESC_5</evt:DescriptionNLSID>
</evt:Attrib>
<evt:Attrib Name="key_column_6" DataType="STRING" isReferenced="true"
  NLSID="ALERT_KEY_COL_NLSID_6" isValueTranslatable="true">
  <evt:DescriptionNLSID>ALERT_KEY_COL_DESC_6</evt:DescriptionNLSID>
</evt:Attrib>
<evt:Attrib Name="key_column_7" DataType="STRING" isReferenced="true"
  NLSID="ALERT_KEY_COL_NLSID_7" isValueTranslatable="true">
  <evt:DescriptionNLSID>ALERT_KEY_COL_DESC_7</evt:DescriptionNLSID>
</evt:Attrib>
<evt:Attrib Name="key_column_1_value" DataType="STRING"
isReferenced="true"
  NLSID="KEY_VALUE_PART_NLSID_1" isValueTranslatable="false">
  <evt:DescriptionNLSID>KEY_VALUE_PART_DESC_1</evt:DescriptionNLSID>
</evt:Attrib>
<evt:Attrib Name="key_column_2_value" DataType="STRING"
isReferenced="true"
  NLSID="KEY_VALUE_PART_NLSID_2" isValueTranslatable="false">
  <evt:DescriptionNLSID>KEY_VALUE_PART_DESC_2</evt:DescriptionNLSID>
</evt:Attrib>
<evt:Attrib Name="key_column_3_value" DataType="STRING"
isReferenced="true"
  NLSID="KEY_VALUE_PART_NLSID_3" isValueTranslatable="false">
  <evt:DescriptionNLSID>KEY_VALUE_PART_DESC_3</evt:DescriptionNLSID>
</evt:Attrib>
<evt:Attrib Name="key_column_4_value" DataType="STRING"
isReferenced="true"
  NLSID="KEY_VALUE_PART_NLSID_4" isValueTranslatable="false">
  <evt:DescriptionNLSID>KEY_VALUE_PART_DESC_4</evt:DescriptionNLSID>
</evt:Attrib>
<evt:Attrib Name="key_column_5_value" DataType="STRING"
isReferenced="true"
  NLSID="KEY_VALUE_PART_NLSID_5" isValueTranslatable="false">
  <evt:DescriptionNLSID>KEY_VALUE_PART_DESC_5</evt:DescriptionNLSID>
</evt:Attrib>
<evt:Attrib Name="key_column_6_value" DataType="STRING"
isReferenced="true"
  NLSID="KEY_VALUE_PART_NLSID_6" isValueTranslatable="false">
  <evt:DescriptionNLSID>KEY_VALUE_PART_DESC_6</evt:DescriptionNLSID>
</evt:Attrib>
<evt:Attrib Name="key_column_7_value" DataType="STRING"
isReferenced="true"
  NLSID="KEY_VALUE_PART_NLSID_7" isValueTranslatable="false">
  <evt:DescriptionNLSID>KEY_VALUE_PART_DESC_7</evt:DescriptionNLSID>
</evt:Attrib>
<evt:Attrib Name="metric_type" DataType="NUMBER" isReferenced="true"
  NLSID="METRIC_TYPE" isValueTranslatable="false">
  <evt:DescriptionNLSID>METRIC_TYPE_DESC</evt:DescriptionNLSID>
</evt:Attrib>
<evt:Attrib Name="num_keys" DataType="NUMBER" isReferenced="true"

```



```

        NLSID="NUM_KEYS" isValueTranslatable="false">
        <evt:DescriptionNLSID>NUM_KEYS_DESC</evt:DescriptionNLSID>
    </evt:Attrib>
    <evt:Attrib Name="unit" DataType="STRING" isReferenced="true"
        NLSID="UNIT_NLSID" isValueTranslatable="true">
        <evt:DescriptionNLSID>UNIT_DESC</evt:DescriptionNLSID>
    </evt:Attrib>
    <evt:Attrib Name="is_thresholdable" DataType="NUMBER" isReferenced="true"
        NLSID="IS_THRESHOLDABLE" isValueTranslatable="false">
        <evt:DescriptionNLSID>IS_THRESHOLDABLE_DESC</evt:DescriptionNLSID>
    </evt:Attrib>
    <evt:Attrib Name="is_remote" DataType="NUMBER" isReferenced="true"
        NLSID="IS_REMOTE" isValueTranslatable="false">
        <evt:DescriptionNLSID>IS_REMOTE_DESC</evt:DescriptionNLSID>
    </evt:Attrib>
    <evt:Attrib Name="is_long_running" DataType="NUMBER" isReferenced="true"
        NLSID="IS_LONG_RUNNING" isValueTranslatable="false">
        <evt:DescriptionNLSID>IS_LONG_RUNNING_DESC</evt:DescriptionNLSID>
    </evt:Attrib>
    <evt:Attrib Name="is_udm" DataType="NUMBER" isReferenced="true"
        NLSID="IS_UDM" isValueTranslatable="false">
        <evt:DescriptionNLSID>IS_UDM_DESC</evt:DescriptionNLSID>
    </evt:Attrib>
    <evt:Attrib Name="is_metric_extension" DataType="NUMBER"
isReferenced="true"
        NLSID="IS_METRIC_EXTENSION" isValueTranslatable="false">
        <evt:DescriptionNLSID>IS_METRIC_EXTENSION_DESC</evt:DescriptionNLSID>
    </evt:Attrib>
</evt:AttributeDef>

    <evt:RefAttribSource><![CDATA[sysman.em_severity.get_metric_alert_ref_
attrs]]></evt:RefAttribSource>

    <evt:SignatureAttribs>
        <evt:SignaturePart>metric_guid</evt:SignaturePart>
        <evt:SignaturePart>key_value</evt:SignaturePart>
    </evt:SignatureAttribs>

    <evt:RuleAttribs>
        <evt:RuleAttrib Name="metric_group"/>
        <evt:RuleAttrib Name="metric_column"/>
        <evt:RuleAttrib Name="key_value"/>
        <evt:RuleAttrib Name="key_column_1_value"/>
        <evt:RuleAttrib Name="key_column_2_value"/>
        <evt:RuleAttrib Name="key_column_3_value"/>
        <evt:RuleAttrib Name="key_column_4_value"/>
        <evt:RuleAttrib Name="key_column_5_value"/>
        <evt:RuleAttrib Name="key_column_6_value"/>
        <evt:RuleAttrib Name="key_column_7_value"/>
    </evt:RuleAttribs>

    <evt:NotifAttribs>
        <evt:NotifAttrib Name="metric_guid"/>
        <evt:NotifAttrib Name="severity_guid"/>
        <evt:NotifAttrib Name="cycle_guid"/>
        <evt:NotifAttrib Name="coll_name"/>
        <evt:NotifAttrib Name="metric_group"/>
        <evt:NotifAttrib Name="metric_column"/>
        <evt:NotifAttrib Name="metric_description"/>
        <evt:NotifAttrib Name="value"/>

```

```

<evt:NotifAttrib Name="key_value" />
<evt:NotifAttrib Name="key_column_1" />
<evt:NotifAttrib Name="key_column_1_value" />
<evt:NotifAttrib Name="key_column_2" />
<evt:NotifAttrib Name="key_column_2_value" />
<evt:NotifAttrib Name="key_column_3" />
<evt:NotifAttrib Name="key_column_3_value" />
<evt:NotifAttrib Name="key_column_4" />
<evt:NotifAttrib Name="key_column_4_value" />
<evt:NotifAttrib Name="key_column_5" />
<evt:NotifAttrib Name="key_column_5_value" />
<evt:NotifAttrib Name="key_column_6" />
<evt:NotifAttrib Name="key_column_6_value" />
<evt:NotifAttrib Name="key_column_7" />
<evt:NotifAttrib Name="key_column_7_value" />
<evt:NotifAttrib Name="num_keys" />
</evt:NotifAttribs>

<evt:Severities>
  <evt:Severity>CRITICAL</evt:Severity>
  <evt:Severity>WARNING</evt:Severity>
</evt:Severities>

</evt:EventClass>

```

Table 9–4 Event Class Attributes for Metric Alerts

Attribute	Description
KEY_VALUE_DESC	Monitored object for the metric corresponding to the Metric Alert event
VALUE_DESC	Value of the metric when the event triggered
METRIC_GROUP_DESC	The name of the metric
METRIC_COLUMN_DESC	The name of the metric column
KEY_COLUMN_1_VALUE	Value of Key Column 1
KEY_COLUMN_2_VALUE	Value of Key Column 2
KEY_COLUMN_3_VALUE	Value of Key Column 3
KEY_COLUMN_4_VALUE	Value of Key Column 4
KEY_COLUMN_5_VALUE	Value of Key Column 5
KEY_COLUMN_6_VALUE	Value of Key Column 6
KEY_COLUMN_7_VALUE	Value of Key Column 7
IS_METRIC_EXTENSION_DESC	Flag to indicate if the metric is metric extension

9.4 Adding Customized Details About the Event

The Incident Details region shows information about the event. It consists of system attributes (such as the message, target name, and when the event was reported) and the class attributes. You can customize the name-value pairs for the class attributes.

Through the event-specific customization XML, you can choose which attributes to show, such as the labels for the name part, and whether you require a link under the

value. For more information, see [Section 9.3, "Creating Event-Specific Customization XML"](#).

Example 9–4 Constructing a Name-Value Pair

```
<evt:DetailUI>
  <evt:UIAttributeList>
    <evt:UIAttrib Name="metric_name">
      <evt:URL PageType="sdkcore-dummy-published-page-id">
        <evt:URLParam Name="target" Value="^TARGET:sys_target_name^"/>
        <evt:URLParam Name="type" Value="^TARGET:sys_target_type^"/>
        <evt:URLParam Name="metric" Value="^metric_name^"/>
        <evt:URLParam Name="metricColumn" Value="^metric_column^"/>
        <evt:URLParam Name="ctxType" Value="Hosts"/>
      </evt:URL>
    </evt:UIAttrib>
  </evt:UIAttributeList>
</evt:DetailUI>
```

[Example 9–4](#) constructs a name-value pair under the Incident Details region. The name is the translated value of `metric_name`, which is an event class attribute. The value part is the value of `metric_name`, with a link to the `METRIC_DETAILS` page with specified URL parameters.

Note: For the `evt:URL` tag, you must use an EDK published page id as the `pageType`. At design-time, you cannot validate the link navigation so if you are using this API, then you must verify that the link works correctly on the Incident Manager UI.

The URL parameters in the links can be:

- Event attributes: Must be enclosed in carets (^). For example, `^metric_name^`.
For information about event attributes, see [Section 9.3.2, "About Events"](#).
- Target attributes: Must be prefixed with `TARGET`. For example, `^TARGET:sys_target_type^`.

Possible target attributes:

- `sys_target_name`
- `sys_target_type`
- `sys_target_owner`
- `sys_target_version`
- `sys_target_lifecycle_status`
- `sys_target_guid`

For more information about these attributes, see [Table 9–2](#).

- Event context attributes: Must be prefixed with `EVENT`. For example, `^EVENT:evt_context_attrib_name^`.

Event context attributes enable the event publisher to provide additional event information to event attributes and are defined as name-value pairs.

- Constants: Must be specified as literal strings. For example, `byDay`.

Note: You can customize the Incident Details region to include class-specific attributes only.

9.5 Providing Content in the Guided Resolution Region

You can make the following customizations to the Guided Resolution region:

1. Customize Repair and Diagnostic links (these links can be added or removed)
2. Add recommendations to the Guided Resolution region
3. Specify a default search phrase for My Oracle Support Knowledge
4. Add areas with text to the Guided Resolution region

The Guided Resolution region provides links to relevant Enterprise Manager pages to help users debug and resolve issues. These context-sensitive links are grouped into multiple areas based on the nature of content in the destination page of the link. The following areas are displayed only if they have content.

- **Diagnose:** This area contains links that can help users diagnose the issue. For example, for an event based on a target, the Diagnose area contains a link called **View topology** that drills down to the Topology Viewer. You can add or remove links in this area using the `evt:DiagLinks` tag (see [Example 9-5](#)).
- **Repair:** This area contains links that can help users resolve the issue. For example, for metric alerts with thresholds, this area might contain a link to **Edit Thresholds**. You can add or remove links to and from this area using the `evt:ActionLinks` tag. This area displays for open events only.
- **Recommendation:** This area contains text content drawn from the metric advisory, if available. You can customize the content of this area using the `evt:Recommendation` tag (see [Example 9-6](#)).
- **Any additional area:** This area contains text content drawn from the metric advisory, if available. You can customize the content of this section using the `evt:Sections` tag ([Example 9-8](#)).

[Example 9-5](#) provides an example of adding a link to the Diagnostics list in the Guided Resolution region using XML:

Example 9-5 Adding a Link to the Diagnostics Subsection

```
<evt:DiagLinks>
  <evt:Add>
    <evt:Link LinkID="diagLink1_example">
      <evt:Label>
        <evt:LocalizedLabel DefaultLabel="Edit Thresholds"
          NLSID="EDIT_METRIC_THRESHOLDS"
          ResourceBundle="oracle.sysman.resources.MntrResourceBundle"/>
      </evt:Label>
      <evt:URL PageType="sdkcore-published-page-id-for-edit-threshold">
        <evt:URLParam Name="target" Value="^[TARGET:sys_target_name^"/>
        <evt:URLParam Name="type" Value="^[TARGET:sys_target_type^"/>
        <evt:URLParam Name="event" Value="doEditTreshhold"/>
        <evt:URLParam Name="metric" Value="^[metric_name^"/>
        <evt:URLParam Name="collName" Value="^[coll_name^"/>
        <evt:URLParam Name="keyValue" Value="^[key_value^"/>
        <evt:URLParam Name="metricColumn" Value="^[column^"/>
      </evt:URL>
    </evt:Link>
```

```
</evt:Add>
</evt:DiagLinks>
```

In [Example 9–5](#), the link text is derived from the `evt:Label` specification and the URL is derived from `evt:URL`.

The URL parameters in the links can be:

- Event attributes: Must be enclosed in carets (^). For example, `^metric_name^`. For information about event attributes, see [Section 9.3.2, "About Events"](#).
- Target attributes: Must be prefixed with TARGET. For example, `^TARGET:sys_target_type^`.

Possible target attributes:

- `sys_target_name`
- `sys_target_type`
- `sys_target_owner`
- `sys_target_version`
- `sys_target_lifecycle_status`
- `sys_target_guid`

For more information about these attributes, see [Table 9–2](#).

- Event context attributes: Must be prefixed with EVENT. For example, `^EVENT:evt_context_attr_name^`.

Event context attributes enable the event publisher to provide additional event information to event attributes and are defined as name-value pairs.

- Constants: Must be specified as literal strings. For example, `byDay`.

9.5.1 Adding Recommendations using XML

The Actions list in the Guided Resolution region enables you to provide some text describing recommended steps that users can follow to diagnose or resolve the issue. Use any of the Label tags to specify the recommendations.

Example 9–6 Adding Recommendations

```
<evt:GuidedResolutionDetails>
  <evt:GuidedResolution>
    <evt:Recommendation ID="reco_foo">
      <evt:Label>
        <evt:LocalizedLabel
          DefaultLabel="Recommendation for my event class"
          NLSID="MY_EVENT_CLASS_NLSID"
          ResourceBundle="oracle.sysman.MyResourceBundle"/>
        </evt:Label>
      </evt:Recommendation>
    </evt:GuidedResolution>
  </evt:GuidedResolutionDetails>
```

You can use other variants of the `evt:Label` tag to construct more complex Recommendations. For example, to substitute the value of an event attribute, such as `alertAction` as the recommendation, use [Example 9–7](#):

Example 9-7 Adding a Complex Recommendation

```
<evt:GuidedResolutionDetails>
  <evt:GuidedResolution>

    <evt:Recommendation ID="reco_unique_id">
      <evt:Label>
        <evt:AttributeValue Name="alertAction" />
      </evt:Label>
    </evt:Recommendation>

  </evt:GuidedResolution>
</evt:GuidedResolutionDetails>
```

9.5.2 Customizing Sections

You can display textual information by adding sections to the Guided Resolution region. For example, while prioritizing an incident about an out-of-the-box Configuration Standard, it might be useful to see the rationale explaining why that Configuration Standard was added. Each added area has a header and some textual content.

To add sections to the Guided Resolution region, you can add specifications to the customization XML similar to [Example 9-8](#):

Example 9-8 Adding Customized Areas

```
<evt:Sections>
  <evt:Add>
    <evt:Section ID = "section_eventclass_1">
      <evt:SectionHeader>
        <evt:Label>
          <evt:LocalizedLabel DefaultLabel="Section Hdr 1"
            NLSID="TRANSLATION_ID"
            ResourceBundle="oracle.sysman.MyResourceBundle" />
        </evt:Label>
      </evt:SectionHeader>
      <evt:SectionText>
        <evt:Label>
          <evt:LocalizedLabel DefaultLabel="Section for my event class"
            NLSID="MY_EVENT_CLASS_NLSID"
            ResourceBundle="oracle.sysman.MyResourceBundle" />
        </evt:Label>
      </evt:SectionText>
    </evt:Section>
  </evt:Add>
</evt:Sections>
```

9.6 Defining a Search String for My Oracle Support Knowledge

The customization framework provides a default search phrase in the following order of preference:

1. Customized value from the plug-in developer
2. ORA error found in the Event Summary
3. Event Summary

The search string can be specified explicitly by using XML as shown in [Example 9–9](#). This search string searches for a metric alert indicating high CPU usage for the plug-in component.

Example 9–9 Defining a Search String

```
<evt:GuidedResolutionDetails>
  <evt:GuidedResolution>
    <evt:SearchPhrase>High CPU Utilization </evt:SearchPhrase>
  </evt:GuidedResolution>
</evt:GuidedResolutionDetails>
```

9.7 Defining Conditions for Customization

To select the events for which customizations are to be applied, you can define conditions using the attributes of an event. To define conditions using attributes from the event payload (for example, system attributes such as target name, target type, or event-class specific attributes such as metric_name for metric alerts, or event-context attributes), use XML.

Conditions can be defined under the evt:ConditionDetails tag. You can specify various event attributes here and they are joined together implicitly using an AND condition.

Note: You can define one condition only under the evt:ConditionDetails tag.

Example 9–10 Defining a Condition

```
<evt:CustomUI AppliesTo ="EVENT" EventClass ="metric_alert" TargetType =host">
  <evt:ConditionDetails>
    <evt:Condition>
      <evt:Attrib Name="metric_name" Value="Load"/>
      <evt:Attrib Name="metricColumn" Value="cpuUtil"/>

    </evt:Condition>
  </evt:ConditionDetails>
  ...
</evt:CustomUI>
```

The following operators are supported to define conditions:

- EQ: Equals. This is the default operator
- NE: Not equals
- ISNULL: Is null
- ISNOTNULL: Is not null
- CONTAINS: Can be contained in the string (as a substring)
- BEGINSWITH: Begins with (for example, the event name begins with Tablespace)
- IN: In a predefined set of values (separated by pre-defined delimiter comma (,))
- NOT_IN: Not in a predefined set of values. Use for exclusion

9.8 Registering Customizations

The event-specific customization XML files are located in the `$PLUGIN_ORACLE_HOME/sysman/metadata/events/custmzn` directory. In the shipped version of the product, these XMLs are registered as part of the plug-in installation.

If you are creating the XMLs for the first time in a view that is already set up or to test changes, use the Metadata Registration Service (MRS) to register XML for the event class. For more information about the MRS, see [Section 13.7, "Updating Deployed Metadata Files Using the Metadata Registration Service \(MRS\)"](#).

```
emctl register oms metadata -service eventSpecificCustmzn -file XML filename
-pluginId plugin_name -sysman_pwd sysman -debug
```

For example:

```
emctl register oms metadata -service eventSpecificCustmzn -file metric_alert_host_
load.xml -pluginId acme.demo.hostsample -sysman_pwd sysman -debug
```

You must restart the Oracle Management Service (OMS) after registering the event-specific customization XML, using the `emctl` command from the OMS home directory (`OMS_HOME`).

```
OMS_HOME>emctl stop oms
OMS_HOME>emctl start oms
```

You might encounter the following errors when you register your event-specific customization XML:

- Syntax error in the XML
For information about the correct syntax and an example of the XML, see [Section 9.3, "Creating Event-Specific Customization XML"](#).
- Incorrect values for attribute names
For information about attribute names, see [Section 9.3.2, "About Events"](#).
- Incorrect credentials
Incorrect credentials will not allow you to connect to the Management Repository. Ensure that you are using authorized credentials.

9.9 Testing Incident Manager After Customization

Test the Incident Manager UI by publishing an event that matches the condition and then make sure that there is an incident created for it:

1. To access Incident Manager, from the **Enterprise** menu, select **Monitoring**, then select **Incident Manager**.

The **Incident Manager: My open incidents and problems** page appears.

2. From **Views**, select either of the following:
 - **All open incidents** to find the incident
 - **Events without incidents** if you did not create an incident

Using Derived Associations

Effective management of IT infrastructure requires knowledge of the relationships between IT entities. Best practices such as those described by ITIL (Information Technology Infrastructure Library) rely on capturing and using such relationships.

Enterprise Manager Cloud Control 12c extends the kinds of relationships being supported and adds a declarative mechanism by which these relationships can be maintained. It also determines the membership of entities in a system based on relationships. Based on accurate relationships, various Enterprise Manager applications and components can support customer uses such as:

- Dependency analysis.
For instance, to understand the impact (to applications and infrastructure) of shutting down a host.
- Topology viewer.
- Change management.
For instance, tracking the source of cloned databases.
- End-to-end performance analysis, in which interdependencies between application components must be known in order to analyze and isolate issues.
- Change tracking of relationships, such as changes in the way VM resources are allocated.

This chapter covers the following:

- [Introduction to Derived Associations](#)
- [Understanding Enterprise Manager Association Concepts](#)
- [About Association Derivation Rules Management](#)
- [Ensuring Performance](#)
- [Using Overlapping Associations](#)
- [Frequently Asked Questions](#)

10.1 Introduction to Derived Associations

As an plug-in developer, you are responsible for defining those association types that apply to your managed entity types and for verifying that the correct associations (association instances) are present.

A manageable entity is an entity that Enterprise Manager is capable of managing. This implies that the entity is exposed in some form to end users in the Cloud Control application, and has well-defined attributes and semantics.

As a plug-in developer, you are responsible for the following steps with regard to derived associations:

1. Identify all associations that need to be represented for any managed entities (MEs) that you own.

This generally includes any containment or dependency associations between an ME you own and any other MEs. For each kind of association identified, you may need to coordinate with the owner of the related ME type to determine who should be responsible for assuring that association instances of that type are kept up to date. Some associations (in particular, `hosted_by` and `managed_by`) are automatically maintained by Enterprise Manager, so association derivation rules should not be used for these.

2. Understand the set of out-of-box association types that are shipped with Enterprise Manager and ensure the use of the most appropriate type.

For more information, see [Section 10.2.1, "About Out-of-Box Association Types"](#).

3. Ensure that association derivation rules are used to (declaratively) describe the associations that are to exist based on configuration data that resides in the repository.

Rules are triggered by configuration collections (where target property changes are also treated as a configuration collection).

For more information, see [Section 10.3.1, "Using Association Derivation Rules Syntax and Semantics"](#).

4. You need to coordinate with the owners of other ME's regarding association maintenance, as associations with your ME types often involve other plug-in's ME types.

Decide which plug-in will package the rules. The plug-in that owns the rule must ensure that it specifies all other needed plug-ins as prerequisites to ensure that all target types and their ECM metadata is present prior to rule installation.

Advanced activation expressions, as described in [Section 10.3.10, "Understanding Activation Expressions"](#) can be used if it is not possible to assure all needed target types are present. Rule triggers should reside in plug-ins that define target types specified by the triggers. However, if target types are known to exist before the plug-in installation, the triggers can reside along with the rule.

10.1.1 Assumptions and Prerequisites

This chapter assumes you are familiar with the following:

- Association types in general, association type hierarchy, concepts of allowed pairs of manageable entity types for association types, forward and concrete (versus abstract) association type, and the semantics of the Enterprise Manager out-of-box association types.
- Target model, target properties, and target components.
- Enterprise Configuration Management configuration collections, including treatment of target properties as configuration data.
- Plug-in development overview, including how to package a plug-in and its XML files.

- Topology viewer (to view your associations).

10.2 Understanding Enterprise Manager Association Concepts

In Enterprise Manager, the concept of a relationship is internally referred to as an association. An association (association instance) represents a relationship between two managed entities and specifies three values, namely, source, destination, and association type. For instance, in “database1 exposed_by listener1”, database1 is the source, listener1 is the destination, and “exposed_by” is the association type.

This section describes association derivation rules, which provide a concise declarative means of defining association types. Association derivation (so called because the existence of associations is derived from collected data) provides a mechanism by which developers can cause association instances to be created and removed based on data collected from a target.

The association derivations are based on the data collected using configuration collections and present in the Enterprise Manager repository. The association derivation mechanism allows you to keep the association consistent with the collected configuration data and to determine associations centrally based on all known data (instead of being done by agent logic, which has access to less data).

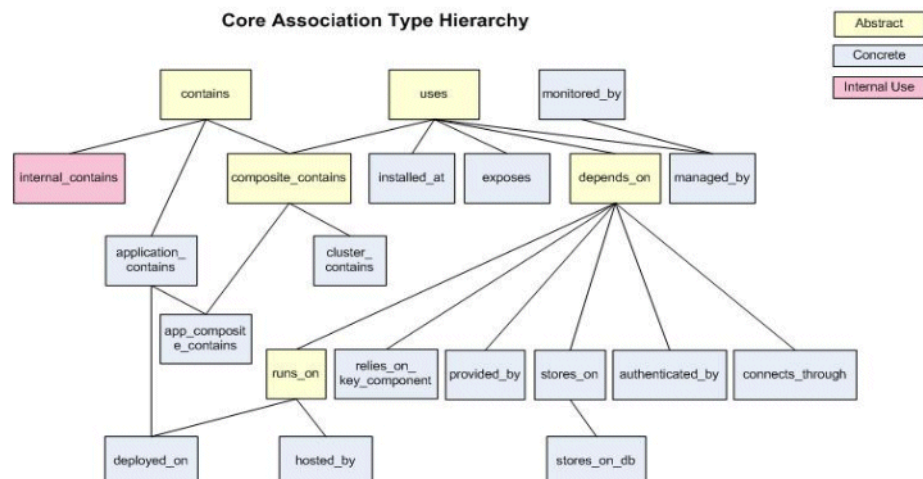
10.2.1 About Out-of-Box Association Types

Enterprise Manager provides a common set of association types that should meet the needs of most plug-in developers and you are encouraged to become familiar with these association types and use them if applicable.

The cardinality specifies the cardinality for the overall association type. An allowed_pairs (constraint) should not specify conflict cardinality, but may specify more specific cardinality. An abstract association type can not have association instance created for it.

The following diagram shows the core association type hierarchy. For more information on out-of-box association types, see [Appendix A, "Out-of-Box Associations"](#)

Figure 10–1 Core Association Type Hierarchy



10.2.2 Using Association Derivation

To use association derivation, complete the following:

1. Specify the logic to run after the collection of target configuration.

The logic derives a set of association instances in the form of triples that specify the source managed entity GUID, association type, and destination managed entity GUID. For instance, the association derivation logic for targets of type `oracle_listener` could return triples that represent associations between the listener and each database for which it listens.

2. Create and run a `SELECT` statement that contains the logic used to derive the triples.

Each returned row contains association type, source, and destination columns and represents an association that should exist.

3. Register the derivation logic against an Enterprise Configuration Management snapshot type.

After every snapshot collection, the registered logic is invoked. Input to the logic is the GUID of the target for which the data was collected.

When the association derivation logic for snapshot S of target T is executed, the derived associations replace the previously derived associations for snapshot S of target T. For example, if associations A1 and A2 were collected yesterday and only A1 is collected today, then A2 is effectively deleted.

10.2.3 About Automated Discovery and Promotion of Associations

One option for adding associations to Enterprise Manager is to provide a discovery script which discovers targets and the associations between them, and the discovery script is then scheduled to run on a selected set of agents by the end-user. The targets and associations discovered by this type of script are automatically promoted, that is they are automatically added to Enterprise Manager. This approach is useful for associations that are between targets that are managed by your plug-in and therefore the specific target identification is known (that is you create the targets on both ends of the association). If these associations are to other targets not included in your plug-in, then typically a derived association rule is used to specify how to locate the "external" target.

A guided discovery process may be used if some interaction with the end-user or administrator is necessary to filter the information discovered by the script, or if some amount of post-processing is necessary to compare it to other information already known to Enterprise Manager.

10.2.4 Understanding Association Creation During Guided Discovery

This approach is similar to the automated discovery approach described in the previous section in that you provide a discovery script that can be run by an Enterprise Manager agent. That discovery script may return any number of related targets and the associations between them. The difference is that in the guided discovery case, you provide a user interface that the end-user interacts with to drive the execution of the discovery script and then process the results returned from it. This processing takes the output from the discovery script and may further filter it or present it to the end-user to allow them to add important information to it.

Guided discovery may also interact with the Enterprise Manager system using target services to obtain information about targets already known to Enterprise Manager to

perform incremental updates to the topology of targets discovered. This approach is also used for cases where the associations to be created are between targets that are managed by your plug-in and therefore the specific target identification is known. That is you create the targets on both ends of the association, but some additional intervention is needed before those associations are added to Enterprise Manager.

10.2.5 Using Associations Derived from a System Stencil

This approach is used solely for creating system membership associations between a system target and its members. The system target and its members are typically all part of a single plug-in, as you must have knowledge of the types of associations that exist between the system target and its members in order to form the system topology.

The system stencil defines the set of association paths that should be considered when forming the system membership. In this way, the plug-in can traverse complex association paths to locate targets that should be treated as members of the system. This is important in cases where a system member is not directly associated with the system target by some other "native" association.

If the plug-in does not include a target type that you wish to be treated as a system, then this approach can be ignored.

10.2.6 Associations Derived from Rule

This approach for creating associations is particularly suited to cases where the destination target of the association is not part of the plug-in but is known to be managed by Enterprise Manager. For example, assume that the configuration of your target included a connection to an Oracle database that was used to store information related to your target operation (such as an application store). The configuration of your target knows something about the database that it uses, likely some connection related details such as host-port-sid or host-service.

You would like to represent this association between your target and the database in Enterprise Manager so that if Enterprise Manager is managing the database, the end-user can see this relationship and traverse it to obtain other information about that database and manage it (if appropriate and allowed).

Because you do not know if Enterprise Manager is managing the database and the identifying information you have is not the Enterprise Manager database target name, but instead the connect information, you can construct a derivation rule that maps the connection information in your target's configuration to that of a database in Enterprise Manager.

This approach is very useful for cases where you wish to construct this type of association between a target that is part of your plug-in and some external target, particularly some Enterprise Manager stack component like Oracle Fusion Middleware or the Oracle Database.

10.3 About Association Derivation Rules Management

Enterprise Manager Cloud Control 12c extends the use of associations by Enterprise Manager components and enhances the overall association framework. It introduces new consumers of associations, including the topology viewer.

Association framework enhancements include the treatment of associations as configuration data. Enterprise Configuration Management features such as change tracking and saved snapshots now apply to associations as well as to traditional

configuration data. Associations can now specify source and destination target components, as well as target GUIDs.

The following sections provide detailed instructions on the use and management of derivation rules:

- [Using Association Derivation Rules Syntax and Semantics](#)
- [Understanding XML Metadata File Syntax and Semantics](#)
- [Using Rule Semantics](#)
- [Maintaining Performance](#)
- [About Regular Query and Trigger Patterns](#)
- [Diagnosing Issues](#)
- [Useful Examples](#)
- [Applying the Mechanical Steps of Integration](#)
- [Understanding Activation Expressions](#)
- [About Debugging](#)

10.3.1 Using Association Derivation Rules Syntax and Semantics

The following sections describe the contents of a rule, including name, query, triggers, and database objects that can be referenced by rule queries.

Name

A rule is identified by a unique rule name that must be unique across all plug-ins. Oracle recommends that you use a suitable prefix to avoid name conflicts. For example, a company symbol or name followed by the plug-in name.

Query

The primary component of a rule is the rule query, which identifies a set of associations. Each row returned by the query represents an association. The SQL must return four columns whose names and types must be:

- `assoc_type`
(VARCHAR2(64)): the association type
- `source_me_guid`
(RAW(16)): a managed entity GUID
- `dest_me_guid`
(RAW(16)): a managed entity GUID
- `derivation_target_guid, derivation_target_guid2, derivation_target_guid3`
(RAW(16))

Often unnecessary, these are one or more optional target GUID columns that identify targets involved in deriving the association (other than the source or destination).

- These cannot be a target component ID, but must be a target GUID.
- Columns should be used in order.

This means that queries returning `derivation_target_guid2` must also return `derivation_target_guid`. Queries returning `derivation_target_guid3` must also return `derivation_target_guid` and `derivation_target_guid2` columns.

The need in some cases for a derivation target guid is illustrated by the case in which the collection for a target determines associations between two other targets. For instance, the collection for a Siebel Enterprise System determines associations between its member targets.

In this case, the derivation target GUID is the target GUID of the Siebel Enterprise System target, but the source and destination are other targets. Similar cases exist for Oracle E-Business Suite and Oracle WebLogic Server, where configuration information is collected from a single source, such as the Oracle WebLogic Server domain admin server, and used to derive associations between the domain members.

Each row returned by the query must specify a valid association instance and must use a concrete (not abstract) forward (not inverse) association type. Valid association instances must specify managed entities that are valid for the specified association type. For example, a `hosted_by` association must specify a destination that is a host target. Inverse association types must not be returned. For example, do not use `host_for`, which is the inverse of `hosted_by` and would be logged as an error.

Note that the rule query returns a repository-wide set of associations, but associations are populated incrementally on behalf of one target at a time. When the rule is evaluated, it is from the perspective of a single target. At evaluation time, the framework wraps the query with an outer query. For example:

```
"SELECT ... FROM <query> WHERE derivation_target_guid = <initiatingTarget> AND ..."
```

Note: You do not need to specify a `DISTINCT` keyword (at the outmost level) in your rule query as the framework will eliminate the duplicates by itself when it wraps your query with its own query.

Query size should not exceed 2000 characters. (It is planned to extend this to 4000 characters in a future release).

DB Objects Referenced by Rule Queries

For security reasons, the `SYSMAN_RO` user will execute your rule query. Therefore, only objects accessible by this user are allowed to be referenced. For objects created outside of your plug-in you can reference views exposed by the Extensibility Development Kit (EDK) at your plug-in level, including those prefixed with `MGMT$`.

For objects created in your plug-in, you can reference `CM$` views auto-generated by the Enterprise Configuration Management framework for your target type collections. You can also reference views prefixed with a `DA_` prefix and packages with invokers rights with a `DA_` prefix.

Your query should not rely on associations unless you ensure that they are present by the time the query is executed (for example, when corresponding triggers fire – see below). For derived associations, the order of executions is not deterministic because the order in which configurations arrive and then corresponding associations are derived is arbitrary. An example of an association that *can* be used is a `hosted_by` association, which appears during target discovery and is not a derived association.

Triggers

Triggers are usually provided in addition to the rule query. A trigger specifies a table that, when changed, may impact the associations returned by the rule query. Generally there are multiple triggers because the rule query often refers to data from multiple tables and because changes to either target's data can affect the association rows returned by the query.

Two triggers may not always be needed as it may be the case that the data for one of the two targets does not change. For instance, an association rule that determines its destination based on (immutable) identity properties of one of the targets is only affected by changes to the source target's configuration. Even in that case, it may still be desirable to specify two triggers. If the destination target can appear after the source and this appearance causes the immediate creation of a new association, the trigger is needed.

A trigger specifies the following:

- A snapshot table

A change to the table (due to upload of new data) will fire the trigger. You should only include tables that affect the set of associations because needless firing of triggers impacts performance. A table is identified by target type, snapshot type, and table name.

- Column ID flag

This indicates whether the source, destination, or a derivation target guid should be used to identify associations affected by the newly uploaded configuration data. In other words, depending on this column value, associations for the source, destination, or a derivation target will be replaced with a new set of associations computed for that target when the trigger table data changes. Possible values include `source`, `destination`, `derivationTarget`, `derivationTarget2`, or `derivationTarget3`.

When the trigger fires, the association derivation framework will effectively replace all currently existing associations where the given target is a source, destination, or derivation (depending on the flag) with newly computed associations. To compute the new set of associations, the rule query is executed with the corresponding column bound to the target id. This simplified explanation assumes that associations only exist because of this single rule and it would be slightly changed for a target components case.

For example, a rule query accesses data from a listener configuration table and a database configuration table and returns associations of the form `<database> exposed by <listener>`. One trigger specifies the database configuration table and a column id flag of source, because a change to the configuration table for a database may affect rows where the database is the source of the association. Similarly, a second trigger specifies the listener configuration table and a column id flag of the destination, because a change to the configuration table for a listener may affect rows where the listener is the destination of the association.

If multiple rules can be triggered for a snapshot table, then the order in which the triggers execute is non-deterministic. This means that the developer cannot make any assumptions about the order.

The table name (TN) specified in a trigger can actually be the name of a base table, view, or synonym. In all cases, the underlying tables of TN are identified. A trigger is created for each such table that is an Enterprise Configuration Management table.

10.3.2 Understanding XML Metadata File Syntax and Semantics

To create or update a rule, you edit an XML metadata file that defines the rule (or set of rules) and then import it into the repository. The metadata import is done when a repository is created or upgraded. It is also done when a plug-in is added, upgraded, or removed.

Schematically, you specify the following information for each rule in the file:

- Name
- Query
- 0 to n triggers with
 - Fully qualified snapshot type (which includes target type).
 - Metric table of that snapshot (view or synonym that refers to such a table)
Normally, this view is used by the rule query.
 - Column flag (source, destination, derivationTarget, derivationTarget2, or derivationTarget3).
 - Optional details (used to specify target property names for triggers based on target properties).

The XML semantics are designed to describe the latest state of a rule and its triggers, no matter what the prior state was. So you only need the latest XML specification of a given rule to know how the rule and its triggers are specified for a plug-in. Moreover, one plugin will not be able to directly affect triggers of another plug-in. However, if a rule is removed, triggers referencing the rule from the other plug-ins will not be useful.

The following outlines the rule query specification semantics:

- A non-empty query implies that you need to add or, if needed, overwrite a prior rule query for a rule that had been registered by the same plug-in.
- Specifying no query implies removing the rule, if the rule query had been registered by the same plugin, and no change otherwise.

Rule Location

Once a rule R with its query is located in a file F within a plug-in P , its corresponding XML Rule element can never be removed from that file or from that plug-in (although its attributes and subelements can be modified). Rule R will always be owned by plug-in P . If the rule does need to be removed, a Rule element with no query sub-element must remain in file F indicating that rule R had been removed. If it is important to move rule R to another file or plug-in, you must rename the rule (to $R2$ for example), remove rule R using the above syntax in file F , and add $R2$ in the new location. This will effectively remove R and create a new rule $R2$.

Plug-in P that owns a rule R should be chosen carefully to ensure that all target types needed by the rule query are present by the time plug-in P is installed. Rule R should rely only on targets of types that are always present in Enterprise Manager (for example, host), targets of types defined by plug-in P , or targets of types defined in plug-ins that are prerequisites of plug-in P . If it is impossible to choose such a plug-in, consider using the advanced feature of activation expressions discussed later.

Trigger Specification Semantics

In terms of trigger specification semantics, you should replace triggers in the same plug-in with a new specified set of triggers. Alternatively, you can just remove any pre-existing triggers if the newly specified set is empty.

Trigger Location

Normally triggers are defined as part of the rule definition in the same plug-in. This way, when the rule changes, corresponding triggers can also change if needed. However, in some cases it is preferable or only possible to place triggers in plug-ins defining the target types of the triggers. For example, a rule that computes association between targets and their corresponding Oracle Home targets cannot list all possible target types and corresponding triggers.

Instead, plug-ins owning target types that want to use the rule specify the triggers for the relevant target types. Therefore, a plug-in that owns an `oracle_ias` target type will have a trigger for this rule with `oracle_ias` listed as the target type in the trigger. Such triggers are changed or removed along with the corresponding plugin.

Once the rule's XML is listed in some file *F* in the plug-in that owns `oracle_ias`, it cannot move to another file (even if it specifies only triggers for the rule). Note that in our example, the rule query itself is not target type specific, as it only depends on the Oracle Home target type and not other specific target types. Therefore, the rule is defined in the plug-in *P* that is installed prior to the plug-ins (such as `oracle_ias` plugin). This way, when the trigger is imported into Enterprise Manager, it finds the rule already present.

The following points should also be taken into consideration:

- You should always have a rule and its triggers specified in at most one file for a given plug-in.

For example, if some of the triggers for a rule (defined in a different plug-in *P1*) are specified in two files for plug-in *P2*, an import of the second file would overwrite the triggers that were specified in the first file. The order of import of the two files is not guaranteed.
- An error will result if a query is specified for a rule that has been registered by a different plug-in.

In other words, plug-ins that have not specified a rule query can only specify a new set of triggers in their context, but cannot overwrite the query. Only one plug-in effectively owns the rule query.
- An error will result if there is a trigger specification for a rule that does not exist or is being removed (by not specifying the rule query).
- To effectively disable all triggers in all plug-ins for a given rule, the plug-in author can just remove the rule, using the syntax mentioned previously. If needed, the author can also create a rule with a different name as a replacement.
- To replace triggers, but not the rule query, in the plug-in that had specified the rule query in a prior release, specify the same query again and a set of new triggers in the next plug-in version.

Note: In terms of performance, specifying textually the same query will result in the best upgrade performance, since the framework will not need to recompute all associations for the query.

The syntax for rule definitions is as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema"
elementFormDefault="qualified" attributeFormDefault="unqualified"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema">
  <xs:simpleType name="YesNo">
    <xs:annotation><xs:documentation>
      Type definition for the Yes/No attribute value.
    </xs:documentation></xs:annotation>
    <xs:restriction base="xs:NMTOKEN">
      <xs:enumeration value="Y"/>
      <xs:enumeration value="N"/>
    </xs:restriction>
  </xs:simpleType>
  <xs:simpleType name="NameDef">
    <xs:restriction base="xs:string">
      <xs:pattern value="[A-Za-z][A-Za-z0-9_]*"/>
    </xs:restriction>
  </xs:simpleType>
  <xs:simpleType name="TriggerKind">
    <xs:restriction base="xs:string">
      <xs:enumeration value="C"/>
      <xs:enumeration value="H"/>
    </xs:restriction>
  </xs:simpleType>
  <xs:simpleType name="ColumnID">
    <xs:restriction base="xs:string">
      <xs:enumeration value="source"/>
      <xs:enumeration value="destination"/>
      <xs:enumeration value="derivationTarget"/>
      <xs:enumeration value="derivationTarget2"/>
      <xs:enumeration value="derivationTarget3"/>
    </xs:restriction>
  </xs:simpleType>
  <xs:complexType name="RuleContentWFlags">
    <xs:simpleContent>
      <xs:extension base="xs:string">
        <xs:attribute name="source_comp" type="YesNo" use="optional">
          <xs:annotation> <xs:documentation>
            Can source entity be a target component? Default: No
          </xs:documentation> </xs:annotation>
        </xs:attribute>
        <xs:attribute name="dest_comp" type="YesNo" use="optional">
          <xs:annotation> <xs:documentation>
            Can destination entity be a target component? Default: No
          </xs:documentation> </xs:annotation>
        </xs:attribute>
      </xs:extension>
    </xs:simpleContent>
  </xs:complexType>
  <xs:complexType name="RuleType">
    <xs:annotation> <xs:documentation>
      Rule definition.
    </xs:documentation> </xs:annotation>
    <xs:sequence minOccurs="0">
      <xs:choice>
        <xs:element name="query" type="RuleContentWFlags" minOccurs="0"
          maxOccurs="1">
          <xs:annotation> <xs:documentation>
            Query that returns 1 row per association. Must return
```

```
        columns named ASSOC_TYPE, SOURCE_ME_GUID, DEST_ME_GUID, and
        optionally, one or more of DERIVATION_TARGET_GUID,
        DERIVATION_TARGET_GUID2, DERIVATION_TARGET_GUID3.
        Returning DERIVATION_TARGET_GUID[N] column
        implies the query also returns DERIVATION_TARGET_GUID and all
        DERIVATION_TARGET_GUID[K] for all K between 2 and N.
    </xs:documentation> </xs:annotation>
</xs:element>
</xs:choice>
<xs:element name="trigger" minOccurs="0" maxOccurs="unbounded">
  <xs:complexType>
    <xs:sequence>
      <xs:element name="targetType" type="xs:string" minOccurs="1"
        maxOccurs="1"/>
      <xs:element name="snapshotType" type="xs:string" minOccurs="0"
        maxOccurs="1"/>
      <xs:element name="table" type="xs:string" minOccurs="0"
        maxOccurs="1">
        <xs:annotation> <xs:documentation>
          Name of an ECM table or more likely a view whose query
          relies on ECM snapshot table(s). The table(s), when
          uploaded,
          should trigger evaluation of the rule. (The fully
          qualified name includes target type and snapshot
          type.)
        </xs:documentation> </xs:annotation>
      </xs:element>
      <xs:element name="idColumn" type="ColumnID" minOccurs="1"
        maxOccurs="1">
        <xs:annotation> <xs:documentation>
          Indicates whether source, destination, or a derivation
          target
          should be used to identify associations affected by the
          newly
          uploaded configuration data. In other words, depending on
          this
          column value, associations for the source, destination, or
          a derivation target will be replaced with new set of
          associations computed for that target, when the trigger
          table data changes.
          ColumnID type definition contains allowed values.
        </xs:documentation> </xs:annotation>
      </xs:element>
      <xs:element name="details" type="xs:string" minOccurs="0"
        maxOccurs="1">
        <xs:annotation> <xs:documentation>
          Additional details for the trigger. Currently used for
          target properties table, in which case, it contains comma
          separated list of property names that should fire the
          trigger. Absence
          of property names indicates that any property change would
          fire the trigger (for the given target type).
          Note: white space is ignored.
        </xs:documentation> </xs:annotation>
      </xs:element>
    </xs:sequence>
    <xs:attribute name="kind" type="TriggerKind" use="optional">
      <xs:annotation> <xs:documentation>
        Kind of the trigger. "C" (configuration load trigger) by
        default.
      </xs:documentation> </xs:annotation>
    </xs:attribute>
  </xs:complexType>
</xs:element>
```

```

        Other allowed value: "H" (host change trigger)
      </xs:documentation> </xs:annotation>
    </xs:attribute>
  </xs:complexType>
</xs:element>
</xs:sequence>
<xs:attribute name="name" type="NameDef" use="required">
  <xs:annotation> <xs:documentation>
    Name of rule, which must be unique. Recommendation: Use a
    prefix that identifies your plugin.
  </xs:documentation> </xs:annotation>
</xs:attribute>
<xs:attribute name="activation_expr" type="xs:string" use="optional">
  <xs:annotation> <xs:documentation>
    Optional activation expression. If not present or empty, implies
    that the rule is always active. Else, the value is a Boolean
    activation expression which must produce true if and only if
    the rule should be active.
  </xs:documentation> </xs:annotation>

```

The expression can use (case insensitive) "AND", "OR", and parenthesis. Operands of the expression are target types. Each occurrence of a target type evaluates to "true" if and only if the target type is present in EM.

Note that a number of target types do not need to be listed in the expression because they are always going to be present whenever the rule is installed and present in EM installation. These include:

- target types installed with the plugin where the rule resides (i.e. target types in the plugin which owns the rule)
- target types in other plugins on which the plugin owning the rule depends
- target types always installed with EM (like host)

Thus, in many cases, if this option is used, a single target type, as in example 1 below, may suffice.

Examples:

(1) "oracle_ovm"

This simple expression implies that the rule should be active only if oracle_ovm target type is installed at EM (in addition to any other target types that are already known to be present when this rule is installed).

(2) "oracle_ovm and (oracle_oam_cluster or oracle_oim_cluster)"

This expression implies that the rule should be active only if oracle_ovm target type is present and either oracle_oam_cluster or oracle_oim_cluster is also present.

```

  </xs:documentation> </xs:annotation>
</xs:attribute>
</xs:complexType>
<xs:element name="Rules">
  <xs:complexType>
    <xs:sequence minOccurs="1" maxOccurs="unbounded">
      <xs:element name="Rule" type="RuleType" />
    </xs:sequence>
  </xs:complexType>
</xs:element>
</xs:schema>

```

for example:

```
<?xml version="1.0" encoding="UTF-8"?>
<Rules xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
  <Rule name="ora_listensFor">
    <query>
      SELECT ...
      fill in query
    </query>
    <trigger>
      <targetType>oracle_database</targetType>
      <snapshotType>db_config</snapshotType>
      <table>CM$DB_CONFIG_TABLE</table>
      <idColumn>destination</idColumn>
    </trigger>
  </Rule>
</Rules>
```

10.3.3 Using Rule Semantics

The following algorithm depicts a simplified form of the semantics for rule R, where the trigger specifies a query Q and a flag FC that corresponds to a column name of either source, destination, or derivation target GUID.

When updates to a snapshot table are uploaded for some target *t* with a GUID of *t_guid*, for each trigger that specifies the modified table, execute these statements:

```
DELETE FROM MGMT_ASSOC_INSTANCES
WHERE <FC> = t_guid AND RULE_ID = R

INSERT INTO MGMT_ASSOC_INSTANCES
(SELECT a.*, R
 FROM (<Q>) a
 WHERE <FC> = t_guid )
```

The actual implementation differs from the above example for the following reasons:

- The actual implementation will not delete and then add the same association as this would be inefficient.

Rather it will compare the current and new set, making changes only where needed. Moreover, an actual delete statement also deletes any associations that specify a target component of the target.

- It is possible that an association may be asserted by more than one origin
This is managed using an origins table, whose contents are rolled up into the MGMT_ASSOC_INSTANCES table.

- Validity testing is performed at various points.

For example, to test that the association type is valid for the provided source and destination MEs.

10.3.4 Maintaining Performance

Because the evaluation of derivation rules may be frequent, any poor performance of the rule queries can be problematic. Rule authors must ensure that any needed indexes are present and that they test query performance based on the specific queries that are generated for each trigger.

In particular, testing of the rule query must be done for each trigger because each trigger causes the execution of a different query. Note how rule query return values are bound to a given target globally unique identifier (GUID) depending on your triggers.

You must have indexes that will make use of these bindings. Furthermore, queries should be written in such a way that they would not prevent the push of bindings from outside into your queries.

10.3.5 About Regular Query and Trigger Patterns

The following sections outline the regular patterns you would normally see in queries and triggers. You should check whether your queries and triggers adhere to these patterns and if not, document the reasons why (since such cases normally represent the exceptions from the rules of thumb).

Query Patterns

The following outline common query patterns:

1. The derivation target should be non-null when the ECM configuration of such target is used to derive associations between two other entities.

One known case for the use of derivation targets is associations with a database system. The database instance target provides configuration, while associations are made with the corresponding database system target. In such cases, the database instance must act as a derivation target for the associations to the corresponding database system target.

2. The query must only access objects such as CM\$, other views that access Enterprise Configuration Management data, or views that access target information, such as MGMT_TARGETS.

For more information on the objects that can be accessed, see [Section 10.3.1, "Using Association Derivation Rules Syntax and Semantics"](#).

3. Association types must be forward and concrete.

Trigger patterns

The following outline common trigger patterns:

1. The number of triggers will often be equal to the number of non-target-entity views in the FROM clause. In other words if the mgmt_targets, mgmt_target_properties, mgmt\$target, mgmt\$target_properties, and other such views are disregarded.

Each Enterprise Configuration Management view will correspond to one trigger. One exception is when a table may be triggered from more than one target type (for example, oracle_database and rac_database). In this case, multiple triggers for the same Enterprise Configuration Management view could be supplied.

2. The table name of the trigger must be based on Enterprise Configuration Management metadata tables for the snapshot specified by the trigger's target type and snapshot type.

Normally, it should be one of the objects in the rule query FROM clause (for example, a CM\$ view).

3. A view specified in the trigger is joined (perhaps indirectly) in the query with a target (or target component) entity.

The entity id will be returned as source, destination, or derivation target in the select clause. The idColumn will match this.

For example, association between targets A and B is dependent on a join between the cm\$Aconfig and cm\$Bconfig tables, where data from the cm\$Aconfig table comes from target A and data from the cm\$Bconfig table comes from target B. The trigger for the cm\$Aconfig table will have an idColumn matching target A (for example, source) and the trigger for the cm\$Bconfig table will have an idColumn that matches where target B GUID is returned (for example, destination).

4. The trigger target type must match the target type of the target returned by the query in the column specified by idColumn.

More generally, the target type of the target returned by the rule query could be a subtype or a target component of the trigger target type.

5. If trigger relies on target properties, specific property names should be identified in the details tag.

10.3.6 Diagnosing Issues

To help diagnose issues and understand how associations were derived, the framework records information about how associations were derived when in debug mode. For more information on debug, see [Section 10.3.11, "About Debugging"](#). It also includes additional sanity (error) checking. For instance, one test checks that the derivation target GUID is that of a real and current target.

10.3.7 Useful Examples

While the following examples include the use of target properties to illustrate their proper employment, Oracle does not recommend relying on target properties. Instead, configuration data should be properly modeled using ECM tables. For more information, see [Section 10.7.2, "Are there guidelines for when to use target properties?"](#).

When reviewing these examples, it is helpful to remember the following concepts:

- Every target type has an Enterprise Configuration Management snapshot type called `orcl_tp_config`.

It includes a snapshot table referenced by the `GC$TARGET_PROPERTIES` view, which if needed should be used by the triggers. Current data for target properties can be accessed through the `MGMT_TARGET_PROPERTIES` and `MGMT$TARGET_PROPERTIES` EDK objects.
- Every Enterprise Configuration Management snapshot table will by default have a view for accessing the current configuration data.

Its name will be that of the table, with the prefix `CM$`. Most rule queries will refer to configuration tables through their `CM$` views.

10.3.7.1 Host on a Virtual Machine

A 'deployed_on' association type is used to represent the fact that a host target is deployed on a virtual machine target.

The Query below returns all associations between a host and associated virtual machines based on matching their MAC addresses. Triggers are defined so that they trigger the rule whenever the corresponding configuration view (that includes the MAC address) changes. The rule described below would reside in the plug-in defining

virtual machine (host target type is guaranteed to be present on any EM installation). Both triggers can be included in the rule and would belong to the plug-in defining virtual machine.

```
<Rule name="...">
  <query>
    select 'deployed_on' as assoc_type,
           host.target_guid as source_me_guid,
           guest.cm_target_guid as dest_me_guid
    from mgmt$hw_nic host,
         cm$vt_vm_vnic guest
    where guest.mac_address = host.mac_address_std
  </query>
  <trigger>
    <targetType>host</targetType>
    <snapshotType>ll_host_config</snapshotType>
    <table>MGMT$HW_NIC</table>
    <idColumn>source</idColumn>
  </trigger>
  <trigger>
    <targetType>oracle_vm_guest</targetType>
    <snapshotType>ovm_guest_config</snapshotType>
    <table>CM$VT_VM_VNIC</table>
    <idColumn>destination</idColumn>
  </trigger>
</Rule>
```

10.3.7.2 Target installed_at Oracle Home

The Oracle Home target type includes the `INSTALL_LOCATION` target property that contains the name of the directory in which the Oracle Home resides. For all target types that are installed in Oracle homes, there is an `OracleHome` target property that specifies the same value as `INSTALL_LOCATION`. Whenever a target's `OracleHome` value matches the `INSTALL_LOCATION` value and both targets reside on the same host, an `installed_at` association exists.

Both a target's `OracleHome` and a home's `INSTALL_LOCATION` are subject to change. It is also possible for a target or home to be created that in turn matches up with a home or target. However, the value of a target's host is immutable.

Query

Returns all associations between Oracle Home targets and the targets that are installed in them.

```
<Rule name="...">
  <query>
    select 'installed_at' as assoc_type,
           t.target_guid as source_me_guid,
           o.target_guid as dest_me_guid
    from mgmt_targets t,
         mgmt_targets o,
         mgmt_target_properties tp,
         mgmt_target_properties op
    where o.target_type = 'oracle_home' and
          t.host_name = o.host_name and
          tp.target_guid = t.target_guid and
          tp.property_name = 'OracleHome' and
          op.target_guid = o.target_guid and
          op.property_name = 'INSTALL_LOCATION' and
          tp.property_value = op.property_value
```

```
</query>
  <trigger>
    <targetType>oracle_home</targetType>
    <snapshotType>orcl_tp_config</snapshotType>
    <table>GC$TARGET_PROPERTIES</table>
    <idColumn>destination</idColumn>
    <details>INSTALL_LOCATION</details>
  </trigger>
</Rule>
```

Trigger 2

The following trigger for the same rule would reside in the plug-in that defines oracle_database target type:

```
<Rule name="...">
  <trigger>
    <targetType>oracle_database</targetType>
    <snapshotType>orcl_tp_config</snapshotType>
    <table>GC$TARGET_PROPERTIES</table>
    <idColumn>source</idColumn>
    <details>OracleHome</details>
  </trigger>
</Rule>
```

Trigger 3-n

This is the same as trigger 2 only with another target type that has an OracleHome property. These triggers would reside with plug-ins that define corresponding target types. This trigger has the same characteristics as trigger 2, except it uses a different target type that has an Oracle_Home property.

10.3.7.3 Listener and Database

An `exposed_by` association type is used to represent the fact that a database is exposed by a listener to applications. One way that this association can be created is based on the ports for which the listener is configured.

Query

Returns all associations between a database and listener on the same machine such that the ports match. Both triggers can reside with the rule in the plug-in that defines Oracle database and Oracle listener target types.

```
<Rule name="...">
  <query>
    select 'exposed_by' AS assoc_type,
           oradb.target_guid AS source_me_guid,
           oralsnr_ports.cm_target_guid AS dest_me_guid
    from mgmt_targets oradb,
         mgmt_target_properties oradbprops1,
         mgmt_target_properties oradbprops2,
         cm$listener_ports oralsnr_ports
    where oradb.target_type = 'oracle_database'
           and oradb.target_guid = oradbprops1.target_guid
           and oradbprops1.property_name = 'MachineName'
           and oradbprops1.property_value = oralsnr_ports.machine_name
           and oradbprops1.target_guid = oradbprops2.target_guid
           and oradbprops2.property_name = 'Port'
           and oradbprops2.property_value = oralsnr_ports.listener_port
  </query>
  <trigger>
```

```

        <targetType>oracle_database</targetType>
        <snapshotType>orcl_tp_config</snapshotType>
        <table>GC$TARGET_PROPERTIES</table>
        <idColumn>source</idColumn>
        <details>MachineName</details>
    </trigger>
    <trigger>
        <targetType>oracle_listener</targetType>
        <snapshotType>listener_config</snapshotType>
        <table>CM$LISTENER_PORTS</table>
        <idColumn>destination</idColumn>
    </trigger>
</Rule>

```

10.3.8 Applying the Mechanical Steps of Integration

After you decide on the ME/association model and write the rules, proceed with the implementation as follows:

1. If your rules need Enterprise Configuration Management configuration data not yet present, add new Enterprise Configuration Management metrics or extend the existing ones.

If needed, you should add new Enterprise Configuration Management tables or columns. You must also make sure that the default collection schedule specifies `<Schedule OFFSET_TYPE="INCREMENTAL">`. Failure to do so will delay the loading of the configuration such that, for instance, newly discovered targets may not get associations for thirty minutes or more.

2. The Association types framework has the concept of allowed pairs indicating which target types are allowed to be associated by a given association type. If you are creating associations between ME types that are not listed as allowed pairs for the respective association type, add the needed pairs.
3. Create one or more files to define the association derivation rules. Syntax errors, such as failing to conform to the XSD, are passed through as Java exceptions. You may want to use JDeveloper or another tool to confirm that you have created a valid document.
4. Test the rule files by importing them using the following command:

```
emctl register oms metadata -sysman_pwd sysman -pluginId <your.plugin.id>
-service derivedAssocs -file <fileName>
```

Validity testing is performed, so diagnostics may result.

5. Package the files into your plug-in.

Place them so that they are imported at repository creation or upgrade time in accordance with the conventions defined by the metadata framework. If part of a plug-in, place the files in a location similar to:

```
<stage_dir/plugin_dist>/oms/metadata/derivedAssoc
```

6. Test the derivation rules with cases that exercise every rule trigger that you specified.

One option is to initiate the upload of the Enterprise Configuration Management configuration data and check that the associations are properly established. Alternatively, you can directly call the PL/SQL procedure that will trigger the rules:

```
DECLARE
    temp GC$DERIV_ASSOC_CHANGE_LIST := GC$DERIV_ASSOC_CHANGE_LIST();
BEGIN
    GC$ECM_CONFIG.run_assoc_deriv_rule(
        p_target_guid => hextoraw('CC70BC294B82E7E9A95DFC257CFA6459'), --
        Updated target/ME guid
        p_rule_name => '...', -- your rule name
        p_column_flag => 'D', -- column flag specifying the
        perspective from which to fire the rule. Possible values:
        S|D|T|U|V (implying source,destination, derivation target,
        derivation target 2, or derivation target 3, respectively)
        p_change_list => temp);
    COMMIT;
    -- examine p_change_list if needed
END;
```

Note: Test the performance of your queries for each trigger after the corresponding output of the query has been bound, as described in [Section 10.4, "Ensuring Performance"](#).

Use the import utility to make rule changes and try again.

10.3.9 Special Triggers: Host Name Change Triggers

With this release, a new kind of trigger is supported for when the host name of a target changes. If your query relies on the host_name column of the mgmt_targets table, this trigger can be useful as it will fire when the host_name column changes, for example, upon relocation of the target.

The trigger syntax specifies a “kind” attribute of the “trigger” element with a value of “H” (which stands for “Host change” trigger). Trigger sub-elements will specify target type, which the trigger applies, and idColumn, which identifies the perspective from which to evaluate the rule. Possible values for idColumn are the same as those for regular triggers.

For example:

```
<trigger kind="H">
  <targetType>oracle_database</targetType>
  <idColumn>source</idColumn>
</trigger>
```

This specifies that the rule (which the trigger is part of) has to be reevaluated from the source perspective whenever oracle_database target’s host_name column changes.

In general, the same rules apply to host change triggers as to regular triggers, including applicable trigger patterns (such as trigger pattern 4, which indicates that the column returned by the query corresponding to idColumn should be type or subtype of the targetType element). The rules related to trigger lifecycle and regular trigger placements in files and plug-ins also apply to host name change triggers.

10.3.10 Understanding Activation Expressions

Note that this feature is available starting with the Enterprise Manager 12c PS1 platform release. As described previously, rules are normally owned by the plug-ins that require all target types needed by the rule query to be present by the time the rule is installed. However, on rare occasions you may encounter a case where two or more plug-ins needed by a rule query are independent and any one of the plug-ins may exist without the presence of the other. In other words, it may not be possible to

specify that one plug-in is a prerequisite of another for a given rule query that relies on configuration tables and data from target types of both plug-ins.

For such cases, you can specify an activation expression in a rule that will indicate when the rule should be active. Note that the rule is still owned by (at most) a single plug-in and the rule query can only be specified in one plug-in that will in the future be responsible for changing or removing the rule. However, the rule could be inactive for as long as not all needed target types are present on the system.

In terms of syntax, you specify activation expression using an attribute in the Rule element where the rule's query is specified:

```
<Rule name="..." activation_expr="..."> ...
```

Normally, when the `activation_expr` attribute is not present, it implies that the rule should always be active. If it is present, its value is a Boolean expression which must produce true if and only if the rule should be active. The expression can use (case insensitive) "AND", "OR", and parenthesis. Operands of the expression are target types. Each occurrence of a target type evaluates to "true" if and only if the target type is present in Enterprise Manager.

Note that a number of target types do not need to be listed in the expression because they are always going to be present whenever the rule is installed and present in Enterprise Manager installation. These include:

- target types installed with the plug-in where the rule resides (target types in the plug-in that owns the rule).
- target types in other plug-ins on which the plug-in owning the rule depends.
- target types always installed with Enterprise Manager (like host).

Therefore, in many cases when activation expression is used, a single target type as described in Example 1 below may suffice:

1. Example 1: oracle_ovm

This simple expression implies that the rule should be active only if `oracle_ovm` target type is installed in Enterprise Manager (in addition to any other target types that are already known to be present when this rule is installed).

This kind of activation expression could be expected in a rule with a query that relies on configuration tables of `oracle_ovm` and `oracle_xyz` (for example) target types. Assuming these target types belong to different and independent plug-ins, if the rule is placed in a plug-in owning `oracle_xyz` target type, its activation expression would be `oracle_ovm`.

2. Example 2: oracle_ovm and (oracle_oam_cluster or oracle_oim_cluster)

This expression implies that the rule should be active only if the `oracle_ovm` target type is present and either `oracle_oam_cluster` or `oracle_oim_cluster` is also present.

Please note that activation expressions should be used very carefully and rarely, since their usages are error prone due to lack of checks prior to rule activation. For example, any typo in a target type or any logical expression error may result in the rule never being activated or not being activated in correct cases. Enterprise Manager cannot check for validity of target types because it will assume unknown target types in the expression may get installed in the future.

The following describes how the activation expression feature interacts with other derived association features:

- During a new release of a rule XML, if rule query is unchanged but the activation expression is changed, the activation expression is updated and the rule is activated or deactivated if needed. If the rule is activated or deactivated, the rule's association instances are reevaluated or removed, respectively.
- Whenever Enterprise Manager adds or removes a target type (due to installation or deinstallation of a plugin for example), Oracle will reevaluate relevant activation expressions and activate or deactivate corresponding rules accordingly.

Note that target type addition is performed before any corresponding targets and their associations or data are added. Target type deletion is done after target instances and their associations are removed. Therefore, we do not reevaluate corresponding association instances for the affected rules. By the time the rule is activated due to the addition of a target type, no associations should exist for such a rule.

Similarly, when the rule is deactivated due to the removal of a target type, the associations are also removed because all targets of that target type are removed. This logic applies to all known cases, including when the target types in the expression are those of source, destination, or one of the derivation targets. Thus, there is no reevaluation of the rule upon target type addition or removal.

- Note that there is a difference between the quarantine feature and the activation expressions. The quarantine feature is controlled by the end-users or administrators to decide which rule evaluations to turn off. On the other hand, activation expressions are controlled by the rule authors and a given Enterprise Manager setup (for example, the presence or absence of involved target types).

Rules that were never activated cannot be quarantined. Otherwise, all other combinations are supported. For example, if a quarantined rule is deactivated and then activated again using an activation expression, it stays quarantined.

Similarly, if an active rule gets quarantined by an administrator and later becomes deactivated due to a target type removal, administrators can still unquarantine it so that if it ever gets activated, it will start computing associations.

- When a rule is being removed, rules activation expression and its activation status are ignored. In other words, a rule can be removed even if it is inactive.

Note that the following feature is available starting with the Enterprise Manager 12c PS2 platform release. This enhancement allows target type version specification as part of the target type expressions. In other words, not only can you list a target type, such as "oracle_ovm" as part of the activation expression, but also the version information can be included in square brackets after the target type specification. This is useful when, for example, a table is defined stating a given version (also known as metaver) of the target type metadata. In such a case, you would want to specify that only when the target type of a given version or above is installed should the rule relying on the table be active (see third bulled below).

In general, there are four kinds of version specifications that are supported:

- "target_type[version], e.g. "oracle_ovm[3.7]"
Indicates that this part of the expression is true only if the specified version of the target type is present. In above example, if oracle_ovm target type version 3.7 is not installed, the expression will evaluate to false even if other versions of oracle_ovm are present.
- "target_type[version1-version2], e.g. "oracle_ovm[3.7-5.2]"
Indicates that this part of the expression is true only if a version of the target type is present that is between the indicated versions, including the two specified

versions. Thus, in above example, the expression would be true if oracle_ovm target type of version 3.7, 3.8, 4.5, 5.0, or 5.2 is present.

If, on the other hand, none of the installed versions of oracle_ovm target type fall into the range between 3.7 and 5.2, "oracle_ovm[3.7-5.2]" would evaluate to false.

- "target_type[version-], e.g. "oracle_ovm[3.7-]"

Indicates that this part of the expression is true only if a version of the target type is present that is the same as the specified version or greater than the specified version. This is the most commonly used variation and the specified version would normally be the one where a table used by the rule query is introduced into a target type collection. In above example, the expression will evaluate to true if and only if a version of 3.7 or higher for target type "oracle_ovm" is installed at Enterprise Manager.

- "target_type[-version], e.g. "oracle_ovm[-5.2]"

Indicates that this part of the expression is true only if a version of the target type is present that is the same as the specified version or less than the specified version. Thus, in above example, the expression will evaluate to true if and only if a version of 5.2 or less for target type "oracle_ovm" is installed on Enterprise Manager.

Note: "There are no spaces between target type and the opening square bracket.

Version specification is optional. You can use just target type specification implying that any version of the target type would satisfy that part of the expression.

For example:

```
"oracle_ovm[3.7-] and (oracle_oam_cluster or oracle_oim_cluster[-2.5] or oracle_oim_cluster[2.8-])"
```

This expression implies that the rule should be active only if oracle_ovm target type of version 3.7 or higher is present, and either oracle_oam_cluster (of any version) or oracle_oim_cluster of versions 2.5 and below or 2.8 and above is also present.

10.3.11 About Debugging

You can begin debugging by initiating the configuration collections used to fire your triggers. These collections will occur when a target is used for the first time or whenever you make changes to the configuration data contained in the tables or views specified in your triggers. You can restart the agent to recollect the data. Make sure that the data in your configuration tables changes as expected before checking whether the triggers fired.

You should also see if your rule query produces all the desired associations across your development Enterprise Manager repository. You can manually run the GC\$ECM_CONFIG.RUN_ASSOC_DERIV_RULE PL/SQL procedure in this release to check whether the desired associations will be created if a rule fires during the configuration change.

If the associations you expect are not getting created, try the following:

- Check that the association instances are present in MGMT_ASSOC_INST_ORIGINS.

It should return a row for each association that should exist based on the specified rule and the collection for your target. If not, check that you did not enter an incorrect query or specify the wrong flag on the trigger.

You can try variants on the last line. For example:

```
WHERE source_me_guid = :y
WHERE dest_me_guid = :y
WHERE derivation_target_guid[N] = :y
```

Once you have logged the rule query, the log will reflect the rows in the MGMT_ASSOC_INST_ORIGINS table that need to be changed, followed by the actual association rows in MGMT_ASSOC_INSTANCES.

10.4 Ensuring Performance

The rules you provide may be fired frequently, depending on the triggers you define and the change frequency for the corresponding configuration tables. Poor performance of frequently triggered rules can adversely affect overall OMS operation.

A rule query is mapped to more specific queries. The query that gets executed depends on the column flag of the trigger (source, destination, or derivationTarget[N] column). As noted previously, the rule query <Q> is mapped to a query of this form:

```
SELECT a.*,
       FROM (<Q>) a
       WHERE <FC> = ?
```

where the parameter is the GUID of the target for which the Enterprise Configuration Management collection fired the trigger and <FC> is the source/destination/derivationTarget[N] column of your query specified by the trigger.

As you can see, the overall query (<Q>) will be filtered based on one of the target GUIDs that it returns. This means that query plans will generally start with data for that target and perform joins from there. Your query plans must push the "<FC> = ?" predicate all the way down and start evaluation with this predicate. Normally, they contain many nested loops that on the deepest level perform above binding and then get to the rest of the data starting from there via indexed lookups. Normally, there should be no full table scans of potentially large data tables (or hash joins).

Consider the example in [Section 10.3.7.3, "Listener and Database"](#). When the first trigger fires, the query will bind a database target and look for associated listener targets. The only way the Enterprise Manager repository can find the other end of the association (for example, the listener targets) is via "cm\$listener_ports oralsnr_ports" joins on the machine_name and listener_port columns. If the table under view cm\$listener_ports can be large, rule author should ensure an index exists starting with these two columns to quickly locate the listener targets instead of performing a full table scan on the table.

For each trigger, you must ensure that supporting indexes are present and that you test the performance of your queries after surrounding them in the query, as described earlier. However, if for example you have multiple triggers for the source column flag, you may have to test the performance only once as the generated query will be the same for both queries.

10.4.1 Using Custom Configuration Specifications for Data Collection

Skip this note if you are not familiar with Custom Configuration Specifications (CCS). CCS tables are generic so that they can accommodate a variety of data and so tend not to perform well for querying. However, properly modeled ECM tables are specific to the data being collected and can perform well for critical code paths, like derived associations computation code. Therefore, you should not use CCS collected data for derived association rules. Instead, collect the data required for derived associations separately using standard ECM collections.

10.5 Using Overlapping Associations

It is possible for more than one rule to derive the same association, although usually you should avoid creating such overlapping rules. This section describes what happens when an association is derived by multiple rules and includes suggestions on when to avoid this and how.

10.5.1 Overlap Between Associations Derived by Rules

When more than one rule derives the same association, that association continues to exist until each rule no longer derives it.

Sometimes, this is what you want. For instance, suppose each of two application target types has knowledge of both the Oracle WebLogic Server on which it runs and the database it accesses. Based on that knowledge, each has a way to derive an association between the Oracle WebLogic Server and the database. If either rule derives the association, that association is real and should exist. Only when both rules no longer derive the association can you be sure that the association no longer exists.

The "exists when any rule derives it" semantics may not be what you intend. Consider two rules that could be defined for the `installed_on` association between the database and Oracle home. Both access the same data, but one is triggered by a property change to the Oracle home and the other by a change to the database. As soon as either rule determines the relationship is gone, the association should be deleted. In such a case, you should use a single rule with two triggers.

Suppose you did not take care to write only one rule in such cases. You may think that this mistake is not serious as, after all, the association will soon be deleted. But this is not so, and the bogus association may exist indefinitely. If in the example described above the association was derived using two rules, then the database is upgraded and its OracleHome property gets changed. The association with the old home should be removed, but this will not happen until the other rule is fired. However, nothing about the Oracle Home target has changed, so its rule is not triggered and the association remains. Indeed, it is often the case that only one target is changed and the other remains unchanged for a long period of time.

As a general rule of thumb, associations based on data from a specific set of tables should be derived using a single rule with multiple triggers.

Unless there are different reasons for asserting an association exists, you should only use one rule. In such cases, the associations returned by derivation rules should be disjoint. Another way to state this is that for those associations, the set of all rows returned by all rule queries must specify no duplications. An association is identified by source, destination, and association type. So this means that the combination of these three values should be unique.

10.6 Creating Associations for Composite and System Target Types

There are several types of associations that must be considered when constructing either a composite target or a system target, and there are several ways in which these associations can be added to EM. The following describes each of these types of associations, how they are used by the EM framework and the typical approach to how they are created.

10.6.1 Composite Membership and the Containment Association

The first important association type is the "contains" association. This association type is typically added between a composite target and each of its members. The presence of this association is necessary in order to populate the target navigator (tree) for a composite target. The set of targets that are members of a composite (associated with it through a "contains" association) can be retrieved using the `getCompositeMembers()` method of the `Target` class.

These containment associations are most often created during discovery, using either a fully automated approach or through the guided discovery process. In either case, the discovery scripts provided with the plug-in are used to identify the set of containment associations between the composite and its members. Other non-containment associations may also be discovered; however, they will not affect the membership of the composite and will not affect the contents of the target navigator.

If the composite target is also to be treated as a system, it is strongly recommended that the system stencil include rule paths that represent the type of containment associations created in this way. This ensures that the target navigator and system membership display member targets consistently.

10.6.2 Other Non-Composite Associations (Composite Topology)

In addition to the discovery of containment associations, you may wish to represent other types of associations between the members of your composite topology. These associations may have meaning only to your target administrators and may be used by your plug-in code to perform other operations.

These associations are typically created using a discovery script, either as part of fully automated discovery or through the guided discovery process.

10.6.3 System Membership Associations

System membership is constructed by the Enterprise Manager framework based on the system stencil. This stencil defines the set of native associations that should be considered when identifying the system members. These native associations are typically either containment or other non-composite associations.

If these associations are found during the evaluation of the system stencil, the destination targets are added as system members.

10.6.4 Associations to External Targets

Up to this point all of the associations discussed have been between targets that are assumed to be part of the same plug-in, and therefore your plug-in code including discovery and UI has intimate direct knowledge of the configuration and topology of these targets.

However, in some cases your target configuration may include associations to other targets not included in your plug-in, such as an Oracle Database used as an

application or backing store for your targets. The configuration of your target knows something about the database that it uses, likely some connection related details such as host-port-sid or host-service.

You would like to represent this association between your target and the database in Enterprise Manager so that if Enterprise Manager is managing the database, the end-user can see this relationship and traverse it to obtain other information about that database and manage it (if appropriate and allowed).

Because you do not know if Enterprise Manager is managing the database and the identifying information you have is not the Enterprise Manager database target name, but instead the connect information, you can construct a derivation rule that maps the connection information in your target's configuration to that of a database in Enterprise Manager.

10.6.5 Regarding the Timing of Association Creation

Because the creation and deletion of targets and associations can be initiated from various sources (such as automated discovery, guided discovery, derived association rules, and system stencil rules), there are cases where the topology of a composite entity in Enterprise Manager may not appear in sync with the reality of that entity. As a plug-in developer, it is important that you are aware of this and account for it in information you present to end-users whenever possible.

One typical situation that may occur is that the discovery of targets occurs, configuration information is collected, and this is followed by the modification of associations as derivation rules are processed by the Enterprise Manager association framework. In this scenario, the user will first see the topology of the entity, including any association added during discovery. However, the additional associations created by derivation rules will not appear until sometime later.

10.7 Frequently Asked Questions

This section addresses three of the most frequently asked questions:

1. [Which tables can I reference in a rule query?](#)
2. [Are there guidelines for when to use target properties?](#)
3. [What is the relationship between discovered and derived associations?](#)

10.7.1 Which tables can I reference in a rule query?

In most cases, your query will just reference configuration (Enterprise Configuration Management) tables using the CM\$ views, and so will your triggers. For a more complete list of objects that can be referenced, see [Section 10.3.1, "Using Association Derivation Rules Syntax and Semantics"](#).

If you refer to other tables and if that data may change independently of Enterprise Configuration Management table changes, then the associations may not be updated when needed. If you have a use case in which a non Enterprise Configuration Management table is referenced where changes to that table must trigger rule evaluation, contact your Oracle representative.

Another consideration is the component in which the table is located. If the table your rule references is not part of the Enterprise Manager core EDK, your plug-in must account for the dependency on that table's plug-in. For example, you must ensure that any object that you reference already exists in the repository using a plug-in dependency mechanism.

10.7.2 Are there guidelines for when to use target properties?

Target properties are being treated as configuration data and there is an Enterprise Configuration Management snapshot table that is populated for each target type. Some care should be taken in using data from this table:

- Many target properties are set at discovery time and never modified.
- Querying name/value pair data can be awkward and take longer than queries on other tables where the data is more structured.
 - If the data is available from both the target properties table and an Enterprise Configuration Management snapshot table, you should use the latter.
 - If you need to add collection of configuration data, you should do so in an Enterprise Configuration Management table, not as a new row in the target properties table.

In general, the use of target properties should be avoided and data should be collected and modelled using standard ECM mechanisms.

However, a rule may need to refer to target properties if, for example, the target has no Enterprise Configuration Management collections that can be added. If an association to such a target is to be created, there must be some way to identify it (for example, the rule must refer to its target properties).

If you must use target properties, you should reference `MGMT_TARGET_PROPERTIES` in your rule query. You can also reference `MGMT$TARGET_PROPERTIES` in the rule query if the view already performs the join you need to do. However, in triggers you must use the `GC$TARGET_PROPERTIES` view for the `orcl_tp_config` snapshot type.

`MGMT_TARGET_PROPERTIES` should be used in queries because it is more efficient, but may include some properties not available in the `GC$` view. Triggering is only available based on property changes in the `GC$` view. For example, the `GC$` view only includes those properties that are properly registered with Enterprise Manager.

10.7.3 What is the relationship between discovered and derived associations?

This is another example of overlapping associations (for more information, see [Section 10.5, "Using Overlapping Associations"](#)). For instance, you may have discovery logic that discovers an association between targets T1 and T2, plus a rule that derives the same association. Oracle recommends that you do not write two sets of logic to create the same association. In this case it is suggested that:

- If a derivation rule is needed because the association may change, you should just write the derivation rule.
- If the association that is discovered will not change until the source or destination is removed, then discovering the association is fine and may be simpler or more efficient.

If you do write two sets of logic to create the same association (discovery logic and derivation rule), then the discovered association will remain and the derivation logic will also assert the existence of that association. If the rule evaluation later determines that the association should no longer exist, the rule's assertion will be removed, but the association will continue to exist unless you manually delete the discovered association.

Defining Target Discovery

The discovery of targets in Enterprise Manager can be accomplished in several different ways including automated discovery scripts, manual addition of targets by specifying target properties, and manual addition of targets using guided discovery.

Automatic target discovery is the process by which targets are located and added to Enterprise Manager. Automatic discovery begins when the Oracle Management Agent starts up after installation. Targets located on the server where the Management Agent is running are discovered and sent to the Management Repository as targets that are not yet managed. The end user can choose which targets to monitor by promoting these targets as targets managed by Enterprise Manager.

This chapter contains the following sections:

- [Introduction to Defining Target Discovery](#)
- [Creating Discovery XML](#)
- [Creating the Discovery Script](#)
- [Packaging Discovery XML and Discovery Content](#)
- [Setting Up and Testing Discovery](#)
- [Manually Adding Targets](#)
- [Configuring and Promoting Targets for Monitoring by Enterprise Manager](#)
- [Examples for Using Generic Discovery Framework](#)
- [Configuring Automatic Discovery For Plug-ins](#)

11.1 Introduction to Defining Target Discovery

As a plug-in developer, you are responsible for the following steps within the discovery process:

1. Create discovery metadata.

Use the Discovery XML Schema Definition (XSD) for guidelines about creating a discovery metadata XML file.

Within the metadata:

- Define the discovery modules using the `DiscoveryModule` element.
- Define discovery parameters (if required) using the `DiscoveryInput` element.

For information about creating discovery metadata, see [Section 11.2, "Creating Discovery XML"](#).

2. Create the discovery script using Perl.

The discovery script enables the Management Agent to automatically discover all the target types belonging to a plug-in.

For information about creating the discovery script, see [Section 11.3, "Creating the Discovery Script"](#).

3. Identify additional Perl modules or JAR files that are required for discovery.
4. Bundle the discovery metadata and contents into the plug-in staging directory (*plugin_stage*).

- a. Save the discovery XML in the *plugin_stage/oms/metadata/discovery* directory.

- b. Save the discovery content in the *plugin_stage/discovery* directory

For information about packaging, see [Section 11.4, "Packaging Discovery XML and Discovery Content"](#).

5. Repackage (if necessary) and deploy the plug-in.

After the plug-in archive is created and the plug-in is deployed to the Management Server, the end user can initiate discovery using the discovery configuration UI for the discovery modules that are registered.

For information about deploying the plug-in, see [Chapter 13, "Validating, Packaging, and Deploying the Plug-in"](#).

6. Configure automatic discovery:

- a. Log in to Enterprise Manager.

- b. Select **Setup**, then select **Add Target**, and then select **Configure Auto Discovery**.

The Configure Auto Discovery page appears. For information about the options on this page, see the Enterprise Manager online help.

For more information about configuring automatic discovery, see [Section 11.9, "Configuring Automatic Discovery For Plug-ins"](#).

7. Test discovery results.

For information about testing discovery, see [Section 11.5, "Setting Up and Testing Discovery"](#).

11.2 Creating Discovery XML

Oracle provides a Discovery XSD so you can write discovery metadata XML to register with the discovery framework. Registering with discovery framework enables the discovery pages to launch discovery of the target types belonging to a plug-in.

For more information about the Discovery XSD, see the Extensibility Development Kit (EDK) specifications.

The following section and [Example 11-4](#) provide examples of discovery XML.

11.2.1 Generic Discovery Integration Example

In this example, discovery requires no information entered and promotion of the target does not require any special logic.

There are no special requirements for configuring the target except that you must have access to the UI.

[Example 11-1](#) provides the discovery integration XML for this discovery.

Note: You are not restricted on the naming of the discovery XML file. However, the standard convention is `plugin_discovery.xml`.

Example 11-1 Discovery Integration XML

```
<?xml version="1.0" encoding="UTF-8"?>
  <EmTargetDiscovery
    xmlns="http://www.oracle.com/EnterpriseGridControl/disc_metadata"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"

xsi:schemaLocation="http://www.oracle.com/EnterpriseGridControl/disc_metadata
discovery.xsd ">
  <DiscoveryInfo>
    <AutomaticDiscovery>
      <DiscoveryModule name="simple_disc_plugin"
resourceBundlePkg="oracle.sysman.simpleplugin.rsc.simplePluinMsg">
        <Display NLSID="SIMPLE_DISC_MODULE">
          <NlsValue>simple_disc_plugin</NlsValue>
        </Display>

        <SupportedAgentOsList>
          <SupportedAgentOs>2000</SupportedAgentOs>
        </SupportedAgentOsList>
      <BasicDiscoveryInfo>
        <DiscoveryScript>SimplePluginDisc.pl</DiscoveryScript>
        <DiscoveryCategory>SIMPLE_PLUGIN_DISC</DiscoveryCategory>
      </BasicDiscoveryInfo>
      <TypesDiscovered>
        <TargetType>simple_plugin_target_type1</TargetType>
        <TargetType>simple_plugin_target_type2</TargetType>
      </TypesDiscovered>
    </DiscoveryModule>
  </AutomaticDiscovery>
</DiscoveryInfo>
</EmTargetDiscovery>
```

After the previous XML is registered, the discovery framework can launch discovery of `simplePluginType` using the `SimplePluginDisc.pl` discovery script.

11.2.2 Discovery Script Example

For example, you can have the following content in the `SimplePluginDisc.pl` script to discover two target instances:

```
#all the discovery scripts get emdRoot and hostname of the agent as arguments to
the script.
my ($emdRoot, $hostName) = @ARGV;
#Discovery root is sent to the script as env variable.
my $discovery_root = $ENV{DISC_ROOT};
print "<Targets>\n";
print "<Target TYPE=\"simple_plugin_target_type1\" NAME=\"smpl_tgt1\">\n";
print "<Property NAME=\"Prop1\" VALUE=\"prop1_foo\"/>\n";
print "    <Property NAME=\"Prop2\" VALUE=\"prop2_value_bar\"/>\n";
print "  </Target>\n";
```

```

print "  \<Target TYPE=\"simple_plugin_target_type2\" NAME=\"smp1_tgt2\" \>\n";
print "    \<Property NAME=\"Prop1\" VALUE=\"value_foo\" /\>\n";
print "      \<Property NAME=\"Prop2\" VALUE=\"value_bar\" /\>\n";
print "    \</Target\>\n";
print "  \</Targets\>\n";

```

This script produces the following output:

```

<Targets>
  <Target type="simple_plugin_target_type1" NAME="smp1_tgt1" >
    <Property NAME="Prop1" VALUE="prop1_foo" />
    <Property NAME="Prop2" VALUE="prop2_value_bar" />
  </Target>
  <Target type="simple_plugin_target_type2" NAME="smp1_tgt2" >
    <Property NAME="Prop1" VALUE="value_foo" />
    <Property NAME="Prop2" VALUE="value_bar" />
  </Target>
</Targets>

```

After your plug-in is deployed on the Management Server, the discovery module is listed in the Discovery UI. Users can configure discovery of this module on one or more Management Agents. This causes the discovery content of the plug-in to be deployed on the Management Agent and subsequently causing the discovery script to be run at the Management Agent.

For example, when you run this discovery script through autodiscovery, the targets will be generated and sent to the Management Server as Not Yet Managed targets. Using the Discovery Results UI, you can promote these two targets as managed targets by Enterprise Manager.

Note: There are methods for writing debug information such as EMD_PERL_INFO, EMD_PERL_DEBUG, and EMD_PERL_ERROR, which can be accessed through the Perl package emdcommon.pm. You can then find the trace information (written through Perl methods) in the Oracle Management Agent trace file (emagent_perl.trc) in the Management Agent log directory.

11.2.3 Overview of the Discovery Metadata Elements

Table 11–1 describes the key elements that define the discovery metadata:

Table 11–1 Key Elements in a plugin_discovery.xml File

Element	Description
EmTargetDiscovery	The root element for the file
DiscoveryInfo	Specifies one or more autodiscovery modules for this plug-in. Each auto discovery module is associated with a discovery script that is run on the Management Agent
AutomaticDiscovery	Specifies the autodiscovery module
DiscoveryModule	This is an element within the autodiscovery module. It includes the name attribute, which defines the name of the discovery module.
SupportedAgentOSList	Specifies the list of Management Agent platforms on which this discovery is supported. It includes the SupportedAgentOS attribute, which defines the platforms. The valid value is 2000 (All platforms).

Table 11–1 (Cont.) Key Elements in a *plugin_discovery.xml* File

Element	Description
BasicDiscoveryInfo	Specifies the discovery script to be run and an optional category name. It includes the following attributes: <ul style="list-style-type: none"> DiscoveryScript: Defines the name of the discovery script DiscoveryCategory: Defines the category name (optional)
TypesDiscovered	Specifies the list of target types that can be discovered using this discovery module. It includes the TargetType attribute, which defines the target type name.
DiscoveryInput	Specifies the information to be entered by the user during discovery. The information entered by the user is available as environment variables in the discovery script running at the Management Agent

11.3 Creating the Discovery Script

After the discovery XML is registered, the discovery framework uses a discovery script to launch discovery.

To create a discovery script:

1. Use Perl to write the top-level discovery script. This script can call Java, Shell, and so on.
2. By default, three variables are provided to the discovery Perl script:
 - a. The discovery framework provides *emdRoot* and the host name of the Management Agent as arguments to the script:

For example:

```
my ($emdRoot, $hostName) = @ARGV;
```

- b. The discovery root directory is sent to the script as an environment variable.

For example:

```
my $discovery_root = $ENV{DISC_ROOT};
```

- c. If there are discovery inputs, then they are made available as environment variables.

For example:

```
my $crs_home = $ENV{'CRS_HOME'};
```

Ideally, the output of the discovery script returns all the name-value pairs that are required to add a target of a target type. If not, the user must provide any remaining values at target promotion time.

[Example 11–5](#) provides an example of a discovery Perl script.

11.3.1 Discovered Targets DTD

The discovery Perl script must produce output which conforms to the following Document Type Definition (DTD):

Example 11–2 Discovered Targets DTD

```
<!ELEMENT Targets (Target*) >
```


Discovery and monitoring scripts are separate. Both are parts of your plug-in. You can run discovery without monitoring content. The lifecycle of both are managed by the discovery or plug-in lifecycle frameworks and is transparent to plug-in developers.

You must package the discovery content required for discovering a particular target type. For example, for an existing database discovery, the discovery content is `oracledb.pl` along with any required utilities for running database discovery.

Note: You must package the discovery metadata file with the Oracle Management Server archive and *not* the Management Agent archive.

Create a discovery directory under *plugin_stage* for installing content for each discovery plug-in. For more information about the directories under *plugin_stage*, see [Section 13, "Validating, Packaging, and Deploying the Plug-in"](#).

Example 11–3 Directory Structure for Installing Discovery Content

```
plugin_stage/discovery/
|
|__other subdirectories created as you specified
```

For Oracle Database discovery, the discovery content might look similar to the following:

```
plugin_stage/discovery/
|
|__oracledb.pl
|
|__utl
|
|__oracledbUtl.pl
|
|__initParamFileUtl.pl
|
|__winRegistry.pl
```

If any custom Perl modules are required for the discovery process, then place the modules under a similar directory structure as shown previously. The discovery root variable provided to the discovery script can be used to load this Perl module. For example, if your perl modules are placed under the *plugin_stage/discovery/utl/pm* directory, then you can load them from the discovery script as follows:

```
my ($emdRoot,$hostName,$crsHome) = @ARGV;
my $discovery_root = $ENV{DISC_ROOT};
require "$discovery_root/utl/pm/propertiesFileParser.pm";
require "$discovery_root/utl/pm/Targets.pm";
```

Note: Perl content that will be used by the discovery module and for other purposes, such as administration of the targets, should be packaged with the discovery bundle as well as the plug-in bundle.

11.4.2.1 Java Content Required by Discovery Scripts

Some discovery scripts (such as for Oracle Fusion Middleware) use JAR files in the process of discovery. If the JAR files are discovery-specific only, then they should be in discovery area.

If there are Java methods written to perform discovery, then they should be moved to a separate class file where possible, to avoid shipping content not required by discovery.

Oracle recommends separating discovery content and management content so that discovery can be performed independent of the management content present. This helps you to decide whether you want to install the plug-in to manage the discovered targets, if any. The discovery content should be lightweight and include the files necessary for discovery only.

For example, you can place the JAR containing discovery-specific code of the particular discovery module under the following directory:

```
plugin_stage/discovery/lib
```

Note: The individual discovery script is responsible for constructing the class path.

You can create your own directory structure in the discovery content area for a particular discovery module. The top level Perl scripts responsible for each discovery module, such as Fusion Middleware, construct the class path before running the Java utilities that they use. Because the JAR files specific to a discovery module will be in their own discovery content area, the discovery Perl script can construct the required class path easily. Again, the code responsible for performing discovery only should be separated out and installed in the discovery content area specific to the particular discovery module.

11.5 Setting Up and Testing Discovery

For testing purposes, you can run discovery:

1. Log in to Enterprise Manager.

```
https://em_host:em_port/em/
```

2. From the console, select **Setup**, then **Add Target**, and then **Configure Auto Discovery**.

The Configure Auto Discovery page appears.

3. In the Configure Auto Discovery tables, under Discovery Module, select **Multiple Target-Type Discovery on Single Host**.

The Target Discovery (Agent Based) page appears.

4. Select the required agent host name and click **Run Discovery Now**.

When target discovery is complete, a Completed Successfully window appears.

11.6 Manually Adding Targets

In addition to automatic discovery, Cloud Control allows you to manually add hosts as well as a wide variety of Oracle software and components as managed targets.

When you add a target manually, you do not need to go through the process of

discovery by adding the target directly. Discovering targets in this way eliminates the need to consume resources on the agent to perform discovery when it is not needed.

You must be able to specify the properties of a target to be managed and create an Enterprise Manager managed target.

Not all target types can be manually added. During registration with the discovery framework, the target type owner indicates whether a target type can be manually added or not.

See the following sections for instructions:

- [Manually Adding Host Targets](#)
- [Manually Adding Non-Host Targets](#)

11.6.1 Manually Adding Host Targets

A wizard guides you through the process of manually deploying a Management Agent to a new host target.

For instructions on installing a Management Agent, see "Installing Oracle Management Agent" in the *Enterprise Manager Cloud Control Basic Installation Guide*.

11.6.2 Manually Adding Non-Host Targets

A configuration page or wizard based on target type metadata listing all the instance properties required to manage target is displayed.

You can specify a name for the target and provide the required configuration information.

To add targets manually to Enterprise Manager:

1. Log in into Enterprise Manager.
2. From the **Setup** menu, select **Add Target**, then select **Add Targets Manually**.

Enterprise Manager displays the Add Targets Manually page.

3. Under the Add Targets Manually page, go to the Add Targets Manually sub-section and select one of the following options:

- Add Non-Host Targets Using Guided Process

From the Target Types list, select one of the target types to add, such as **Oracle Cluster and High Availability Service**, **Oracle Database Machine**, or **WebLogic Domain Discovery**, then click **Add Using Guided Discovery**. This process will also add related targets.

- Add Non-Host Targets by Specifying Target Monitoring Properties

From the Target Type list, select one of the target types to add, such as Fusion J2EE Application, Applications Utilities, or Supplier Portal, and from the Monitoring Agent list, select the required Monitoring Agent, then click **Add Manually**.

4. After you select the target type, you will follow a wizard specific to the target type to add the target.

Upon confirmation, the target becomes a managed target in Enterprise Manager. Enterprise Manager accepts the information, performs validation of the supplied data where possible and starts monitoring the target.

See Also: For more information about adding targets manually to Enterprise Manager, see Oracle Enterprise Manager Online Help.

11.7 Configuring and Promoting Targets for Monitoring by Enterprise Manager

Discovery of targets on Enterprise Manager managed hosts provides you with a list of targets such as new databases, and SQL servers which are not yet managed by Enterprise Manager. This helps you to determine if any of the new targets found are candidates for monitoring and managing by Enterprise Manager.

The Enterprise Manager UI allows you to review discovered unmanaged targets and promote targets to be managed by Enterprise Manager for monitoring.

See Also: For more information about configuring and promoting targets, see Oracle Enterprise Manager Online Help.

11.8 Examples for Using Generic Discovery Framework

You can use the generic discovery UI to launch and schedule discovery to run periodically at the Management Agent. For the generic discovery UI to launch discovery, you must specify the inputs that are required through the registration XML. Enterprise Manager uses this information to construct the generic discovery UI.

11.8.1 Discovery Integration Example Requiring User Input

In this example, discovery requires the user to enter information. The user is prompted to enter values for CRS_HOME and CRS_HOME1. The user does not have to enter a value for CRS_HOME but a value for CRS_HOME1 is mandatory

There are no special requirements for configuring the target except that you must have access to the UI.

Note: While supported, Oracle does not recommend the use of required discovery inputs because it eliminates the benefits of automatic discovery. You can use optional discovery inputs as hints to optimize the discovery process.

Example 11-4 Discovery Integration XML With Discovery Parameters

```
<?xml version="1.0" encoding="UTF-8"?>
  <EmTargetDiscovery
    xmlns="http://www.oracle.com/EnterpriseGridControl/disc_metadata"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"

    xsi:schemaLocation="http://www.oracle.com/EnterpriseGridControl/disc_
metadata discovery.xsd ">
    <DiscoveryInfo>
      <AutomaticDiscovery>
        <DiscoveryModule name="simple_disc_plugin"
resourceBundlePkg="oracle.sysman.simpleplugin.rsc.simplePluinMsg">

          <Display NLSID="SIMPLE_DISC_MODULE_MSG_ID">
            <NlsValue>simple_disc_plugin</NlsValue>
```



```

        </Display>
        <SupportedAgentOsList>
            <SupportedAgentOs>2000</SupportedAgentOs>
        </SupportedAgentOsList>

        <BasicDiscoveryInfo>
        <DiscoveryScript>simple_plugin_disc.pl</DiscoveryScript>
        <DiscoveryCategory>SIMPLE_PLUGIN_DISC</DiscoveryCategory>
        </BasicDiscoveryInfo>
        <TypesDiscovered>
            <TargetType>simple_plugin_target_type1</TargetType>
            <TargetType>simple_plugin_target_type2</TargetType>
        </TypesDiscovered>

        <!-- optional discovery hint -->
        <DiscoveryInput name="CRS_HOME" isRequired="false">
        </DiscoveryInput>

        <!-- a required discovery input -->
        <DiscoveryInput name="CRS_HOME1" isRequired="true">
        </DiscoveryInput>

    </DiscoveryModule>
</AutomaticDiscovery>
</DiscoveryInfo>
</EmTargetDiscovery>

```

After the previous XML is registered, the discovery framework can launch discovery using the `simple_plugin_disc.pl` discovery script

The `simple_plugin_disc.pl` script can access the discovery parameters as illustrated in the sample `simple_plugin_disc.pl` script.

Example 11–5 Sample Discovery Script

```

my $discovery_root = $ENV{DISC_ROOT};

#Inputs passed to the script.
my $crs_home = $ENV{'CRS_HOME'};
my $crs_home1 = $ENV{'CRS_HOME1'};

#add the logic here to find the targets.

print "\<Targets\>\n";
print " \<Target TYPE=\"simple_plugin_target_type1\" NAME=\"smpl_tgt1\">\n";
print "   \<Property NAME=\"Prop1\" VALUE=\"prop1_foo\"/>\n";
print "   \<Property NAME=\"Prop2\" VALUE=\"prop2_value_bar\"/>\n";
print " \</Target\>\n";
print " \<Target TYPE=\"simple_plugin_target_type2\" NAME=\"smpl_tgt2\">\n";
print "   \<Property NAME=\"Prop1\" VALUE=\"value_foo\"/>\n";
print "   \<Property NAME=\"Prop2\" VALUE=\"value_bar\"/>\n";
print " \</Target\>\n";
print " \</Targets\>\n";

```

Note: For guidelines about the output of the discovery Perl script, see [Section 11.3.1, "Discovered Targets DTD"](#).

11.9 Configuring Automatic Discovery For Plug-ins

You can configure automatic discovery for targets in a plug-in to run at the Management Agent-side. Currently, automatic discovery is scheduled to run every day.

Configuration is done from the Oracle Management Server where your metadata resides.

Note: For plug-ins deployed before the Management Agents are installed or plug-ins deployed to new hosts, discovery runs every 24 hours automatically (only if discovery does not require any user inputs).

You can configure parameters from the Enterprise Manager UI after the Management Agents are installed or deployed.

After discovery, the targets are sent to Enterprise Manager. The end user can then review the targets and choose which targets to monitor by promoting the targets as targets managed by Enterprise Manager.

See Also: For more information about configuring automatic discovery, see Oracle Enterprise Manager Online Help.

Adding Compliance Standards

The Oracle Enterprise Manager Compliance Management solution provides the capability to define, customize, and manage compliance frameworks and compliance standards.

This chapter contains the following sections:

- [Introduction to Adding Compliance Standards](#)
- [About the Compliance Standard Rules](#)
- [Defining Compliance Standards](#)
- [Defining a Compliance Framework](#)
- [Defining Compliance Content](#)
- [Removing Compliance Content](#)
- [Supporting Translation](#)
- [Packaging Compliance XML](#)
- [Setting Up and Testing Compliance Standards and Rules](#)
- [More Compliance Examples](#)
- [Publishing Compliance Content Using Self Update](#)

12.1 Introduction to Adding Compliance Standards

As a plug-in developer, you are responsible for the following steps when adding compliance standards:

1. Define compliance standard rules.

Compliance standard rules can be either of the following:

- Repository check-based rules
- Real-time monitoring rules.

For information about defining compliance standard rules, see [Section 12.2, "About the Compliance Standard Rules"](#).

2. Define a compliance standard.

For more information, see [Section 12.3, "Defining Compliance Standards"](#).

3. Define a compliance framework.

For more information, see [Section 12.4, "Defining a Compliance Framework"](#).

4. Package the compliance standard rules, standards, and framework as metadata XML.

For more information, see [Section 12.8, "Packaging Compliance XML"](#).

5. Set up and test the compliance content.

For more information, see [Section 12.9, "Setting Up and Testing Compliance Standards and Rules"](#).

6. Deploy the plug-in.

For information about deploying plug-ins, see [Chapter 13, "Validating, Packaging, and Deploying the Plug-in"](#).

12.1.1 Assumptions and Prerequisites

This chapter assumes you are familiar with the following:

- Plug-in development overview, including how to package a plug-in and its XML files.
- If you are installing compliance data with the `emctl register oms metadata -service gccompliance` command, it will require an `EM_COMPLIANCE_UTIL.trigger_rule_dependency_job` callback. You must enter the following commands through SQL*Plus as the SYSMAN user:

```
begin EM_COMPLIANCE_UTIL.trigger_rule_dependency_job;  
end;  
/
```

Note: This is necessary *only* if you are using the `emctl register oms metadata -service gccompliance` command to install compliance content.

If you are installing the plug-in, then you do not have to enter the previous SQL.

For more information about the `emctl register oms metadata` command, see [Section 13.7, "Updating Deployed Metadata Files Using the Metadata Registration Service \(MRS\)"](#).

12.2 About the Compliance Standard Rules

This section provides a description of the following:

- [Defining Repository Check-based Rules](#)
- [Defining Real-time Monitoring Rules](#)

12.2.1 Defining Repository Check-based Rules

A repository check-based rule checks the configuration state of one or more targets. A rule is compliant if the test fails to identify a violation. In other words, the test determines that the configuration item is in the required state or has the prescribed value. Any rule that uncovers a violation is noncompliant.

The data source that is evaluated by a rules test condition can be based on a repository query. A rules test condition can be implemented using a threshold condition based on

the underlying metrics or queries column value, or SQL expression, or a PL/SQL function. (The policies are similar to Oracle Enterprise Manager 10g Release 5).

The key points in this area include:

- Defining Compliance Standard Rules, Compliance Standards, and Compliance Frameworks
- Replacing out-of-box policy groups (10.2.x/11.10) with Compliance Standards that you create referring to Compliance Standard Rules
- Mapping your compliance standards to the appropriate Compliance Frameworks
- Defining Oracle Business Intelligence Publisher (BI Publisher) reports for compliance

[Example 12-1](#) provides the syntax for defining repository rules and [Example 12-2](#) on page 12-4 provides an example of a repository rule definition.

Note: For the complete compliance XML Schema Definitions (XSDs), see the following JAR file:

\$ORACLE_HOME/sysman/jlib/gccomplianceCommon.jar

See Also: For additional examples, see [Section 12.10, "More Compliance Examples"](#).

Example 12-1 Repository Rule Definition Syntax

```
<xsd:complexType name="RuleT">
  <xsd:sequence>
    <xsd:element name="DisplayName" type="std:DisplayString256Def"
minOccurs="0"/>
    <xsd:element name="TargetType" type="std:Name256Def"/>
    <xsd:element name="IsSystem" type="std:BooleanDef" minOccurs="0"
default="false"/>
    <xsd:element name="IsHidden" type="std:BooleanDef" minOccurs="0"
default="false"/>
    <xsd:element name="evaluateAlways" type="std:BooleanDef" default="false"
minOccurs="0"/>
    <!-- E.g. target version, platform based filter -->
    <xsd:element ref="std:TargetPropertyFilter" minOccurs="0"/>
    <xsd:element name="Description" type="std:DisplayString800Def"
minOccurs="0"/>
    <xsd:element name="Impact" type="std:DisplayString800Def" minOccurs="0"/>
    <xsd:element name="Recommendation" type="std:DisplayString800Def"
minOccurs="0"/>
    <xsd:element name="FixLinkList" type="std:FixLinkListT" minOccurs="0"/>
    <xsd:element name="ViolationContextList" type="std:ViolationContextListT"/>
    <xsd:element name="CheckSource" type="std:CheckSourceT" minOccurs="1"
maxOccurs="1"/>

    <xsd:element name="Severity" default="MinorWarning" minOccurs="0">
      <xsd:simpleType>
        <xsd:restriction base="xsd:string">
          <xsd:enumeration value="MinorWarning"/>
          <xsd:enumeration value="Warning"/>
          <xsd:enumeration value="Critical"/>
        </xsd:restriction>
      </xsd:element>
    </xsd:sequence>
  </xsd:complexType>
```

```

</xsd:simpleType>
</xsd:element>
<xsd:element name="LifeCycleStatus" default="Development" minOccurs="0">
  <xsd:simpleType>
    <xsd:restriction base="xsd:string">
      <xsd:enumeration value="Development"/>
      <xsd:enumeration value="Production"/>
    </xsd:restriction>
  </xsd:simpleType>
</xsd:element>
<xsd:element name="KeywordList" type="std:KeywordListT" minOccurs="0"/>
<xsd:element name="UrlLink" type="std:String4000Def" minOccurs="0"/>
<xsd:element name="ViolationMessage" type="std:DisplayString800Def"
minOccurs="0"/>
  <xsd:element name="ClearViolationMessage" type="std:DisplayString800Def"
minOccurs="0"/>
    <xsd:element name="Author" type="std:Name256Def" minOccurs="0"/>
    <xsd:element name="LastUpdatedBy" type="std:Name256Def" minOccurs="0"/>
  </xsd:sequence>
  <xsd:attribute name="name" type="std:Name256Def" use="required"/>
  <xsd:attribute name="oms_version" type="std:Name32Def" use="required"/>
</xsd:complexType>

```

[Example 12-2](#) is defined for `oracle_database` target_type, which is part of the database plug-in.

You can define a rule for any target type registered with Enterprise Manager.

Example 12-2 Sample Rule

```

<Rule xmlns="http://www.oracle.com/DataCenter/ConfigStd" oms_version="12.1.0.1.0"
name="sample_rule1">
  <DisplayName nlsid="SAMPLE_RULE_1_NAME">Sample Rule 1</DisplayName>
  <TargetType>oracle_database</TargetType>
  <IsSystem>true</IsSystem>
  <TargetPropertyFilter>
    <PropertyItem>
      <PropertyName>orcl_gtp_operating_system</PropertyName>
      <ValueList>
        <Value>Windows</Value>
      </ValueList>
    </PropertyItem>
    <PropertyItem>
      <PropertyName>orcl_gtp_target_version</PropertyName>
      <ValueList>
        <Value>8.1.6+</Value>
      </ValueList>
    </PropertyItem>
  </TargetPropertyFilter>
  <Description nlsid="SAMPLE_RULE_1_DESC">Checks for use of a single control
file</Description>
  <Impact nlsid="SAMPLE_RULE_1_IMPACT">The control file is one of the most
important files in an Oracle database. It maintains many physical characteristics
and important recovery information about the database. If you lose the only copy
of the control file due to a media error, there will be unnecessary down time and
other risks.</Impact>
  <Recommendation nlsid="SAMPLE_RULE_1_RECO">Use at least two control files that
are multiplexed on different disks.</Recommendation>
  <ViolationContextList>
    <Column type="String" name="FILE_LIST">
      <DisplayLabel nlsid="SAMPLE_RULE_1_COL_1">FILE_LIST</DisplayLabel>
    </Column>
  </ViolationContextList>

```

```

        <IsHidden>false</IsHidden>
        <IsKey>false</IsKey>
    </Column>
    <Column type="Number" name="CONTROL_FILE_COUNT">
        <DisplayLabel nlsid="SAMPLE_RULE_1_COL_2">CONTROL_FILE_
COUNT</DisplayLabel>
        <IsHidden>false</IsHidden>
        <IsKey>false</IsKey>
    </Column>
</ViolationContextList>
<CheckSource>
    <RepositoryCheckDefinition>
        <Metric>
            <TargetType>oracle_database</TargetType>
            <MetricName>sample_rule1</MetricName>
            <SourceType>SQL</SourceType>
            <Source>select CONTROL_FILE_COUNT, FILE_LIST, TARGET_GUID from MGMT$CS_DB_
CONTROL_FILE_COUNT</Source>
            <MetricColumnList>
                <MetricColumnInfo>
                    <ColumnName>FILE_LIST</ColumnName>
                    <ColumnType>String</ColumnType>
                    <isKey>false</isKey>
                    <ColumnLabel nlsid="SAMPLE_RULE_1_COL_1">FILE_LIST</ColumnLabel>
                </MetricColumnInfo>
                <MetricColumnInfo>
                    <ColumnName>CONTROL_FILE_COUNT</ColumnName>
                    <ColumnType>Number</ColumnType>
                    <isKey>false</isKey>
                    <ColumnLabel nlsid="SAMPLE_RULE_1_COL_2">CONTROL_FILE_
COUNT</ColumnLabel>
                </MetricColumnInfo>
            </MetricColumnList>
        </Metric>
        <ParameterList>
            <RuleParameter>
                <ParamName>CONTROL_FILE_COUNT</ParamName>
                <ParamType>Number</ParamType>
            </RuleParameter>
        </ParameterList>
        <ParameterDefaultSettings>
            <ParamValue>
                <ParamName>CONTROL_FILE_COUNT</ParamName>
                <MinorWarnThreshold>1</MinorWarnThreshold>
            </ParamValue>
        </ParameterDefaultSettings>
        <TestCondition>
            <ThresholdCriteria>
                <ColumnName>CONTROL_FILE_COUNT</ColumnName>
                <TestOperator>EQ</TestOperator>
                <ThresholdValue>1</ThresholdValue>
                <ThresholdType>Number</ThresholdType>
            </ThresholdCriteria>
        </TestCondition>
    </RepositoryCheckDefinition>
</CheckSource>
<Severity>MinorWarning</Severity>
<LifeCycleStatus>Production</LifeCycleStatus>
<KeywordList>
    <Keyword nlsid="CONFIGURATION">Configuration</Keyword>

```

```

    </KeywordList>
    <ViolationMessage nlsid="SAMPLE_RULE_1_VIOL_MSG">The database has an
insufficient number of control files.</ViolationMessage>
    <ClearViolationMessage nlsid="SAMPLE_RULE_1_VIOL_CLEAR_MSG">The database has
sufficient number of control files.</ClearViolationMessage>
    <Author>SYSMAN</Author>
</Rule>

```

Table 12–1 provides a description of the tags used to define a repository rule:

Table 12–1 Key Tags for Defining Repository Rules

Tag	Subtag	Description
DisplayName		The display name of the rule. It provides the nlsid attribute to support the translation of messages.
TargetType		The type of target to which this rule is can be associated
IsSystem		True for out-of-the-box rules. Otherwise, False
IsHidden		False by default. When set to True, the IsHidden rules are not visible in the UI and no events are generated. This element should be set to true for out-of-the-box rules
Description		The description of the rule. It provides the nlsid attribute to support the translation of messages.
Impact		Impact if the rule violates (when rule is noncompliant). It provides the nlsid attribute to support the translation of messages
ViolationContextList		Violation context defines a violation to a rule uniquely. Violation context lists columns from <Source> Query, which will be visible as a part of the violation. Each column must mark as key or non-key. The mandatory target_guid column from <Source> query is implicitly added to the violation context and should not be included in the violation context explicitly.
	Column	Metric Column name. Uses Attributes name and type <ul style="list-style-type: none"> ■ DisplayLabel: Display name for column. It provides the nlsid attribute to support the translation of messages ■ IsHidden: True, if this column should not be displayed as a part of a violation context. Otherwise, False. ■ IsKey: True, if this column is a key
CheckSource		Defines the data source for Rule evaluation.
..RepositoryCheckDefinition		Defines data source for Repository Rule.

Table 12–1 (Cont.) Key Tags for Defining Repository Rules

Tag	Subtag	Description
	Metric	<p>Defines data source query.</p> <ul style="list-style-type: none"> ■ MetricName: Name of metric ■ SourceType: SQL. The only supported SourceType ■ Source: This is a SQL query written over MGMT\$_% views. (MGMT\$_% views are provided with Enterprise Manager). If required, this SQL query can be written over other provided views that they have a direct SELECT privilege to the MGMT_VIEW user in Enterprise Manager. For information about Enterprise Configuration Management views, see Chapter 6, "Collecting Target Configuration Data". ■ MetricColumnList: List of columns in Source query ■ MetricColumnInfo: Metric column ColumnName: Metric Column name ColumnType: Metric Column Type isKey: True, if column is a key column. Otherwise, False. ColumnLabel: Column display name
ParameterList		<p>List of parameters</p> <p>Note: If parameters are specified and used in a where clause, then you can customize the parameter value at compliance standard rule and target type level or compliance standard rule and target instance level. This enables the user to customize or control the check definition behavior per target instance or at the target type level.</p>
RuleParameter		Parameter definition
	ParamName	Name of parameter
	ParamType	parameter Type (String or number)
ParameterDefaultSettings		Default values for parameters
	ParamValue	<p>Define a default value for a parameter:</p> <ul style="list-style-type: none"> ■ ParamName: Name of parameter ■ CritThreshold/WarnThreshold/MinorWarnThreshold: Parameter default value. For critical severity, use CritThreshold. For warning severity, use WarnThreshold. For minor warning severity, use MinorWarnThreshold.
TestCondition		<p>The TestCondition tag operates over the data source fetched by running the Metric's <Source> Query. Any data source row that satisfies the condition is a violation to the rule</p>

Table 12–1 (Cont.) Key Tags for Defining Repository Rules

Tag	Subtag	Description
	ThresholdCriteria	Requires a column from Source Query, a threshold value, and operator (=,<,>, and so on) to relate the column value and threshold value.
	SqlWhereClauseCriteria	Requires a SQL condition over the columns from <Source> query. Optionally, this condition can include one or more parameters.
Severity		Severity of Rule (Critical, Warning, MinorWarning)
LifeCycleStatus		Lifecycle status of rule, (Development or Production)
UrlLink		Detail URL for the Rule, containing details about the rule
ViolationMessage		Message recorded with violation (for rule). Used for notifications. It provides the nlsid attribute to support the translation of messages
ClearViolationMessage		Message recorded with clearing of violation (for rule). Used for notifications. It provides the nlsid attribute to support the translation of messages.
KeywordList		List of keywords associated with the Rule.
	Keyword	Keywords applicable to the compliance standard
Author		Rule Author.

12.2.2 Defining Real-time Monitoring Rules

You can use a real-time monitoring rule to monitor any action that can happen against a file, a database object, or a Microsoft Windows Registry key. It can also monitor the starting and stopping of processes, and the logging in, logging out, and switching user (su) activity of users. The real-time aspect of the monitoring means that it captures the exact time the action occurred along with the user that performed the action.

Results from real-time monitoring can be reconciled with a Change Management system such as BMC Remedy. This reconciliation can automatically determine if an action was supposed to happen (authorized) or not (unauthorized). If a customer does not have a Change Management server, this audit status annotation can be made manually through the UI.

A major part of any IT compliance initiative is to ensure that your IT operations staff are making changes and managing the environment according to corporate policies. By reconciling what is happening in the environment to the customer's change management process, real-time monitoring helps to identify out-of-policy actions that will either lead to a high risk environment, or a compliance control that will fail audits.

Creating a real-time monitoring rule involves the following steps. These are explained further below.

- Choose the target type and entity type being monitored. A rule can also be limited based on certain target type properties (OS, target version, hardware platform, and target lifecycle)

- Choose one or more target type facets to monitor
- Choose one or more observations to watch for
- Choose zero or more facets to filter the results that are monitored
- Choose the change management reconciliation options

Integration points in this area include:

- Defining (one or more) facets for your target types. Facets define the low level artifacts that will be monitored
- Defining new compliance standard rules for new or existing compliance standards
- Mapping your compliance standard rules and compliance standards to the out-of-box compliance frameworks that are related to industry standard frameworks.
- Creating connectors to support new ticketing systems (including definition of custom region). This can also be used to extend out-of-box change reconciliation support (currently limited to Remedy 7.x). For information about the process for creating new connectors, see the *Oracle Enterprise Manager Connector Integration Guide*.

12.2.2.1 What Entity Types Can I Monitor?

When you define a real-time monitoring rule, the first thing you have to decide is what entity type on a host to monitor. For Oracle Enterprise Manager Cloud Control 12c, the following entity types can be monitored with Real-time Monitoring Rules:

- OS File
- OS Process
- OS User
- Microsoft Windows Registry
- Microsoft Active Directory User
- Microsoft Active Directory Computer
- Microsoft Active Directory Group
- Oracle Database Table
- Oracle Database View
- Oracle Database Procedure
- Oracle Database User
- Oracle Database Index
- Oracle Database Sequence
- Oracle Database Function
- Oracle Database Package
- Oracle Database Library
- Oracle Database Trigger
- Oracle Database Tablespace
- Oracle Database Materialized View
- Oracle Database Cluster

- Oracle Database Link
- Oracle Database Dimension
- Oracle Database Profile
- Oracle Database Public DB Link
- Oracle Database Synonym
- Oracle Database Public Synonym
- Oracle Database Segment
- Oracle Database Type
- Oracle Database Role
- Oracle Database SQL Query Statement

These entity types are fixed by the capabilities of the current release and cannot be extended. However, you can use them when creating facets and Real-Time monitoring rules.

In addition to facets defining what can be monitored, there is a set of entities that can be used for filtering also. The following list includes the most commonly used filtering entity types:

- OS Process
- OS User
- Oracle Database User
- Time Window
- Host

When you create a Real-time monitoring rule, choose what to monitor (that is, what files). Then choose if you want to use filtering so that only actions performed by certain users, or at certain periods of time are monitored.

12.2.2.2 About Real-time Monitoring Facets

Target Type Facets are used to specify the list of entities to monitor. These facets can be used again at a later time in any number of rules. They can be created on their own, or created inline with a Real-time Monitoring rule.

In the case of OS File monitoring, a facet could be a list of distinct single files, patterns with wildcards that would include many files, or simply an entire directory.

These patterns can also include parameters with a default, but can be overridden as required for each target.

The following are some examples of facets that may be defined for a HOST target type and an OS FILE entity type:

User Credential Files

- /etc/passwd
- /etc/shadow
- /etc/mail/trusted-users

Network Configuration Files

- /etc/hosts

- /etc/resolv.conf
- /etc/hosts.*
- /etc/defaultrouter
- /etc/nsswitch.conf
- /etc/netmasks
- {app_install_directory}/network/config

Here are some examples of facets that might be defined for a HOST target type and an OS PROCESS entity type. These might be monitored in real-time because any of these processes started on a production server could lead to a significant security risk.

Network Configuration Tools

- ifconfig
- xhost

The following table provides a list of hypothetical facets that you might create for your given target type. The facet name can be anything you choose. For some plug-in developers, there might be many more facets than these limited examples. For each facet, there is a description of the included patterns.

Target-Type Facet	Description
Log files	List each log file the target type has. Customers want to monitor when regular users modify a log file (not a system user)
Binary Files	List each binary the target type has. Rules can be created to monitor if a binary is tampered with or when a binary is patched. Instead of listing each individual binary, it can also list a whole directory, but exclude frequently changing files
Library Files	List each library the target type has. Rules can be created to monitor if a library is tampered with or when a library is patched. Instead of listing each individual library, it can also list a whole directory, but exclude frequently changing files
General Configuration Files	List any configuration files that are user changeable normally, but the user might want to capture changing.
Security Key Files	List any files that store certificates, keys, and so on. This can be a whole directory also, but exclude files that change regularly. This is to monitor if any users read the files in an attempt to get the content of the certificates.
Security Configuration Files	List any files that configure how security works in the target type, such as encryption configuration, and so on
Application Users	List the typical application users (that is, Oracle, root), and so on. Users can use this facet to filter monitoring changes where they do not care if the application user makes the change
Utility Processes	Any utility processes that normally run during a maintenance period, but should not be run during production
Registry Keys	Any Microsoft Windows registry keys that affect the configuration of the target

Target-Type Facet	Description
Configuration Tables	Any database tables that store configuration data.

12.2.2.3 Creating Real-time Monitoring Facets

This section provides an overview of the XML tags used in creating a real-time monitoring facet and an example of XML fragment showing facet creation. Facets can be created on their own as shown in this example, or inline with a real-time monitoring rule creation.

[Table 12–2](#) provides descriptions of the tags used to define a Real-time monitoring facet:

Table 12–2 Key Tags Used to Define a Real-Time Monitoring Facet

Tag	Subtag	Description
Name		The internal name of the facet. This must be unique across all facets that exist and is not visible on the UI.
DisplayName		The display name of the facet. It provides the nlsid attribute to support the translation of messages.
TargetType		The type of target to which this rule can be associated.
EntityType		The entity type for which you are creating the facet (such as <code>osfile</code> , <code>osprocess</code> , <code>osuser</code> , and so on)
IsSystem		True, for out-of-the-box rules. Otherwise, False.
Description		The description of the facet. It provides the nlsid attribute to support the translation of messages.
Author		The Enterprise Manager user that is the author of the facet.
LastUpdatedBy		The Enterprise Manager user that last updated the facet. This should be same as the author for your initially created data.
SourcePattern/GeneralPattern:		Container holding the pattern definition that makes up the facet

Table 12–2 (Cont.) Key Tags Used to Define a Real-Time Monitoring Facet

Tag	Subtag	Description
	Patterns/Pattern	<p>Collection of patterns that define the facet. A single facet can be made up of include and exclude patterns.</p> <ul style="list-style-type: none"> Value: An actual pattern. This pattern can include wildcards and parameters. Parameters are specified or bounded by { and }. Parameters must have a default value which is defined further down in the XML. The entity type determines the limitations on how wildcards are used. The product documentation outlines these limitations per entity type. Description: Description of the pattern. It provides the nlsid attribute to support the translation of messages. IsIncluded: Whether this pattern is an include pattern or exclude pattern. The notion of include or exclude is useful for wildcards. You can have a pattern which includes an entire directory, then you can exclude subdirectories or individual files under that included directory. A value of 0 indicates that this pattern is an exclude pattern. 1 indicates an include pattern.
	Parameters/Parameter	<p>Collection of pattern default values for each parameter introduced in the patterns. Parameters are not shared across facets. If you use the same parameter name in two facets, each facet must define its own default value.</p> <ul style="list-style-type: none"> Name: The parameter name used in the patterns defined above in the XML Value: The default value for this parameter. Users can override this parameter value per target when associating a Compliance Standard to one or more targets where this facet is in use. Description: Description of the parameter. It provides the nlsid attribute to support the translation of messages. isActive: Whether this parameter is currently in use in the list of patterns. This should always be 1 as you would not define new parameters without using them in the patterns

Example 12–3 Sample Facet Definition

```

<Facet xmlns="http://www.oracle.com/DataCenter/ConfigStd" is_time_window="0">
  <Name>network_configuration_files</Name>
  <DisplayName nlsid="SAMPLE_FACET_DNAME">Networking configuration
files</DisplayName>
  <TargetType>host</TargetType>
  <EntityType>osfile</EntityType>
  <IsSystem>1</IsSystem>
  <Description nlsid="SAMPLE_FACET_DESC">Files on a standard UNIX operating
system that contain configuration relevant to the networking
operations.</Description>

```

```

<Author>SYSMAN</Author>
<LastUpdatedBy>SYSMAN</LastUpdatedBy>
<SourcePattern>
  <GeneralPattern>
    <Patterns>
      <Pattern>
        <Value>{ETCDIR}/hosts</Value>
        <Description nlsid="SAMPLE_FACET_PATTERN_1_DESC">Contains IP
to hostname mappings</Description>
        <IsIncluded>1</IsIncluded>
      </Pattern>
      <Pattern>
        <Value>{ETCDIR}/resolv.conf</Value>
        <Description nlsid="SAMPLE_FACET_PATTERN_2_DESC">Contains
local name resolution mappings.</Description>
        <IsIncluded>1</IsIncluded>
      </Pattern>
      <Pattern>
        <Value>{ETCDIR}/appsecurity/*</Value>
        <Description nlsid="SAMPLE_FACET_PATTERN_3_DESC">All files in
a directory used for my custom application.</Description>
        <IsIncluded>1</IsIncluded>
      </Pattern>
      <Pattern>
        <Value>{ETCDIR}/appsecurity/sample.conf</Value>
        <Description nlsid="SAMPLE_FACET_PATTERN_4_DESC">Excluding one
file that is not a production configuration file that does not need to be
monitored.</Description>
        <IsIncluded>0</IsIncluded>
      </Pattern>

    <Parameters>
      <Parameter>
        <Name>ETCDIR</Name>
        <Description nlsid="SAMPLE_FACET_PARAMETER_1_DESC">Location
where all base Unix configuration files sit.</Description>
        <Value>/etc</Value>
        <IsActive>1</IsActive>
      </Parameter>
    </Parameters>
  </GeneralPattern>
</SourcePattern>
</Facet>

```

12.2.2.4 Creating Real-time Monitoring Facets for Time Windows

Time windows are a special type of facet that is used for filtering real-time monitoring. Typically, the Enterprise Manager end user creates time window facets since they are specific to their own operations schedules, but this document includes the content for reference purposes.

[Table 12–3](#) provides a description of the tags of a time window facet:

Table 12–3 Key Tags Used to Define a Time Window Facet

Tag	Subtag	Description
Name		The internal name of the facet. This must be unique across all facets that exist and is not visible on the UI.

Table 12–3 (Cont.) Key Tags Used to Define a Time Window Facet

Tag	Subtag	Description
DisplayName		The display name of the facet. It provides the nlsid attribute to support the translation of messages.
TargetType		The type of target to which this rule is associated
EntityType		The entity type for which you are creating the facet. For this example, it is timewindow
IsSystem		True, for out-of-the-box rules. Otherwise, False.
Description		The description of the facet. It provides the nlsid attribute to support the translation of messages.
Author		The Enterprise Manager user that is the author of the facet.
LastUpdatedBy		The Enterprise Manager user that last updated the facet. This should be same as the author for your initially created data.
SourcePattern/SchedulePattern	TZDisplayName	The display name of the time zone in English. For example Greenwich Mean Time (UTC+0).
	Duration	<ul style="list-style-type: none"> ■ DurStartMinute: The minute starting from 00:00 (midnight) when the time window starts. For example, 1439 = 11:59PM ■ DurEndMinute: The minute starting from 00:00 (midnight) when the time window ends ■ DurMinute: Precalculated duration that can be used for describing the time window, especially if it spans two days. A duration must be less or equal to 1440 (24 hours)
	Recurrence	<ul style="list-style-type: none"> ■ RecStartDate: The date that the time window starts ■ RecurrencePattern: <ul style="list-style-type: none"> RecPattern: The type of recurrence. Options are: "Single", "Daily", "Weekly", "Monthly", or "Yearly". RecPatternDays: Represents the days of the week, comma separated values. Sunday = 1, Saturday = 7. RecPatternFrequency: The frequency for repeating if the type of recurrence is to do "Every X of some pattern".

Example 12–4 Sample Time Window Facet Definition

```

<Facet is_time_window="1">
  <Name>general_working_hours</Name>
  <DisplayName>General Working Hours</DisplayName>
  <TargetType>host</TargetType>

```

```

<EntityType>timewindow</EntityType>
<IsSystem>0</IsSystem>
<Description>Define the work hour from 9:00 am to 5:00 pm</Description>
<Author>SYSMAN</Author>
<LastUpdatedBy>SYSMAN</LastUpdatedBy>
<SourcePattern>
  <SchedulePattern>
    <TZDisplayName/>
    <Duration>
      <DurStartMinute>540</DurStartMinute>
      <DurEndMinute>1020</DurEndMinute>
      <DurMinute>480</DurMinute>
    </Duration>
    <Recurrence>
      <RecStartDate>2010-07-26</RecStartDate>
      <RecurrencePattern>
        <RecPattern>WEEKLY</RecPattern>
        <RecPatternDays>1,2,5</RecPatternDays>
        <RecPatternFrequency>1</RecPatternFrequency>
      </RecurrencePattern>
    </Recurrence>
  </SchedulePattern>
</SourcePattern>
</Facet>

```

12.2.2.5 Creating Real-time Monitoring Rules

This section provides an overview of the XML tags used in creating a real-time monitoring rule and an example XML fragment showing rule creation. This XML fragment assumes that the facet has been created already and is referenced in this rule.

[Table 12–4](#) provides a description of the tags used to define a real-time rule:

Table 12–4 Key Tags Used to Define a Real-time Rule

Tag	Subtag	Description
DisplayName		The display name of the rule. It provides the <code>nlid</code> attribute to support the translation of messages.
TargetType		The type of target to which this rule is associated
IsSystem		True, for out-of-the-box rules. Otherwise, False.
Description		The description of the rule. It provides the <code>nlid</code> attribute to support the translation of messages.
Impact		Impact if the rule violates (when rule is noncompliant). It provides the <code>nlid</code> attribute to support the translation of messages
ViolationContext List		<p>Violation context defines a violation to a rule uniquely. Violation context lists columns from <code><Source> Query</code>, which will be visible as a part of the violation. Each column must mark as key or non-key.</p> <p>The mandatory <code>target_guid</code> column from <code><Source> query</code> is implicitly added to the violation context and should not be included in the violation context explicitly.</p>

Table 12–4 (Cont.) Key Tags Used to Define a Real-time Rule

Tag	Subtag	Description
CheckSource	Column	<p>Metric Column name. Uses Attributes name and type:</p> <ul style="list-style-type: none"> ■ DisplayLabel: Display name for column. It provides the nlsid attribute to support the translation of messages ■ IsHidden: True, if this column should not be displayed as a part of a violation context. Otherwise, False. ■ IsKey: True, if this column is a key.
		Defines the data source for Rule evaluation.
	RealTimeMonitoringLogicDefinition	Defines data source for Real-time Monitoring Rule
	EntityType	The type of monitoring performed (that is, <code>osfile</code> , <code>osprocess</code> , <code>osuser</code> , and so on). A full list is available in Section 12.2.2.1, "What Entity Types Can I Monitor?" .
	Facets	<p>The collection of facets to refer to in this rule. Some facets can be monitoring facets and some might be filtering facets.</p> <ul style="list-style-type: none"> ■ Facet Reference: The internal reference name of the facet <p>Name: The internal reference name of the facet that the rule refers to</p> <p>TargetType: The target type of the referenced facet. This should always be the same as the rule target type for your content.</p> <p>EntityType: The entity type of the referenced facet</p> <p>IsFilteredFacet: 0 indicates this facet reference is used to determine what to monitor. 1 indicates this facet reference is a filter.</p> <p>InvertedFilteredFacet: Only applicable if IsFilteredFacet=1. This specifies that the patterns in the facet definition are inverted (1) or not (0). If a filter facet was for "Production Hours" and then it was inverted, then monitoring will only occur outside of the pattern defined for "Production Hours"</p>
	ObservationTypes / ObservationType	<p>The types of observations you want to monitor in real-time.</p> <p>Name: Internal reference name for the observation type you want to have monitored in this rule.</p>

Table 12–4 (Cont.) Key Tags Used to Define a Real-time Rule

Tag	Subtag	Description
	Settings	<ul style="list-style-type: none"> CMSetting: Settings related to the performance of change management reconciliation. The auto_authorized=0 attribute indicates manual reconciliation only. 1 indicates integration with a change management server. Typically, you cannot use this field because the connector would not exist yet. If you create a rule without CM settings, a customer can override the rule and set their own custom CM settings after associating the rule with a Compliance Standard. CMConnector: The connector the rule should use for automatic reconciliation. AnnotateAuthObservation: Indicates whether the Change Management connector should annotate authorized observations into the requests that made the observations authorized.
	AdvancedSetting	<p>Advanced rule settings</p> <ul style="list-style-type: none"> GroupSetting: Settings about how observation bundles will be closed. Observation bundles collect a series of actions that happen against the same rule, by the same user, and on the same target over a short period of time. ObsGroupIdleTimeout: The timeout period (in minutes) after the last user action before a bundle will be closed. ObsGroupMaxAge: The maximum duration (in minutes) of an observation bundle. ObsGroupMaxObservations: The maximum number of observations in an observation bundle. GenerateEventByManualAuth: If you are using manual reconciliation, then you have the option of generating informational level events when observations occur. 1 indicates that an informational event will be created for each observation group. 0 indicates that no informational event will be created.
	Options	<p>Additional options that can be configured based on the entity type. Some entity types will not have options.</p> <ul style="list-style-type: none"> Option (Name/Value): A single name or value pair option setting.
Severity		Severity of the rule (Critical, Warning, or MinorWarning)
LifeCycleStatus		Lifecycle status of the rule (Development or Production)
UrlLink		Detail URL for the rule, containing details about the rule

Table 12–4 (Cont.) Key Tags Used to Define a Real-time Rule

Tag	Subtag	Description
ViolationMessage		Message recorded with violation (for the rule). Used for notifications. It provides the nlsid attribute to support the translation of messages.
ClearViolationMessage		Message recorded with clearing of violation (for the rule). Used for notifications. It provides the nlsid attribute to support the translation of messages.
KeywordList		List of keywords associated with the rule
	Keyword	Keywords applicable to the compliance standard
Author		Rule author

Example 12–5 Sample Rule Definition

```

<Rule xmlns="http://www.oracle.com/DataCenter/ConfigStd" Name="monitor_critical_
os_config_files">
  <DisplayName nlsid="SAMPLE_RULE_NAME">Monitor critical OS configuration
files</DisplayName>
  <TargetType>host</TargetType>
  <IsSystem>True</IsSystem>
  <Description nlsid="SAMPLE_RULE_DESC">Monitor several critical
configuration areas of a Linux host to ensure no configuration changes are
happening out of bounds. Monitoring is only done during production
hours.</Description>
  <Impact nlsid="SAMPLE_RULE_IMPACT">Capturing real-time changes to these files
may indicate a serious security issue.</Impact>
  <Recommendation nlsid="SAMPLE_RULE_RECO">Ensure that change management policy
documents how and when changes should be made in production.
Create compensating controls to address these out of bound
issues.</Recommendation>
  <ViolationContextList/>
  <CheckSource>
    <RealTimeMonitoringLogicDefinition>
      <EntityType>osfile</EntityType>
      <Facets>
        <FacetReference>
          <Name>network_configuration_files</Name>
          <TargetType>host</TargetType>
          <EntityType>osfile</EntityType>
          <IsFilteredFacet>0</IsFilteredFacet>
          <InvertFilteredFacet>0</InvertFilteredFacet>
        </FacetReference>
        <FacetReference>
          <Name>maild_configuration_files</Name>
          <TargetType>host</TargetType>
          <EntityType>osfile</EntityType>
          <IsFilteredFacet>0</IsFilteredFacet>
          <InvertFilteredFacet>0</InvertFilteredFacet>
        </FacetReference>
        <FacetReference>
          <Name>sshd_configuration_files</Name>
          <TargetType>host</TargetType>
          <EntityType>osfile</EntityType>
          <IsFilteredFacet>0</IsFilteredFacet>
          <InvertFilteredFacet>0</InvertFilteredFacet>
        </FacetReference>
      </Facets>
    </RealTimeMonitoringLogicDefinition>
  </CheckSource>
</Rule>

```

```

    <FacetReference>
      <Name>crontab_configuration_files</Name>
      <TargetType>host</TargetType>
      <EntityType>osfile</EntityType>
      <IsFilteredFacet>0</IsFilteredFacet>
      <InvertFilteredFacet>0</InvertFilteredFacet>
    </FacetReference>
    <FacetReference>
      <Name>kernel_configuration_files</Name>
      <TargetType>host</TargetType>
      <EntityType>osfile</EntityType>
      <IsFilteredFacet>0</IsFilteredFacet>
      <InvertFilteredFacet>0</InvertFilteredFacet>
    </FacetReference>
    <FacetReference>
      <Name>production_hours</Name>
      <TargetType>host</TargetType>
      <EntityType>timewindow</EntityType>
      <IsFilteredFacet>1</IsFilteredFacet>
      <InvertFilteredFacet>0</InvertFilteredFacet>
    </FacetReference>
  </Facets>
  <ObservationTypes>
    <ObservationType>
      <Name>osfile_create_suc</Name>
    </ObservationType>
    <ObservationType>
      <Name>osfile_content_modified_suc</Name>
    </ObservationType>
    <ObservationType>
      <Name>osfile_delete_suc</Name>
    </ObservationType>
    <ObservationType>
      <Name>osfile_content_mod_archive_suc</Name>
    </ObservationType>
  </ObservationTypes>
  <Settings>
    <CMSetting auto_authorized="0">
      <CMConnector></CMConnector>
      <AnnotateAuthObservation></AnnotateAuthObservation>
    </CMSetting>
    <AdvancedSetting>
      <GroupSetting>
        <ObsGroupIdleTimeout>15</ObsGroupIdleTimeout>
        <ObsGroupMaxAge>30</ObsGroupMaxAge>
        <ObsGroupMaxObservations>1000</ObsGroupMaxObservations>
      </GroupSetting>
      <GenerateEventByManualAuth>0</GenerateEventByManualAuth>
    </AdvancedSetting>
  </Settings>
  <Options>
    <Option value="10" name="osfile_archivenumber"/>
    <Option value="50000" name="osfile_polling_maxfilealert"/>
    <Option value="100" name="osfile_archive_maxsrcfilealert"/>
  </Options>
</RealTimeMonitoringLogicDefinition>
</CheckSource>
<Severity>MinorWarning</Severity>
<LifeCycleStatus>Development</LifeCycleStatus>
<KeywordList>

```

```

    <Keyword nlsid="CONFIGURATION">Configuration</keyword>
    <Keyword nlsid="SECURITY">Security</keyword>
  </KeywordList>
  <ViolationMessage nlsid="SAMPLE_RULE_VIOL_MSG">Violation due to change in
critical OS configuration files during production hours.</ViolationMessage>
  <ClearViolationMessage nlsid="SAMPLE_RULE_VIOL_CLRMSG">Cleared violation due
to change in critical OS configuration files during production
hours.</ClearViolationMessage>
  <Author>SYSMAN</Author>
</Rule>

```

12.3 Defining Compliance Standards

Compliance Standards are mapped to Compliance Standard Rules (Repository Rules or Real-time Monitoring Rules) in a hierarchical fashion.

[Example 12-6](#) provides the syntax for defining compliance standards and [Example 12-7](#) provides an example of a Compliance Standard Definition.

Note: For the complete compliance XSDs, see the following JAR file:

\$ORACLE_HOME/sysman/jlib/gccomplianceCommon.jar

See Also: For additional examples, see [Section 12.10, "More Compliance Examples"](#).

Example 12-6 Compliance Standard Definition Syntax

```

<xsd:complexType name="StandardT">
  <xsd:sequence>
    <xsd:element name="DisplayName" type="std:DisplayString128Def"
minOccurs="0"/>
    <xsd:element name="TargetType" type="std:Name128Def" minOccurs="1"
maxOccurs="1"/>
    <xsd:element ref="std:TargetPropertyFilter" minOccurs="0"/>
    <xsd:element name="Author" type="std:Name256Def" default="ORACLE"
minOccurs="0"/>
    <xsd:element name="Version" type="xsd:nonNegativeInteger" default="1"
minOccurs="0"/>
    <xsd:element name="LifecycleStatus" default="Development" minOccurs="0">
      <xsd:simpleType>
        <xsd:restriction base="xsd:string">
          <xsd:enumeration value="Development"/>
          <xsd:enumeration value="Production"/>
        </xsd:restriction>
      </xsd:simpleType>
    </xsd:element>
    <xsd:element name="IsHidden" type="std:BooleanDef" minOccurs="0"
default="false"/>
    <xsd:element name="IsSystem" type="std:BooleanDef" minOccurs="0"
default="false"/>
    <xsd:element name="IsAutoEnable" type="std:BooleanDef" minOccurs="0"
default="false"/>
    <xsd:element name="Description" type="std:DisplayString800Def"
minOccurs="0"/>
    <xsd:element name="KeywordList" type="std:KeywordListT" minOccurs="0"/>
  </xsd:sequence>
</complexType>

```

```

        <xsd:element name="ReferenceURL" type="std:String4000Def" minOccurs="0"/>
        <xsd:element name="FrontMatter" type="std:DisplayString800Def"
minOccurs="0"/>
        <xsd:element name="RearMatter" type="std:DisplayString800Def"
minOccurs="0"/>
        <xsd:element name="Notice" type="std:DisplayString800Def" minOccurs="0"/>
        <xsd:element name="Body" type="std:BodyT" minOccurs="0"/>
        <xsd:element name="ExtraInfo" type="xsd:string" minOccurs="0"
maxOccurs="1"/>
    </xsd:sequence>
    <xsd:attribute name="name" type="std:NameDef" use="required"/>
    <xsd:attribute name="oms_version" type="std:Name32Def" use="required"/>
</xsd:complexType>

```

Table 12–5 provides a description of the tags used in defining Compliance Standards:

Table 12–5 Key Tags Used in Defining Compliance Standards

Tag	Subtag	Description
DisplayName		The display name of the compliance standard. It provides the nlsid attribute to support the translation of messages. Note: The nlsid attribute is not applicable to metadata plug-ins.
TargetType		The type of target to which this compliance standard can be associated
Author		Compliance standard author
Version		The version of the compliance standard
LifecycleStatus		Lifecycle status of compliance standard (Development or Production)
IsSystem		True, if the compliance standard is provided out-of-the-box. Otherwise, False.
Description		Description of the compliance standard. It provides the nlsid attribute to support the translation of messages.
IsAutoEnable		If set to True, the compliance standard will be associated with all exiting targets for the defined target type. (Defined using TargetType)
KeywordList		A list of keywords applicable to the compliance standard
	Keyword	Keywords applicable to the compliance standard
ReferenceURL		The reference URL of the compliance standard
FrontMatter		Front matter message. It provides the nlsid attribute to support the translation of messages.
RearMatter		Rear matter message. It provides the nlsid attribute to support the translation of messages.
Notice		Notice message. It provides the nlsid attribute to support the translation of messages.
Body		Body of the compliance standard. Can have one or more of the following listed elements

Table 12–5 (Cont.) Key Tags Used in Defining Compliance Standards

Tag	Subtag	Description
	RuleFolder	<p>Defines a rule folder. A RuleFolder can have the following:</p> <p>RuleFolder</p> <p>RuleReference</p> <p>Include Standard Reference</p> <ul style="list-style-type: none"> ■ DisplayName: The display name of the Rule Folder. It provides the nlsid attribute to support the translation of messages. ■ Description: Description of the Rule Folder. It provides the nlsid attribute to support the translation of messages. <p>Note: The nlsid attribute is not applicable to metadata plug-ins. ■ ReferenceURL: The reference URL of the Rule Folder ■ Importance: Importance of Rule Folder (Low/Normal/High) </p>
	Include	Include another compliance standard reference to the including compliance standard
	RuleReference	Include rule reference to the compliance standard

Example 12–7 Sample Compliance Standard 1

```

<Standard xmlns="http://www.oracle.com/DataCenter/ConfigStd" oms_
version="12.1.0.1.0" name="sample_cs1">
  <DisplayName nlsid="SAMPLE_CS_1_NAME">Sample Compliance Standard
1</DisplayName>
  <TargetType>oracle_database</TargetType>
  <TargetPropertyFilter>
    <PropertyItem>
      <PropertyName>orcl_gtp_target_version</PropertyName>
      <ValueList>
        <Value>8.1.6+</Value>
      </ValueList>
    </PropertyItem>
  </TargetPropertyFilter>
  <Author>SYSTEM</Author>
  <Version>1</Version>
  <LifecycleStatus>Production</LifecycleStatus>
  <IsSystem>true</IsSystem>
  <Description nlsid="SAMPLE_CS_1_DESC">Sample Description</Description>
  <KeywordList>
    <Keyword nlsid="CONFIGURATION">Configuration</Keyword>
  </KeywordList>
  <ReferenceURL>http://sampleurl.com</ReferenceURL>
  <Body>
    <RuleFolder name="sample_RF_1">
      <DisplayName nlsid="SAMPLE_RF_1_NAME">Sample
Rulefolder</DisplayName>
      <Description nlsid="SAMPLE_RF_1_DESC">This includes rules
that checks for use of a single control file</Description>
      <ReferenceURL>http://www.oracle.com/db_rf1</ReferenceURL>
      <Importance>Normal</Importance>
    </RuleReference>
  </Body>
</Standard>

```

```
        <Name>sample_rule1</Name>
        <TargetType>oracle_database</TargetType>
        <Importance>Normal</Importance>
      </RuleReference>
    </RuleFolder>
  </Body>
</Standard>
```

12.4 Defining a Compliance Framework

Note: Although the Compliance Framework term is used throughout this document, the XML API uses the term `Group` or `SubGroup`. This is an internal name used for the XML structure that is not exposed on the Enterprise Manager UI.

[Example 12-8](#) provides the syntax for defining a compliance framework and [Example 12-9](#) on page 12-26 provides an example of a compliance framework definition.

Note: For the complete compliance XSDs, see the following JAR file:

`$ORACLE_HOME/sysman/jlib/gccomplianceCommon.jar`

See Also: For additional examples, see [Section 12.10, "More Compliance Examples"](#).

Example 12-8 Compliance Framework Definition Syntax

```
<xsd:complexType name="StandardGroupT">
  <xsd:sequence>
    <xsd:element name="DisplayName" type="std:DisplayString128Def"
minOccurs="0"/>
    <xsd:element name="Author" type="std:Name256Def" default="ORACLE"
minOccurs="0"/>
    <xsd:element name="Version" type="xsd:nonNegativeInteger" default="1"
minOccurs="0"/>
    <xsd:element name="LifecycleStatus" default="Development"
minOccurs="0">
      <xsd:simpleType>
        <xsd:restriction base="xsd:string">
          <xsd:enumeration value="Development"/>
          <xsd:enumeration value="Production"/>
        </xsd:restriction>
      </xsd:simpleType>
    </xsd:element>
    <xsd:element name="Description" type="std:DisplayString800Def"
minOccurs="0"/>
    <xsd:element name="KeywordList" type="std:KeywordListT" minOccurs="0"/>
    <xsd:element name="ReferenceURL" type="std:String4000Def" minOccurs="0"/>
    <xsd:element name="FrontMatter" type="std:DisplayString800Def"
minOccurs="0"/>
    <xsd:element name="RearMatter" type="std:DisplayString800Def"
minOccurs="0"/>
```

```

        <xsd:element name="Notice" type="std:DisplayString800Def" minOccurs="0"/>
        <xsd:element name="IsHidden" type="std:BooleanDef" minOccurs="0"
default="false"/>
        <xsd:element name="IsSystem" type="std:BooleanDef" minOccurs="0"
default="false"/>
        <xsd:element name="GroupBody" type="std:GroupBodyT" minOccurs="0"/>
        <xsd:element name="ExtraInfo" type="xsd:string" minOccurs="0"
maxOccurs="1"/>
    </xsd:sequence>
    <xsd:attribute name="name" type="std:NameDef" use="required"/>
    <xsd:attribute name="oms_version" type="std:Name32Def" use="required"/>
</xsd:complexType>

```

Table 12–6 provides a description of the tags used in defining a Compliance Framework:

Table 12–6 Key Tags Used in Defining a Compliance Framework

Tag	Subtag	Description
DisplayName		The display name of the compliance framework. It provides the <code>nlsid</code> attribute to support the translation of messages.
Author		Author of the compliance framework
Version		The version of the compliance framework
LifeCycleStatus		The lifecycle status of the compliance framework (Development or Production)
IsSystem		True, if compliance framework is provided out-of-the box. Otherwise, False.
Description		Description of compliance framework. It provides the <code>nlsid</code> attribute to support the translation of messages. Note: The <code>nlsid</code> attribute is not applicable to metadata plug-ins.
KeywordList		List of keywords applicable to compliance framework
	Keyword	Keywords applicable to the compliance standard
ReferenceURL		The reference URL of the compliance framework
FrontMatter		Front matter message. It provides the <code>nlsid</code> attribute to support the translation of messages
RearMatter		Rear matter message. It provides the <code>nlsid</code> attribute to support the translation of messages.
Notice		Notice message. It provides the <code>nlsid</code> attribute to support the translation of messages.
ExtraInfo		Additional information about the compliance framework.
GroupBody		Defines the body of the compliance framework. It can have one or more of the following elements:

Table 12–6 (Cont.) Key Tags Used in Defining a Compliance Framework

Tag	Subtag	Description
	SubGroup	<p>Defines a child framework element. A child framework element can include the following:</p> <p>Child framework</p> <p>Include Standard Reference.</p> <ul style="list-style-type: none"> ■ DisplayName: The display name of the child framework ■ Description: Description of the child framework ■ ReferenceURL: The reference URL of the child framework ■ Importance: Importance of child framework (Low, Normal, or High)
	StandardReference	Includes the compliance standard reference to the compliance framework

Example 12–9 Sample Compliance Framework

```

<StandardGroup xmlns="http://www.oracle.com/DataCenter/ConfigStd" name="sample_
csg" oms_version="12.1.0.1.0">
  <DisplayName nlsid="SAMPLE_CSG_NAME">Sample Compliance
Framework</DisplayName>
  <Author>SYSTEM</Author>
  <Version>1</Version>
  <LifecycleStatus>Production</LifecycleStatus>
  <Description nlsid="SAMPLE_CSG_DESC">Sample Description</Description>
  <KeywordList>
    <Keyword nlsid="SECURITY">Security</Keyword>
  </KeywordList>
  <ReferenceURL>http://sampleurl.com</ReferenceURL>
  <IsHidden>false</IsHidden>
  <IsSystem>true</IsSystem>
  <GroupBody>
    <SubGroup name="SampleSubgroup">
      <DisplayName nlsid="SAMPLE_CSG_SUBGROUP_NAME">Sample Child
Framework</DisplayName>
      <Description nlsid="SAMPLE_CSG_SUBGROUP_DESC">Sample Child
framework Description</Description>
      <ReferenceURL>http://sampleurl.com</ReferenceURL>
      <Importance>Normal</Importance>
      <StandardReference>
        <Name>sample_cs3</Name>
        <Author>SYSTEM</Author>
        <Version>1</Version>
        <Importance>Normal</Importance>
      </StandardReference>
    </SubGroup>
  </GroupBody>
</StandardGroup>

```

12.5 Defining Compliance Content

[Example 12–10](#) provides the syntax for defining compliance content and [Example 12–11](#) provides an example of XML compliance metadata.

See Also: For additional examples, see [Section 12.10, "More Compliance Examples"](#).

Example 12–10 Compliance Content Definition Syntax

```
<xsd:complexType name="ComplianceContentT">
  <xsd:sequence>

    <!-- Cumulative change since the first release.-->
    <xsd:element ref="std:ChangeList" minOccurs="0" maxOccurs="1"/>
    <!-- End Cumulative change since the first release -->

    <!-- Current state of entities -->
    <xsd:element ref="std:Facet" minOccurs="0" maxOccurs="unbounded"/>
    <xsd:element ref="std:Rule" minOccurs="0" maxOccurs="unbounded"/>
    <xsd:element ref="std:Standard" minOccurs="0" maxOccurs="unbounded"/>
    <xsd:element ref="std:StandardGroup" minOccurs="0" maxOccurs="unbounded"/>
    <!-- Current state of entities -->
  </xsd:sequence>
  <xsd:attribute name="oms_version" type="std:Name32Def" use="required"/>
  <xsd:attribute name="name" type="std:Name64Def" use="required"/>
  <!-- content_version of compliance content should be equal to version of last
change tag if any. -->
  <xsd:attribute name="content_version" type="std:Name64Def" use="optional"
default = "12.1.0.0.0"/>
  <xsd:attribute name="IsCompareEnabled" type="std:BooleanDef" use="optional"
default = "true"/>
</xsd:complexType>

<xsd:element name="ComplianceContent" type="std:ComplianceContentT"/>
```

Table 12–7 provides a description of some of the attributes used in defining compliance content:

Table 12–7 Compliance Content Attributes

Attribute	Description
oms_version	Version of Oracle Management Service (OMS)
name	Name of the compliance content
content_version	Version of the compliance content

Table 12–7 (Cont.) Compliance Content Attributes

Attribute	Description
IsCompareEnabled	<p>Specifies whether a rule or compliance standard is updated incrementally or if the entire rule or compliance standard is regenerated.</p> <p>Possible Values:</p> <ul style="list-style-type: none"> ■ True: For each rule and standard tag, the software finds the incremental change automatically and updates the entity incrementally. For example, if only one rule is updated in a compliance standard, only that rule is updated in the compliance standard and then the updated rule is reevaluated for all targets associated to the compliance standard at the time of the rule update (where the rule is a repository rule) ■ False: The user must specify <code><UpdateRule></code> within the <code><ChangeList><Change..>...</ChangeList></Change></code> tags. This causes the rule to be overridden (that is, all attributes and definitions). <p>Similarly, if a compliance standard is updated, it will override the standard completely and and regenerate results (in case of repository check-based standards).</p> <p>Note: If you set <code>isCompareEnabled = false</code>, then you must provide all the changes that were made in each version cumulatively since the compliance content was created. This is very important for metadata consistency.</p> <p>Oracle recommends that you always summarize the changes in each version even if the <code>isCompareEnabled</code> attribute is set to true. Because if you need to switch from <code>isCompareEnabled= true</code> (default) to <code>isCompareEnabled=false</code> at a future date, then all historical changes across different versions of the compliance content will be available to you.</p>

Example 12–11 Sample XML Compliance Metadata

```

<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<!DOCTYPE ComplianceContent [
<!ENTITY rule SYSTEM "SampleRuleThresholdCondition.xml">
<!ENTITY standard SYSTEM "SampleComplianceStandard.xml">
<!ENTITY standardgroup SYSTEM "SampleComplianceStandardGroup.xml">
]>
<ComplianceContent xmlns="http://www.oracle.com/DataCenter/ConfigStd" name="Sample
Compliance Framework" oms_version="11.2.0.1.0">
<ChangeList>
  <Change version="12.2.0.0.0">
    <UpdateRule>
      <RuleName>sample_rule</RuleName>
      <TargetType>oracle_database</TargetType>
    </UpdateRule>
    <UpdateStandardGroup>
      <StandardGroupName>sample_csg</StandardGroupName>
      <StandardGroupAuthor>SYSTEM</StandardGroupAuthor>
      <StandardGroupVersion>1</StandardGroupVersion>
    </UpdateStandardGroup>
  </Change>
</ChangeList>
&rule;
&standard;
&standardgroup;
</ComplianceContent>

```

12.6 Removing Compliance Content

To remove or delete compliance content, enter the following command:

```
emctl deregister oms metadata -sysman_pwd sysman -core -service gccompliance -file
DeleteComplianceContent.xml
```

In the previous command, *DeleteComplianceContent.xml* represents the name of the Delete Compliance Content XML file.

[Example 12–12](#) provides the syntax for defining Delete Compliance Content and [Example 12–13](#) provides an example of a Delete Compliance Content XML file.

Example 12–12 Delete Compliance Content Syntax

```
<!-- delete compliance metadata corresponding to the compliance content name
provided. -->
<xsd:complexType name="DeleteComplianceContentT">
  <xsd:attribute name="name" type="std:Name64Def" use="required"/>
</xsd:complexType>
<xsd:element name="DeleteComplianceContent" type="std:DeleteComplianceContentT"/>
```

Example 12–13 DeleteComplianceContent XML

```
<DeleteComplianceContent xmlns="http://www.oracle.com/DataCenter/ConfigStd"
name="Sample Compliance Framework" />
```

12.7 Supporting Translation

Note: Translation is supported for the Oracle Fusion Middleware plug-in only.

For each *nlsid* attribute in the XML samples, you must specify a Data Loading Format (DLF) map entry. A DLF file contains the English string for each defined *nlsid* attribute. These strings are available for translation.

Example 12–14 Sample DLF File

```
<?xml version="1.0" encoding="UTF-8"?>
<table xml:lang="en" name="MGMT_MESSAGES">

  <!-- lookup-key indicates which columns are used by
TransX to recognize a row as a duplicate -->
  <lookup-key>
    <column name="MESSAGE_ID"/>
    <column name="SUBSYSTEM"/>
    <column name="LANGUAGE_CODE"/>
    <column name="COUNTRY_CODE"/>
  </lookup-key>

  <!-- columns field indicates which columns will be loaded as
part of processing the dataset and which should be
translated by the Translation Group -->
  <columns>
    <column name="MESSAGE_ID" type="string" maxsize="64"/>
    <column name="SUBSYSTEM" type="string" maxsize="64"/>
```

```
<column name="LANGUAGE_CODE" type="string" language="%l"/>
<column name="COUNTRY_CODE" type="string" language="%Cs"/>
<column name="MESSAGE" type="string" maxsize="1000" translate="yes"/>
</columns>

<!-- dataset specifies the data to be loaded into the repository -->
<dataset>

  <row>
    <col name="MESSAGE_ID">SAMPLE_RULE_NAME</col>
    <col name="SUBSYSTEM">POLICY</col>
    <col name="MESSAGE">Sample Rule</col>
  </row>

  <row>
    <col name="MESSAGE_ID">SAMPLE_RULE_DESC</col>
    <col name="SUBSYSTEM">POLICY</col>
    <col name="MESSAGE">Checks for use of a single control file</col>
  </row>

  <row>
    <col name="MESSAGE_ID">SAMPLE_RULE_IMPACT</col>
    <col name="SUBSYSTEM">POLICY</col>
    <col name="MESSAGE">The control file is one of the most important files in an
Oracle database. It maintains many physical characteristics and important
recovery information about the database. If you lose the only copy of the control
file due to a media error, there will be unnecessary down time and other
risks.</col>
  </row>

  <row>
    <col name="MESSAGE_ID">SAMPLE_RULE_RECO</col>
    <col name="SUBSYSTEM">POLICY</col>
    <col name="MESSAGE">Use at least two control files that are multiplexed on
different disks.</col>
  </row>

  <row>
    <col name="MESSAGE_ID">SAMPLE_RULE_COL_1</col>
    <col name="SUBSYSTEM">POLICY</col>
    <col name="MESSAGE">FILE_LIST</col>
  </row>

  <row>
    <col name="MESSAGE_ID">SAMPLE_RULE_COL_2</col>
    <col name="SUBSYSTEM">POLICY</col>
    <col name="MESSAGE">CONTROL_FILE_COUNT</col>
  </row>

  <row>
    <col name="MESSAGE_ID">SAMPLE_RULE_VIOL_MSG</col>
    <col name="SUBSYSTEM">POLICY</col>
    <col name="MESSAGE">The database has an insufficient number of control
files.</col>
  </row>

  <row>
    <col name="MESSAGE_ID">SAMPLE_RULE_VIOL_CLEAR_MSG</col>
    <col name="SUBSYSTEM">POLICY</col>
    <col name="MESSAGE">The database has sufficient number of control files.</col>
  </row>
```



```

</row>

<!-- Standard NLSID Mappings -->

<row>
  <col name="MESSAGE_ID">SAMPLE_CS_NAME</col>
  <col name="SUBSYSTEM">POLICY</col>
  <col name="MESSAGE">Sample Compliance Standard</col>
</row>

<row>
  <col name="MESSAGE_ID">SAMPLE_CS_DESC</col>
  <col name="SUBSYSTEM">POLICY</col>
  <col name="MESSAGE">Sample Description</col>
</row>

<row>
  <col name="MESSAGE_ID">SAMPLE_RF_NAME</col>
  <col name="SUBSYSTEM">POLICY</col>
  <col name="MESSAGE">Sample Rulefolder</col>
</row>

<row>
  <col name="MESSAGE_ID">SAMPLE_RF_DESC</col>
  <col name="SUBSYSTEM">POLICY</col>
  <col name="MESSAGE">This includes rules that checks for use of a single
control file.</col>
</row>

<!-- Standard Group NLSID Mappings -->

<row>
  <col name="MESSAGE_ID">SAMPLE_CSG_NAME</col>
  <col name="SUBSYSTEM">POLICY</col>
  <col name="MESSAGE">Sample Compliance Framework</col>
</row>

<row>
  <col name="MESSAGE_ID">SAMPLE_CSG_DESC</col>
  <col name="SUBSYSTEM">POLICY</col>
  <col name="MESSAGE">Sample Description</col>
</row>

<row>
  <col name="MESSAGE_ID">SAMPLE_CSG_SUBGROUP_NAME</col>
  <col name="SUBSYSTEM">POLICY</col>
  <col name="MESSAGE">Sample Child Framework</col>
</row>

<row>
  <col name="MESSAGE_ID">SAMPLE_CSG_SUBGROUP_DESC</col>
  <col name="SUBSYSTEM">POLICY</col>
  <col name="MESSAGE">Sample Child Framework Description</col>
</row>

</dataset>
</table>

```

Note: If the DLF entry is for a real-time monitoring facet or pattern, then the subsystem is GCCOMPLIANCE_CCC. For all other rules, the subsystem is POLICY.

12.8 Packaging Compliance XML

This section indicates the location of the XML and DLF files.

- XML Files

Store all the XML files in the following directory:

plugin_stage/oms/metadata/gccompliance/

In the previous directory path, *plugin_stage* is the plug-in staging directory.

For more information about the plug-in staging directory, see [Section 13.2, "Staging the Plug-in"](#).

- DLF Files

Store all the DLF files in the following directory:

plugin_stage/oms/rsc/area/gccompliance

In the previous directory path, *plugin_stage* is the plug-in staging directory and *area* represents the subcomponent such as *db* for database or *ecm* for configuration management.

12.9 Setting Up and Testing Compliance Standards and Rules

To test your compliance standards or rules, do the following:

- [Install Compliance Content](#)
- [Test Compliance Standard](#)

12.9.1 Install Compliance Content

To install compliance content:

1. Use the following command to install the compliance content:

```
emctl register oms metadata -sysman_pwd password -core -service gccompliance  
-file ComplianceContent.xml
```

2. Submit the following job for back-end processing:

```
begin em_compliance_util.trigger_rule_dependency_job;end;
```

12.9.2 Test Compliance Standard

To test your compliance standard:

1. Log in to the Cloud Control console.
2. From the **Enterprise** menu, select **Compliance**, then select **Library**.
The Compliance Library page appears.
3. Click **Compliance Standards**.

4. Select the required compliance standard, then click **Associate Targets**.

The Target Association for Compliance Standard: *Compliance Standard Name* page appears, where *Compliance Standard Name* is the name of your selected compliance standard.

5. Click **Add**.

The Search and Select: Targets window appears.

6. Select the target that you want to evaluate, then click **Select**.

7. From the Target Association for Compliance Standard: *Compliance Standard Name* page, click **OK**.

8. Click **Yes** to the Save Association message.

The Compliance Standards page appears.

The previous steps trigger the evaluation, which occurs in a background job.

9. After a few minutes, from the **Enterprise** menu, select **Compliance**, then select **Results**.

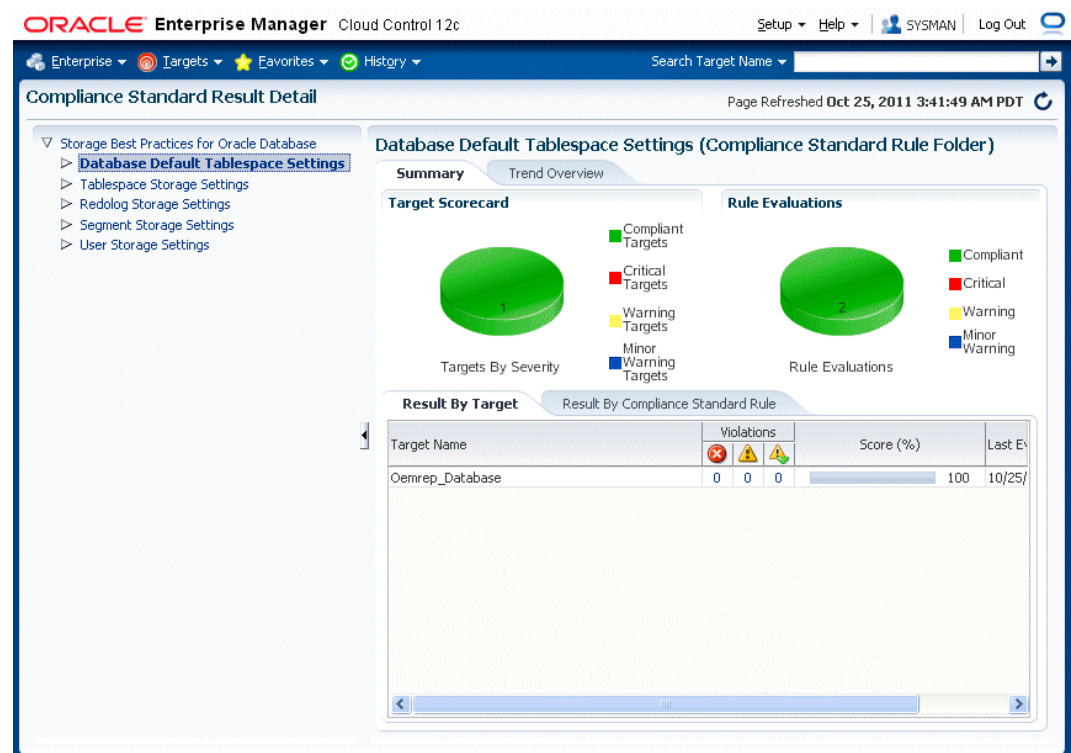
The Compliance Standards Evaluation Results page appears.

10. Select your compliance standard, then click **Show Details**.

The Compliance Standard Result Detail page appears.

11. From the left-hand side of the page, expand *Compliance Standard Name* to view any nodes, then click a node to view the results for that node.

Figure 12–1 Compliance Standard Result Detail



12.9.3 Constraints for Testing

Note the following constraints when you are testing your compliance standards or rules:

- The MGMT_VIEW user must have the SELECT privilege on the views used in the query
- target_guid must be one of the SELECT attributes in the query
- Alias names or select clause names must be less than 64 characters
- Ensure that the standard references from a compliance standard are imported first. Place the standard references first in the compliance content list.
- At least one column from the SELECT clause of the SQL source must be marked as a non-key column in the violation context definition and metric definition.
- The target_guid column must *not* be specified for violation context columns or for metric definitions.
- If the query references views from outside of the enclosing plug-in, then the views must be exposed by the EDK to the plug-in (at the plug-in EDK level).
- If the SQL source query of a repository rule refers to a PLSQL function, then ensure that it refers to global PLSQL functions only, and not package functions (that is, if those PLSQL functions depend on tables whose update triggers a rule evaluation). This is required to generate the list of tables which the rule evaluation outcome depends on correctly. Execute privileges must be granted to the mgmt_view user on this function.
- The target type of the rule included in a compliance standard must be the same as that of the immediate parent standard.
- Key columns of STRING type must contain less than 64 characters.

12.10 More Compliance Examples

This section provides additional examples of compliance content, rules, compliance standards, and compliance framework.

[Example 12-15](#) provides an example of compliance content version 1 and [Example 12-26](#) provides an example of compliance content version 2. Version 1 is the initial version of the compliance content. Note that the content version number is 12.1.0.1.0, while the content version in [Example 12-26](#) is 12.1.0.2.0.

Compliance content contains a `ChangeList` element. The `ChangeList` element describes the changes that have occurred since the first version of compliance content, such as updated rules, standards, and so on.

Example 12-15 Compliance Content Version 1

```
<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<!DOCTYPE ComplianceContent [
<!ENTITY rule1 SYSTEM "SampleRule1.xml">
<!ENTITY rule2 SYSTEM "SampleRule2.xml">
<!ENTITY rule3 SYSTEM "SampleRule3.xml">
<!ENTITY rule4 SYSTEM "SampleRule4.xml">
<!ENTITY rule5 SYSTEM "SampleRule5.xml">
<!ENTITY rule6 SYSTEM "SampleRule6.xml">
<!ENTITY standard1 SYSTEM "SampleComplianceStandard1.xml">
<!ENTITY standard2 SYSTEM "SampleComplianceStandard2.xml">
<!ENTITY standard3 SYSTEM "SampleComplianceStandard3.xml">
```

```

<!ENTITY standardgroup SYSTEM "SampleComplianceFramework.xml">]
]>
<ComplianceContent xmlns="http://www.oracle.com/DataCenter/ConfigStd"
name="SampleComplianceContent" oms_version="12.1.0.1.0" content_
version="12.1.0.1.0">
<ChangeList>
  <!-- ChangeList tag process each of the Change Tag with respect to the version
of the ComplianceContent installed in repository. -->
    <Change version="12.1.0.1.0">
      <!-- AddSubGroupWithinStandardGroup will introduce a subgroup within an
existing compliance framework/standard group in repository. -->
      <!-- AddStandardReferenceToStandardGroup will introduce a reference to a
standard within an existing compliance framework/standard group in repository.
-->
        <AddSubGroupWithinStandardGroup order="2">
          <StandardGroupName>oracle_pci</StandardGroupName>
          <StandardGroupAuthor>ORACLE</StandardGroupAuthor>
          <StandardGroupVersion>1</StandardGroupVersion>
          <SubGroup name="sampleSubgroup1">
            <DisplayName>sub1</DisplayName>
            <ReferenceURL>http://sampleAddedSubgroup.com</ReferenceURL>
            <Importance>High</Importance>
          </SubGroup>
        </AddSubGroupWithinStandardGroup>
        <AddStandardReferenceToStandardGroup>
          <StandardGroupName>oracle_pci</StandardGroupName>
          <StandardGroupAuthor>ORACLE</StandardGroupAuthor>
          <StandardGroupVersion>1</StandardGroupVersion>
          <SubGroupListInfo>
            <SubGroupElem>oracle_pci_ctrlobj_a</SubGroupElem>
          </SubGroupListInfo>
          <StandardReference>
            <Name>sample_cs1</Name>
            <Author>SYSTEM</Author>
            <Version>1</Version>
          </StandardReference>
        </AddStandardReferenceToStandardGroup>
      </Change>
    </ChangeList>
    <!--List of compliance standard rules -->
    &rule1;
    &rule2;
    &rule3;
    &rule4;
    &rule5;
    &rule6;
    <!--List of compliance standards -->
    &standard1;
    &standard2;
    &standard3;
    <!--List of compliance standard groups/frameworks-->
    &standardgroup;
  </ComplianceContent>

```

[Example 12-16](#) provides an example of a compliance rule that checks for use of a single control file

Example 12-16 Sample Rule 1

```

<Rule xmlns="http://www.oracle.com/DataCenter/ConfigStd" oms_version="12.1.0.1.0"

```

```
name="sample_rule1">
  <DisplayName nlsid="SAMPLE_RULE_1_NAME">Sample Rule 1</DisplayName>
  <TargetType>oracle_database</TargetType>
  <IsSystem>true</IsSystem>
  <TargetPropertyFilter>
    <PropertyItem>
      <PropertyName>orcl_gtp_operating_system</PropertyName>
      <ValueList>
        <Value>Windows</Value>
      </ValueList>
    </PropertyItem>
    <PropertyItem>
      <PropertyName>orcl_gtp_target_version</PropertyName>
      <ValueList>
        <Value>8.1.6+</Value>
      </ValueList>
    </PropertyItem>
  </TargetPropertyFilter>
  <Description nlsid="SAMPLE_RULE_1_DESC">Checks for use of a single control
file</Description>
  <Impact nlsid="SAMPLE_RULE_1_IMPACT">The control file is one of the most
important files in an Oracle database. It maintains many physical characteristics
and important recovery information about the database. If you lose the only copy
of the control file due to a media error, there will be unnecessary down time and
other risks.</Impact>
  <Recommendation nlsid="SAMPLE_RULE_1_RECO">Use at least two control files that
are multiplexed on different disks.</Recommendation>
  <ViolationContextList>
    <Column type="String" name="FILE_LIST">
      <DisplayLabel nlsid="SAMPLE_RULE_1_COL_1">FILE_LIST</DisplayLabel>
      <IsHidden>false</IsHidden>
      <IsKey>false</IsKey>
    </Column>
    <Column type="Number" name="CONTROL_FILE_COUNT">
      <DisplayLabel nlsid="SAMPLE_RULE_1_COL_2">CONTROL_FILE_
COUNT</DisplayLabel>
      <IsHidden>false</IsHidden>
      <IsKey>false</IsKey>
    </Column>
  </ViolationContextList>
  <CheckSource>
    <RepositoryCheckDefinition>
      <Metric>
        <TargetType>oracle_database</TargetType>
        <MetricName>sample_rule1</MetricName>
        <SourceType>SQL</SourceType>
        <Source>select CONTROL_FILE_COUNT, FILE_LIST, TARGET_GUID from MGMT$CS_DB_
CONTROL_FILE_COUNT</Source>
      <MetricColumnList>
        <MetricColumnInfo>
          <ColumnName>FILE_LIST</ColumnName>
          <ColumnType>String</ColumnType>
          <isKey>false</isKey>
          <ColumnLabel nlsid="SAMPLE_RULE_1_COL_1">FILE_LIST</ColumnLabel>
        </MetricColumnInfo>
        <MetricColumnInfo>
          <ColumnName>CONTROL_FILE_COUNT</ColumnName>
          <ColumnType>Number</ColumnType>
          <isKey>false</isKey>
          <ColumnLabel nlsid="SAMPLE_RULE_1_COL_2">CONTROL_FILE_
```

```

COUNT</ColumnLabel>
  </MetricColumnInfo>
</MetricColumnList>
  </Metric>
  <ParameterList>
    <RuleParameter>
      <ParamName>CONTROL_FILE_COUNT</ParamName>
      <ParamType>Number</ParamType>
    </RuleParameter>
  </ParameterList>
  <ParameterDefaultSettings>
    <ParamValue>
      <ParamName>CONTROL_FILE_COUNT</ParamName>
      <MinorWarnThreshold>1</MinorWarnThreshold>
    </ParamValue>
  </ParameterDefaultSettings>
  <TestCondition>
    <ThresholdCriteria>
      <ColumnName>CONTROL_FILE_COUNT</ColumnName>
      <TestOperator>EQ</TestOperator>
      <ThresholdValue>1</ThresholdValue>
      <ThresholdType>Number</ThresholdType>
    </ThresholdCriteria>
  </TestCondition>
</RepositoryCheckDefinition>
</CheckSource>
<Severity>MinorWarning</Severity>
<LifeCycleStatus>Production</LifeCycleStatus>
<KeywordList>
  <Keyword nlsid="CONFIGURATION">Configuration</Keyword>
</KeywordList>
  <ViolationMessage nlsid="SAMPLE_RULE_1_VIOL_MSG">The database has an
insufficient number of control files.</ViolationMessage>
  <ClearViolationMessage nlsid="SAMPLE_RULE_1_VIOL_CLEAR_MSG">The database has
sufficient number of control files.</ClearViolationMessage>
  <Author>SYSMAN</Author>
</Rule>

```

Example 12-17 provides an example of a sample compliance rule that checks for use of a single control file.

Example 12-17 Sample Rule 2

```

<Rule xmlns="http://www.oracle.com/DataCenter/ConfigStd" oms_version="12.1.0.1.0"
name="sample_rule2">
  <DisplayName nlsid="SAMPLE_RULE_2_NAME">Sample Rule 2</DisplayName>
  <TargetType>oracle_database</TargetType>
  <IsSystem>true</IsSystem>
<TargetPropertyFilter>
  <PropertyItem>
    <PropertyName>orcl_gtp_operating_system</PropertyName>
    <ValueList>
      <Value>Windows</Value>
    </ValueList>
  </PropertyItem>
  <PropertyItem>
    <PropertyName>orcl_gtp_target_version</PropertyName>
    <ValueList>
      <Value>8.1.6+</Value>
    </ValueList>
  </PropertyItem>

```

```
</PropertyItem>
</TargetPropertyFilter>
<Description nlsid="SAMPLE_RULE_2_DESC">Checks for use of a single control
file</Description>
<Impact nlsid="SAMPLE_RULE_2_IMPACT">The control file is one of the most
important files in an Oracle database.
It maintains many physical characteristics and important recovery information
about the database. If you lose the only copy of the control file due to a media
error, there will be unnecessary down time and other risks.</Impact>
<Recommendation nlsid="SAMPLE_RULE_2_RECO">Use at least two control files that
are multiplexed on different disks.</Recommendation>
<ViolationContextList>
  <Column type="String" name="FILE_LIST">
    <DisplayLabel nlsid="SAMPLE_RULE_2_COL_1">FILE_LIST</DisplayLabel>
    <IsHidden>>false</IsHidden>
    <IsKey>>false</IsKey>
  </Column>
  <Column type="Number" name="CONTROL_FILE_COUNT">
    <DisplayLabel nlsid="SAMPLE_RULE_2_COL_2">CONTROL_FILE_
COUNT</DisplayLabel>
    <IsHidden>>false</IsHidden>
    <IsKey>>false</IsKey>
  </Column>
</ViolationContextList>
<CheckSource>
  <RepositoryCheckDefinition>
    <Metric>
      <TargetType>oracle_database</TargetType>
      <MetricName>sample_rule2</MetricName>
      <SourceType>SQL</SourceType>
      <Source>select CONTROL_FILE_COUNT, FILE_LIST, TARGET_GUID from MGMT$CS_DB_
CONTROL_FILE_COUNT</Source>
      <MetricColumnList>
        <MetricColumnInfo>
          <ColumnName>FILE_LIST</ColumnName>
          <ColumnType>String</ColumnType>
          <isKey>>false</isKey>
          <ColumnLabel nlsid="SAMPLE_RULE_2_COL_1">FILE_LIST</ColumnLabel>
        </MetricColumnInfo>
        <MetricColumnInfo>
          <ColumnName>CONTROL_FILE_COUNT</ColumnName>
          <ColumnType>Number</ColumnType>
          <isKey>>false</isKey>
          <ColumnLabel nlsid="SAMPLE_RULE_2_COL_2">CONTROL_FILE_
COUNT</ColumnLabel>
        </MetricColumnInfo>
      </MetricColumnList>
    </Metric>
    <ParameterList>
      <RuleParameter>
        <ParamName>CONTROL_FILE_COUNT</ParamName>
        <ParamType>Number</ParamType>
      </RuleParameter>
    </ParameterList>
    <ParameterDefaultSettings>
      <ParamValue>
        <ParamName>CONTROL_FILE_COUNT</ParamName>
        <MinorWarnThreshold>1</MinorWarnThreshold>
      </ParamValue>
    </ParameterDefaultSettings>
  </RepositoryCheckDefinition>
</CheckSource>
</Metric>
</TargetPropertyFilter>
</PropertyItem>
```



```

        <TestCondition>
            <ThresholdCriteria>
                <ColumnName>CONTROL_FILE_COUNT</ColumnName>
                <TestOperator>EQ</TestOperator>
                <ThresholdValue>1</ThresholdValue>
                <ThresholdType>Number</ThresholdType>
            </ThresholdCriteria>
        </TestCondition>
    </RepositoryCheckDefinition>
</CheckSource>
<Severity>MinorWarning</Severity>
<LifeCycleStatus>Production</LifeCycleStatus>
<KeywordList>
    <Keyword nlsid="CONFIGURATION">Configuration</Keyword>
</KeywordList>
<ViolationMessage nlsid="SAMPLE_RULE_2_VIOL_MSG">The database has an
insufficient number of control files.</ViolationMessage>
<ClearViolationMessage nlsid="SAMPLE_RULE_2_VIOL_CLEAR_MSG">The database has
sufficient number of control files.</ClearViolationMessage>
<Author>SYSMAN</Author>
</Rule>

```

[Example 12-18](#) provides an example of a compliance rule that checks for use of a single control file.

Example 12-18 Sample Rule 3

```

<Rule xmlns="http://www.oracle.com/DataCenter/ConfigStd" oms_version="12.1.0.1.0"
name="sample_rule3">
    <DisplayName nlsid="SAMPLE_RULE_3_NAME">Sample Rule 3</DisplayName>
    <TargetType>oracle_database</TargetType>
    <IsSystem>true</IsSystem>
<TargetPropertyFilter>
    <PropertyItem>
        <PropertyName>orcl_gtp_operating_system</PropertyName>
        <ValueList>
            <Value>Windows</Value>
        </ValueList>
    </PropertyItem>
    <PropertyItem>
        <PropertyName>orcl_gtp_target_version</PropertyName>
        <ValueList>
            <Value>8.1.6+</Value>
        </ValueList>
    </PropertyItem>
</TargetPropertyFilter>
    <Description nlsid="SAMPLE_RULE_3_DESC">Checks for use of a single control
file</Description>
    <Impact nlsid="SAMPLE_RULE_3_IMPACT">The control file is one of the most
important files in an Oracle database.
It maintains many physical characteristics and important recovery information
about the database. If you lose the only copy of the control file due to a media
error, there will be unnecessary down time and other risks.</Impact>
    <Recommendation nlsid="SAMPLE_RULE_3_RECO">Use at least two control files that
are multiplexed on different disks.</Recommendation>
    <ViolationContextList>
        <Column type="String" name="FILE_LIST">
            <DisplayLabel nlsid="SAMPLE_RULE_3_COL_1">FILE_LIST</DisplayLabel>
            <IsHidden>>false</IsHidden>
            <IsKey>>false</IsKey>
        </Column>
    </ViolationContextList>

```

```

        </Column>
        <Column type="Number" name="CONTROL_FILE_COUNT">
            <DisplayLabel nlsid="SAMPLE_RULE_3_COL_2">CONTROL_FILE_
COUNT</DisplayLabel>
            <IsHidden>false</IsHidden>
            <IsKey>false</IsKey>
        </Column>
    </ViolationContextList>
    <CheckSource>
        <RepositoryCheckDefinition>
            <Metric>
                <TargetType>oracle_database</TargetType>
                <MetricName>sample_rule3</MetricName>
                <SourceType>SQL</SourceType>
                <Source>select CONTROL_FILE_COUNT, FILE_LIST, TARGET_GUID from MGMT$CS_DB_
CONTROL_FILE_COUNT</Source>
                <MetricColumnList>
                    <MetricColumnInfo>
                        <ColumnName>FILE_LIST</ColumnName>
                        <ColumnType>String</ColumnType>
                        <isKey>false</isKey>
                        <ColumnLabel nlsid="SAMPLE_RULE_3_COL_1">FILE_LIST</ColumnLabel>
                    </MetricColumnInfo>
                    <MetricColumnInfo>
                        <ColumnName>CONTROL_FILE_COUNT</ColumnName>
                        <ColumnType>Number</ColumnType>
                        <isKey>false</isKey>
                        <ColumnLabel nlsid="SAMPLE_RULE_3_COL_2">CONTROL_FILE_
COUNT</ColumnLabel>
                    </MetricColumnInfo>
                </MetricColumnList>
            </Metric>
            <ParameterList>
                <RuleParameter>
                    <ParamName>CONTROL_FILE_COUNT</ParamName>
                    <ParamType>Number</ParamType>
                </RuleParameter>
            </ParameterList>
            <ParameterDefaultSettings>
                <ParamValue>
                    <ParamName>CONTROL_FILE_COUNT</ParamName>
                    <MinorWarnThreshold>1</MinorWarnThreshold>
                </ParamValue>
            </ParameterDefaultSettings>
            <TestCondition>
                <ThresholdCriteria>
                    <ColumnName>CONTROL_FILE_COUNT</ColumnName>
                    <TestOperator>EQ</TestOperator>
                    <ThresholdValue>1</ThresholdValue>
                    <ThresholdType>Number</ThresholdType>
                </ThresholdCriteria>
            </TestCondition>
        </RepositoryCheckDefinition>
    </CheckSource>
    <Severity>MinorWarning</Severity>
    <LifecycleStatus>Production</LifecycleStatus>
    <KeywordList>
        <Keyword nlsid="CONFIGURATION">Configuration</Keyword>
    </KeywordList>
    <ViolationMessage nlsid="SAMPLE_RULE_3_VIOL_MSG">The database has an

```

```

insufficient number of control files.</ViolationMessage>
  <ClearViolationMessage nlsid="SAMPLE_RULE_3_VIOL_CLEAR_MSG">The database has
sufficient number of control files.</ClearViolationMessage>
  <Author>SYSMAN</Author>
</Rule>

```

Example 12–19 provides an example of a compliance rule that checks that no unintended ports are left open.

Example 12–19 Sample Rule 4

```

<Rule xmlns="http://www.oracle.com/DataCenter/ConfigStd" oms_version="12.1.0.1.0"
name="sample_rule4">
  <DisplayName nlsid="SAMPLE_RULE_4_NAME">Sample Rule 4</DisplayName>
  <TargetType>host</TargetType>
  <IsSystem>true</IsSystem>
  <Description nlsid="SAMPLE_RULE_4_DESC">Ensure that no unintended ports are
left open</Description>
  <Impact nlsid="SAMPLE_RULE_4_IMPACT">Open ports may allow a malicious user to
take over the host.</Impact>
  <Recommendation nlsid="SAMPLE_RULE_4_RECOMM">Do not open insecure
ports.</Recommendation>
  <ViolationContextList>
    <Column type="Number" name="port">
      <DisplayLabel nlsid="SAMPLE_RULE_4_PORT_COL">Port
Number</DisplayLabel>
      <IsHidden>false</IsHidden>
      <IsKey>true</IsKey>
    </Column>
  </ViolationContextList>
  <CheckSource>
    <RepositoryCheckDefinition>
      <Metric>
        <TargetType>host</TargetType>
        <MetricName>sample_rule4</MetricName>
        <SourceType>SQL</SourceType>
        <Source>SELECT target_guid, port as port, port as dummy FROM
MGMT$ESM_PORTS_LATEST</Source>
        <MetricColumnList>
          <MetricColumnInfo>
            <ColumnName>port</ColumnName>
            <ColumnType>Number</ColumnType>
            <isKey>true</isKey>
            <ColumnLabel nlsid="SAMPLE_RULE_4_LABEL">Port
Number</ColumnLabel>
          </MetricColumnInfo>
        </MetricColumnList>
      </Metric>
      <ParameterList>
        <RuleParameter>
          <ParamName nlsid="SAMPLE_RULE_4_DFLT_PORT_PNAME">DFLT_
PORT</ParamName>
          <ParamType>Number</ParamType>
        </RuleParameter>
      </ParameterList>
      <ParameterDefaultSettings>
        <ParamValue>
          <ParamName>DFLT_PORT</ParamName>
          <MinorWarnThreshold>655</MinorWarnThreshold>
        </ParamValue>
      </ParameterDefaultSettings>
    </RepositoryCheckDefinition>
  </CheckSource>

```

```

        </ParameterDefaultSettings>
        <TestCondition>
            <SqlWhereClauseCriteria>
                <WhereClause>:port &lt; :DFLT_PORT</WhereClause>
            </SqlWhereClauseCriteria>
        </TestCondition>
    </RepositoryCheckDefinition>
</CheckSource>
<Severity>Critical</Severity>
<LifecycleStatus>Production</LifecycleStatus>
<KeywordList>
    <Keyword nlsid="SECURITY">Security</Keyword>
</KeywordList>
<ViolationMessage nlsid="SAMPLE_RULE_4_MESG">The host is in an insecure state.
Port %port% is open.</ViolationMessage>
<ClearViolationMessage nlsid="SAMPLE_RULE_4_CLR_MESG">Port %port% is not
open.</ClearViolationMessage>
<Author>ORACLE</Author>
<LastUpdatedBy>&lt;SYSTEM&gt;</LastUpdatedBy>
</Rule>

```

Example 12-20 provides an example of a compliance rule that checks that no unintended ports are left open.

Example 12-20 Sample Rule 5

```

<Rule xmlns="http://www.oracle.com/DataCenter/ConfigStd" oms_version="12.1.0.1.0"
name="sample_rule5">
    <DisplayName nlsid="SAMPLE_RULE_5_NAME">Sample Rule 5</DisplayName>
    <TargetType>host</TargetType>
    <IsSystem>true</IsSystem>
    <Description nlsid="SAMPLE_RULE_5_DESC">Ensure that no unintended ports are
left open</Description>
    <Impact nlsid="SAMPLE_RULE_5_IMPACT">Open ports may allow a malicious user to
take over the host.</Impact>
    <Recommendation nlsid="SAMPLE_RULE_5_RECOMM">Do not open insecure
ports.</Recommendation>
    <ViolationContextList>
        <Column type="Number" name="port">
            <DisplayLabel nlsid="SAMPLE_RULE_5_PORT_COL">Port
Number</DisplayLabel>
            <IsHidden>false</IsHidden>
            <IsKey>true</IsKey>
        </Column>
    </ViolationContextList>
    <CheckSource>
        <RepositoryCheckDefinition>
            <Metric>
                <TargetType>host</TargetType>
                <MetricName>sample_rule5</MetricName>
                <SourceType>SQL</SourceType>
                <Source>SELECT target_guid, port as port, port as dummy FROM
MGMT$ESM_PORTS_LATEST</Source>
                <MetricColumnList>
                    <MetricColumnInfo>
                        <ColumnName>port</ColumnName>
                        <ColumnType>Number</ColumnType>
                        <isKey>true</isKey>
                        <ColumnLabel nlsid="SAMPLE_RULE_5_LABEL">Port
Number</ColumnLabel>
                    </MetricColumnInfo>

```

```

        </MetricColumnList>
    </Metric>
    <ParameterList>
        <RuleParameter>
            <ParamName nlsid="SAMPLE_RULE_5_DFLT_PORT_PNAME">DFLT_
PORT</ParamName>
            <ParamType>Number</ParamType>
        </RuleParameter>
    </ParameterList>
    <ParameterDefaultSettings>
        <ParamValue>
            <ParamName>DFLT_PORT</ParamName>
            <MinorWarnThreshold>655</MinorWarnThreshold>
        </ParamValue>
    </ParameterDefaultSettings>
    <TestCondition>
        <SqlWhereClauseCriteria>
            <WhereClause>:port < ; :DFLT_PORT</WhereClause>
        </SqlWhereClauseCriteria>
    </TestCondition>
    </RepositoryCheckDefinition>
</CheckSource>
<Severity>Critical</Severity>
<LifecycleStatus>Production</LifecycleStatus>
<KeywordList>
    <Keyword nlsid="SECURITY">Security</Keyword>
</KeywordList>
<ViolationMessage nlsid="SAMPLE_RULE_5_MESG">The host is in an insecure state.
Port %port% is open.</ViolationMessage>
<ClearViolationMessage nlsid="SAMPLE_RULE_5_CLR_MESG">Port %port% is not
open.</ClearViolationMessage>
<Author>ORACLE</Author>
<LastUpdatedBy>&lt;SYSTEM&gt;</LastUpdatedBy>
</Rule>

```

Example 12-21 provides an example of a compliance rule that checks that no unintended ports are left open.

Example 12-21 Sample Rule 6

```

<Rule xmlns="http://www.oracle.com/DataCenter/ConfigStd" oms_version="12.1.0.1.0"
name="sample_rule6">
    <DisplayName nlsid="SAMPLE_RULE_6_NAME">Sample Rule 6</DisplayName>
    <TargetType>host</TargetType>
    <IsSystem>true</IsSystem>
    <Description nlsid="SAMPLE_RULE_6_DESC">Ensure that no unintended ports are
left open</Description>
    <Impact nlsid="SAMPLE_RULE_6_IMPACT">Open ports may allow a malicious user to
take over the host.</Impact>
    <Recommendation nlsid="SAMPLE_RULE_6_RECOMM">Do not open insecure
ports.</Recommendation>
    <ViolationContextList>
        <Column type="Number" name="port">
            <DisplayLabel nlsid="SAMPLE_RULE_6_PORT_COL">Port
Number</DisplayLabel>
            <IsHidden>false</IsHidden>
            <IsKey>true</IsKey>
        </Column>
    </ViolationContextList>
    <CheckSource>

```

```

    <RepositoryCheckDefinition>
      <Metric>
        <TargetType>host</TargetType>
        <MetricName>sample_rule6</MetricName>
        <SourceType>SQL</SourceType>
        <Source>SELECT target_guid, port as port, port as dummy FROM
MGMT$ESM_PORTS_LATEST</Source>
        <MetricColumnList>
          <MetricColumnInfo>
            <ColumnName>port</ColumnName>
            <ColumnType>Number</ColumnType>
            <isKey>true</isKey>
            <ColumnLabel nlsid="SAMPLE_RULE_6_LABEL">Port
Number</ColumnLabel>
          </MetricColumnInfo>
        </MetricColumnList>
      </Metric>
      <ParameterList>
        <RuleParameter>
          <ParamName nlsid="SAMPLE_RULE_6_DFLT_PORT_PNAME">DFLT_
PORT</ParamName>
          <ParamType>Number</ParamType>
        </RuleParameter>
      </ParameterList>
      <ParameterDefaultSettings>
        <ParamValue>
          <ParamName>DFLT_PORT</ParamName>
          <MinorWarnThreshold>655</MinorWarnThreshold>
        </ParamValue>
      </ParameterDefaultSettings>
      <TestCondition>
        <SqlWhereClauseCriteria>
          <WhereClause>:port &lt; :DFLT_PORT</WhereClause>
        </SqlWhereClauseCriteria>
      </TestCondition>
    </RepositoryCheckDefinition>
  </CheckSource>
  <Severity>Critical</Severity>
  <LifeCycleStatus>Production</LifeCycleStatus>
  <KeywordList>
    <Keyword nlsid="SECURITY">Security</Keyword>
  </KeywordList>
  <ViolationMessage nlsid="SAMPLE_RULE_6_MESG">The host is in an insecure state.
Port %port% is open.</ViolationMessage>
  <ClearViolationMessage nlsid="SAMPLE_RULE_6_CLR_MESG">Port %port% is not
open.</ClearViolationMessage>
  <Author>ORACLE</Author>
  <LastUpdatedBy>&lt;SYSTEM&gt;</LastUpdatedBy>
</Rule>

```

Example 12-22 provides an example of a compliance standard that includes rules to check for use of a single control file.

Example 12-22 Sample Compliance Standard 1

```

<Standard xmlns="http://www.oracle.com/DataCenter/ConfigStd" oms_
version="12.1.0.1.0" name="sample_cs1">
  <DisplayName nlsid="SAMPLE_CS_1_NAME">Sample Compliance Standard
1</DisplayName>
  <TargetType>oracle_database</TargetType>

```

```

    <TargetPropertyFilter>
      <PropertyItem>
        <PropertyName>orcl_gtp_target_version</PropertyName>
        <ValueList>
          <Value>Windows</Value>
        </ValueList>
      </PropertyItem>
      <PropertyItem>
        <PropertyName>orcl_gtp_target_version</PropertyName>
        <ValueList>
          <Value>8.1.6+</Value>
        </ValueList>
      </PropertyItem>
    </TargetPropertyFilter>
    <Author>SYSTEM</Author>
    <Version>1</Version>
    <LifeCycleStatus>Production</LifeCycleStatus>
    <IsSystem>true</IsSystem>
    <Description nlsid="SAMPLE_CS_1_DESC">Sample Description</Description>
    <KeywordList>
      <Keyword nlsid="CONFIGURATION">Configuration</Keyword>
    </KeywordList>
    <ReferenceURL>http://sampleurl.com</ReferenceURL>
    <Body>
      <RuleFolder name="sample_RF_1">
        <DisplayName nlsid="SAMPLE_RF_1_NAME">Sample
Rulefolder</DisplayName>
        <Description nlsid="SAMPLE_RF_1_DESC">This includes rules that
checks for use of a single control file</Description>
        <ReferenceURL>http://www.oracle.com/db_rf1</ReferenceURL>
        <Importance>Normal</Importance>
        <RuleReference>
          <Name>sample_rule1</Name>
          <TargetType>oracle_database</TargetType>
          <Importance>Normal</Importance>
        </RuleReference>
      </RuleFolder>
    </Body>
  </Standard>

```

Example 12-23 provides an example of a compliance standard that includes rules to check for open unsecured ports.

Example 12-23 Sample Compliance Standard 2

```

<Standard xmlns="http://www.oracle.com/DataCenter/ConfigStd" oms_
version="12.1.0.1.0" name="sample_cs2">
  <DisplayName nlsid="SAMPLE_CS_2_NAME">Sample Compliance Standard
2</DisplayName>
  <TargetType>host</TargetType>
  <Author>SYSTEM</Author>
  <Version>1</Version>
  <LifeCycleStatus>Production</LifeCycleStatus>
  <IsSystem>true</IsSystem>
  <Description nlsid="SAMPLE_CS_2_DESC">Sample Description</Description>
  <KeywordList>
    <Keyword nlsid="SECURITY">Security</Keyword>
  </KeywordList>
  <ReferenceURL>http://sampleurl.com</ReferenceURL>
  <Body>

```

```

        <RuleFolder name="sample_RF_2">
          <DisplayName nlsid="SAMPLE_RF_2_NAME">Sample Rulefolder</DisplayName>
          <Description nlsid="SAMPLE_RF_2_DESC">This includes rules that checks
for open insecure ports.</Description>
          <ReferenceURL>http://www.oracle.com/db_rf1</ReferenceURL>
          <Importance>Normal</Importance>
          <RuleReference>
            <Name>sample_rule4</Name>
            <TargetType>host</TargetType>
            <Importance>Normal</Importance>
          </RuleReference>
        </RuleFolder>
      </Body>
    </Standard>

```

Example 12-24 provides an example of a compliance standard that includes rules to check for open unsecured ports.

Example 12-24 Sample Compliance Standard 3

```

<Standard xmlns="http://www.oracle.com/DataCenter/ConfigStd" oms_
version="12.1.0.1.0" name="sample_cs3">
  <DisplayName nlsid="SAMPLE_CS_3_NAME">Sample Compliance Standard
3</DisplayName>
  <TargetType>host</TargetType>
  <Author>SYSTEM</Author>
  <Version>1</Version>
  <LifecycleStatus>Production</LifecycleStatus>
  <IsSystem>true</IsSystem>
  <Description nlsid="SAMPLE_CS_3_DESC">Sample Description</Description>
  <KeywordList>
    <Keyword nlsid="SECURITY">Security</Keyword>
  </KeywordList>
  <ReferenceURL>http://sampleurl.com</ReferenceURL>
  <Body>
    <RuleFolder name="sample_RF_3">
      <DisplayName nlsid="SAMPLE_RF_3_NAME">Sample
Rulefolder</DisplayName>
      <Description nlsid="SAMPLE_RF_3_DESC">This includes rules that
checks for open insecure ports.</Description>
      <ReferenceURL>http://www.oracle.com/db_rf1</ReferenceURL>
      <Importance>Normal</Importance>
      <RuleReference>
        <Name>sample_rule5</Name>
        <TargetType>host</TargetType>
        <Importance>Normal</Importance>
      </RuleReference>
    </RuleFolder>
  </Body>
</Standard>
]]

```

Example 12-25 provides an example of a compliance framework.

Example 12-25 Sample Compliance Framework

```

<StandardGroup xmlns="http://www.oracle.com/DataCenter/ConfigStd" name="sample_
csg" oms_version="12.1.0.1.0">
  <DisplayName nlsid="SAMPLE_CSG_NAME">Sample Compliance
Framework</DisplayName>
  <Author>SYSTEM</Author>

```



```

<Version>1</Version>
<LifecycleStatus>Production</LifecycleStatus>
<Description nlsid="SAMPLE_CSG_DESC">Sample Description</Description>
<KeywordList>
  <Keyword nlsid="SECURITY">Security</Keyword>
</KeywordList>
<ReferenceURL>http://sampleurl.com</ReferenceURL>
<IsHidden>false</IsHidden>
<IsSystem>true</IsSystem>
<GroupBody>
  <SubGroup name="SampleSubgroup">
    <DisplayName nlsid="SAMPLE_CSG_SUBGROUP_NAME">Sample Child
Framework</DisplayName>
    <Description nlsid="SAMPLE_CSG_SUBGROUP_DESC">Sample Child
framework Description</Description>
    <ReferenceURL>http://sampleurl.com</ReferenceURL>
    <Importance>Normal</Importance>
    <StandardReference>
      <Name>sample_cs3</Name>
      <Author>SYSTEM</Author>
      <Version>1</Version>
      <Importance>Normal</Importance>
    </StandardReference>
  </SubGroup>
</GroupBody>
</StandardGroup>

```

Example 12-26 provides an example of compliance content.

Example 12-26 Compliance Content Version 2

```

<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<!DOCTYPE ComplianceContent [
<!ENTITY rule1 SYSTEM "SampleRule1.xml">
<!ENTITY rule2 SYSTEM "SampleRule2.xml">
<!ENTITY rule3 SYSTEM "SampleRule3.xml">
<!ENTITY rule5 SYSTEM "SampleRule5.xml">
<!ENTITY rule6 SYSTEM "SampleRule6.xml">
<!ENTITY standard1 SYSTEM "SampleComplianceStandard1.xml">
<!ENTITY standard3 SYSTEM "SampleComplianceStandard3.xml">
<!ENTITY standardgroup SYSTEM "SampleComplianceFramework.xml">
]>
<ComplianceContent xmlns="http://www.oracle.com/DataCenter/ConfigStd"
name="SampleComplianceContent" oms_version="12.1.0.1.0" content_
version="12.1.0.2.0">
<ChangeList>
  <!-- ChangeList tag process each of the Change Tag with respect to the version
of the ComplianceContent installed in repository. -->

  <Change version="12.1.0.1.0">

    <!-- AddSubGroupWithinStandardGroup/AddStandardReferenceToStandardGroup tags
will modify StandardGroup definition. -->
    <!-- AddSubGroupWithinStandardGroup will introduce a subgroup within an
existing compliance framework/standard group in repository. -->
    <!-- AddStandardReferenceToStandardGroup will introduce a reference to a
standard within an existing compliance framework/standard group in repository.
-->

```

```
<AddSubGroupWithinStandardGroup order="2">
  <StandardGroupName>oracle_pci</StandardGroupName>
  <StandardGroupAuthor>ORACLE</StandardGroupAuthor>
  <StandardGroupVersion>1</StandardGroupVersion>
  <SubGroup name="sampleSubgroup1">
    <DisplayName>sub1</DisplayName>
    <ReferenceURL>http://sampleAddedSubgroup.com</ReferenceURL>
    <Importance>High</Importance>
  </SubGroup>
</AddSubGroupWithinStandardGroup>
<AddStandardReferenceToStandardGroup>
  <StandardGroupName>oracle_pci</StandardGroupName>
  <StandardGroupAuthor>ORACLE</StandardGroupAuthor>
  <StandardGroupVersion>1</StandardGroupVersion>
  <SubGroupListInfo>
    <SubGroupElem>oracle_pci_ctrlobj_a</SubGroupElem>
  </SubGroupListInfo>
  <StandardReference>
    <Name>sample_cs1</Name>
    <Author>SYSTEM</Author>
    <Version>1</Version>
  </StandardReference>
</AddStandardReferenceToStandardGroup>
</Change>

<Change version="12.1.0.2.0">

  <!-- Delete will be remove rule/standard from repository if present,
else it will be noop. -->

  <DeleteStandard>
    <StandardName>sample_cs2</StandardName>
    <StandardAuthor>SYSTEM</StandardAuthor>
    <StandardVersion>1</StandardVersion>
  </DeleteStandard>

  <DeleteRule>
    <RuleName>sample_rule4</RuleName>
    <TargetType>host</TargetType>
  </DeleteRule>

  <!-- Entities with Update tag will override definitions if they exist in the
repository. -->
  <!-- Please note that if standard/rule is updated then old results are
replaced by new results based on standard/rule definition after update.
-->

  <UpdateRule>
    <RuleName>sample_rule5</RuleName>
    <TargetType>host</TargetType>
  </UpdateRule>
  <UpdateStandard>
    <StandardName>sample_cs3</StandardName>
    <StandardAuthor>SYSTEM</StandardAuthor>
    <StandardVersion>1</StandardVersion>
  </UpdateStandard>

  <UpdateStandardGroup>
    <StandardGroupName>sample_csg</StandardGroupName>
```

```

        <StandardGroupAuthor>SYSTEM</StandardGroupAuthor>
        <StandardGroupVersion>1</StandardGroupVersion>
    </UpdateStandardGroup>

    <!-- AddSubGroupWithinStandardGroup will introduce a subgroup within an
    existing compliance framework/standard group in repository. -->
    <!-- AddStandardReferenceToStandardGroup will introduce a reference to a
    standard within an existing compliance framework/standard group in repository.
    -->

    <AddSubGroupWithinStandardGroup order="2">
        <StandardGroupName>oracle_pci</StandardGroupName>
        <StandardGroupAuthor>ORACLE</StandardGroupAuthor>
        <StandardGroupVersion>1</StandardGroupVersion>
        <SubGroup name="sampleSubgroup2">
            <DisplayName>sub2</DisplayName>
            <ReferenceURL>http://sampleAddedSubgroup.com</ReferenceURL>
            <Importance>High</Importance>
        </SubGroup>
    </AddSubGroupWithinStandardGroup>
    <AddStandardReferenceToStandardGroup>
        <StandardGroupName>oracle_pci</StandardGroupName>
        <StandardGroupAuthor>ORACLE</StandardGroupAuthor>
        <StandardGroupVersion>1</StandardGroupVersion>
        <SubGroupListInfo>
            <SubGroupElem>oracle_pci_ctrlobj_a</SubGroupElem>
        </SubGroupListInfo>
        <StandardReference>
            <Name>sample_cs3</Name>
            <Author>SYSTEM</Author>
            <Version>1</Version>
        </StandardReference>
    </AddStandardReferenceToStandardGroup>
</Change>

</ChangeList>
<!--List of compliance standard rules -->
&rule1;
&rule2;
&rule3;
&rule5;
&rule6;
<!--List of compliance standards -->
&standard1;
&standard3;
<!--List of compliance standard groups/frameworks -->
&standardgroup;
</ComplianceContent>

```

12.11 Publishing Compliance Content Using Self Update

If you want to publish compliance content without having to deploy the plug-in, than use the Self Update console.

To publish and apply compliance content from the Self Update console:

1. Create a compliance content JAR file from the XML content using the following command:

```
-jar cvfM compliancecontent.jar compliance_content_files
```

Note: Similarly, multiple DLF files can be combined in a JAR file.

2. Create a manifest file to specify the name of the compliance content, label, and the version of the compliance content to be published. This manifest file specifies `compliancecontent.jar` and `compliancecdf.jar` in order respectively.

Example 12–27 Sample Manifest File

```
<?xml version="1.0" encoding="utf-8"?>
<tns:EntityInstance
xmlns:tns="http://www.oracle.com/EnterpriseGridControl/SelfUpdateManifest"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance" EntityType="param"
EntityTypeVersion="12.1.0.1.0" Vendor="Oracle" Maturity="TEST">

<tns:Description><![CDATA[<param>]]></tns:Description>

<tns:AttributeList>
<tns:Version>12.1.0.1.0</tns:Version>
<tns:Attribute Name="name" Value="<param>" Label="display_name"/>
</tns:AttributeList>
<tns:Readme><![CDATA[ <param>]]>
</tns:Readme>
<tns:CustomParamList/>
<tns:DependsOn/>
<tns:ArchiveList>
<tns:Archive Filename="param"/>
<tns:Archive Filename="param"/>
</tns:ArchiveList>
<tns:CustomData/>
</tns:EntityInstance>
```

3. Create a SAR (self update archive) file from the manifest file, `compliancecontent.jar`, and `compliancecdf.jar` using the following command:

```
edkutil prepare_update -manifest manifest_file_name -archivedir directory_
containing_compliancecontent.jar_and_compliancecdf.jar -out sar_file_name
```

Note: Before you import the SAR file into Enterprise Manager, make sure that the Software Library is configured.

For more information, see [Section 13.5.1.1, "Setting Up the Software Library"](#).

4. Import the SAR file into Enterprise Manager using the following command:

```
emcli import_update -omslocal -file=complete_path_to_sar_file
```

Note: Using the `-omslocal` flag means that the file must be placed on the Oracle Management Server (OMS) file system.

5. Log in to Enterprise Manager. From the **Setup** menu, select **Extensibility**, and then **Self Update**.

The Self Update page appears.

6. From the **Status** area, check that Downloaded Updates is set to 1 for Compliance Content.

7. In the Type column, click **Compliance Content**.

The Self Update: Compliance Content page appears.

8. Select the row with downloaded in the Status column, then **Apply**. Follow the steps in the wizard that appears.
9. From the Actions list, select **Apply** and check that the Status column reads succeeded.
10. Verify the imported compliance content from the Compliance Library. To view the Compliance Library, from the **Enterprise** menu, select **Compliance**, then select **Library**.

Validating, Packaging, and Deploying the Plug-in

This chapter contains the following sections:

- [Introduction to Validating, Packaging, and Deploying the Plug-in](#)
- [Staging the Plug-in](#)
- [Validating the Plug-in](#)
- [Creating the Plug-in Archive](#)
- [Importing and Deploying the Plug-in Archive into Enterprise Manager](#)
- [Adding a Target Instance](#)
- [Updating Deployed Metadata Files Using the Metadata Registration Service \(MRS\)](#)

13.1 Introduction to Validating, Packaging, and Deploying the Plug-in

As a plug-in developer, you are responsible for the following steps within the validation, packaging, and deployment process:

1. Create the staging directory (*plugin_stage*):
The staging directory structure defines the location of files as they will be deployed to Oracle Management Service and Management Agents.
For more information, see [Section 13.2, "Staging the Plug-in"](#).
2. Validate the plug-in.
Use the `empdk validate_plugin` command to validate the content of the plug-in once you have designed and developed it. This command verifies that the XML metadata files are compliant.
For more information, see [Section 13.3, "Validating the Plug-in"](#).
3. Create the Oracle Plug-in Archive (OPAR) file.
The plug-in archive is the standard way of distributing the plug-in for importing and deploying the plug-in across different installations of the Enterprise Manager Cloud Control.
For more information, see [Section 13.4, "Creating the Plug-in Archive"](#).
4. Import the OPAR into Enterprise Manager.

Use the `emcli import_update` command to import the plug-in into Enterprise Manager.

For more information, see [Section 13.5, "Importing and Deploying the Plug-in Archive into Enterprise Manager"](#).

5. Deploy the plug-in.

You must deploy a plug-in on the Oracle Management Service before it is used for monitoring targets.

For more information, see [Section 13.5.3, "Deploying the Plug-in on Oracle Management Service \(OMS\)"](#).

6. Add target instances for the deployed plug-in to monitor.

For more information, see [Section 13.6, "Adding a Target Instance"](#).

7. Use the Metadata Registration Service (MRS) to deploy updated metadata files.

The MRS allows you to upload one or more updated metadata files to the Oracle Management Service and Management Agents where your plug-in is deployed. The MRS registers the updated metadata files with Enterprise Manager, and overwrites the existing metadata with your updates.

For more information, see [Section 13.7, "Updating Deployed Metadata Files Using the Metadata Registration Service \(MRS\)"](#).

13.2 Staging the Plug-in

After you have created the plug-in files, the next step is to stage the plug-in in preparation for validation and packaging. The staging directory structure defines the location of files as they will be deployed to Oracle Management Service and Management Agents.

[Example 13–1](#) provides an example of the staging directory structure and [Table 13–1](#) describes the archive directory structure. Files are placed in the archive based on whether they are deployed to Oracle Management Service, Management Agents, or both. When the plug-in is deployed to an OMS instance or a Management Agent, the requisite files are copied to their respective directory locations.

Example 13–1 Plug-in Directory Structure

```

plugin_stage/
|
| plugin.xml
| agent/
|   |
|   | plugin_registry.xml
|   | default_collection/
|   |   |
|   |   | target_type.xml
|   | metadata/
|   |   |
|   |   | target_type.xml
|   | scripts/
|   |   |
|   |   | scripts
| oms/
|   |
|   | metadata/
|   |   |
|   |   |

```



```

default_collection/
    |
    target_type.xml
derivedAssoc/
    |
    derivedAssoc_rule.xml
discovery/
    |
    discovery.xml
gccompliance/
    |
    ComplianceContent_name.xml
jobTypes/
    |
    job_type.xml
mpcui/
    |
    mpcui.xml
reports/
    |
    report.xml
snapshotlive/
    |
    target-type_ecmdef.xml
targetType/
    |
    target_type.xml
discovery/
    |
    discovery scripts

```

Note: Use of the specified subdirectory names within the archive are not required, but are recommended by Oracle.

Table 13–1 File Locations in Plug-in Archive Structure

File	Directory	Notes
plugin.xml	plugin_stage/	Required. This file defines generic plug-in metadata that is deployed to Oracle Management Service. Place it at the root level within the archive structure. For more information, see Section 2.4, "Creating the plugin.xml File" .
plugin_registry.xml	plugin_stage/agent/	Required. This file defines metadata describing the plug-in used by the Management Agent. It must be placed at the top level of the /agent directory. For more information, see Section 2.5, "Creating the plugin_registry.xml File" .

Table 13–1 (Cont.) File Locations in Plug-in Archive Structure

File	Directory	Notes
<i>target_type.xml</i>	<i>plugin_ stage/oms/metadata/targetT ype/ plugin_ stage/agent/metadata/</i>	Required. This file defines metrics to be collected or computed for the target type. An identical copy of this file must be placed in both the /oms and /agent directories. For more information, see Section 3.3, "Creating the Target Type Metadata File" .
<i>default_ collections.xml</i>	<i>plugin_ stage/oms/metadata/default _collection/ plugin_ stage/agent/default_ collection/</i>	Required. This file defines metric collection parameters such as metric data collection frequency and default metric alert thresholds. An identical copy of this file must be placed in both the /oms and /agent directories. For more information, see Section 3.5, "Creating the Default Collection File" . Note: Oracle recommends that you name the default collections metadata file with the same file name as the target type metadata file.
<i>target-type_ ecmdef.xml</i>	<i>plugin_ stage/oms/metadata/snapsho tlive/</i>	Optional. Defines configuration data collection. For more information, see Section 6.3.1, "Defining Configuration Collection Tables" .
<i>job_type.xml</i>	<i>plugin_ stage/oms/metadata/jobType s/</i>	Optional. Place all job type definition files in the /jobTypes directory. For more information, see Chapter 7, "Adding Job Types" .
<i>report.xml</i>	<i>plugin_ stage/oms/metadata/reports /</i>	Optional. Put report definition files in the /reports directory.
<i>derivedAssoc_ rule.xml</i>	<i>plugin_ stage/oms/metadata/derived Assoc/</i>	Optional. Place the metadata file that defines the association derivation rules (or set of rules) in this directory. For more information, see Chapter 10, "Using Derived Associations" .
<i>ComplianceConte nt_name.xml</i>	<i>plugin_ stage/oms/metadata/gccompl iance/</i>	Optional. <i>ComplianceContent_ name.xml</i> contains references to compliance standards, rules, and frameworks. This directory can contain compliance rule, compliance standard, and compliance framework XML files. For more information, see Section 12.8, "Packaging Compliance XML" .

Table 13–1 (Cont.) File Locations in Plug-in Archive Structure

File	Directory	Notes
<i>compliance.dlf</i>	<i>plugin_ stage/oms/rsc/area/gccompl iance/</i>	Optional. Place all Data Loading Format (DLF) map entry (DLF) files associated with compliance rules or standards definitions in this directory.
<i>mpcui.xml</i>	<i>plugin_ stage/oms/metadata/mpcui/</i>	Optional. Place all management user interface metadata files in this directory. For more information, see Chapter 8, "Defining a Management User Interface" .
<i>discovery.xml</i>	<i>plugin_ stage/oms/metadata/discove ry/</i>	Optional. Place discovery metadata files in this location. For more information, see Section 11.4, "Packaging Discovery XML and Discovery Content" .
discovery script file(s)	<i>plugin_stage/discovery/</i>	Optional. Place the Perl scripts and JAR files (if any) that are required to perform automatic discovery in this location. For more information, see Section 11.4, "Packaging Discovery XML and Discovery Content" .
script file(s)	<i>plugin_ stage/agent/scripts/</i>	Optional. Put any scripts that will be deployed to Management Agents, such as metric collection scripts invoked by fetchlets, in this location. Use of the <code>/scripts</code> directory is not required but is recommended, as it allows use of the <code>%scriptsDir%</code> token from metric query descriptors defined in the <code>target-type.xml</code> file and in job type command references.

13.3 Validating the Plug-in

Plug-in validation should be done throughout the development cycle, and should certainly be done before packaging the plug-in. The `empdk validate_plugin` command can be used to validate the content of the plug-in once you have designed and developed it to verify that the XML metadata files are compliant. The tool is run against the specified plug-in staging directory and generates a report of any violations found. The format of the generated report can be specified using the `-format` option.

Usage is as shown below.

```
empdk validate_plugin -stage_dir staging directory
                        [-tmp_dir temporary working location]
                        [-out_dir output directory]
                        [-format (html|text|xml)]
                        [-conn_desc] - not used by external developers
                        [-repos_user Enterprise Manager repository owner]
                        [-debug [file to output debug information to]]
```

The following example validates the plug-in source files in the specified staging directory, and generates the validation report as a text file in the current working directory:

```
edk\bin>empdk validate_plugin -stage_dir C:\plugin_staging -format text
```

Table 13–2 provides the options that can be used to validate the plug-in.

Table 13–2 Options for Validating the Plug-in

Option	Description
-tmp_dir	Specify a temporary location to extract the plug-in files into. If not specified, it defaults to the current directory.
-out_dir	The directory the validation report file will generated into. If not specified, the report file will be generated in the current working directory.
-debug	Specify a file name where you want to store the debug information. If not specified, the default log file (validateplugin.logtime) will be created in the out directory and will store warning and error message only. If specified, then it will store all the debugging information to that log file
-format	The format the validation report will be generated in. If not specified, the report will be generated as a text file.

13.4 Creating the Plug-in Archive

After you have created the plug-in stage directory and validated the plug-in, the next step is to create an Oracle Plug-in Archive (OPAR) file. The OPAR file plays an important role at various stages of the plug-in lifecycle. It serves the following:

- The plug-in archive is the standard way of distributing the plug-in for importing and deploying the plug-in across different installations of the Enterprise Manager Cloud Control.
- You must test the plug-in being developed on an Enterprise Manager Cloud Control installation.

A plug-in is created by adding the files previously discussed to an OPAR using the Enterprise Manager Extensibility Development Kit (EDK). For more information about the EDK, see the *Oracle Enterprise Manager Cloud Control Extensibility Programmer's Guide*.

To create an OPAR, at the command prompt, enter the `empdk create_plugin` command. For more information about the `create_plugin` verb, see the command line help

The `empdk create_plugin` command syntax is as follows:

```
empdk create_plugin -stage_dir staging_dir -conn_desc repository_connection_string
-repos_user username [-repos_password repos_password]
-out_dir output_directory [-debug] [-force]
```

For example:

```
edk\bin>empdk create_plugin -stage_dir C:\pluginstagdir -conn_desc
myhost.us.example.com:25055:$ORACLE_SID -repos_user sysman -out_dir /tmp/plugins
```

Table 13–3 provides the options that can be used to create an OPAR:

Table 13–3 Options for Creating an OPAR

Option	Description
-tmp_dir	This option enables the command to create a temporary directory while executing. You can specify the path you want to use for this by providing a value following the option. Specify an existing directory or else you will receive an error. If not specified, then the out directory will be used for temporary location. If no out directory is specified, the current directory is the default.
-out_dir	The directory in which the plug-in archive (*.opar) file will be created. If not specified, the plug-in archive will be created in the current directory.
-debug	Specify a file name where you want to store the debug information. If not specified, the default log file (createplugin.logtime) will be created in the out directory and will store only warning and error messages. If specified, then it will store all the debugging information to that log file. This debugging information can be used to identify issues you may encounter while creating a plug-in. You can append the log created when you are filing a support request for a create plug-in related issue.
-force	If the out directory contains an OPAR with the same name, then you will be prompted to specify whether to overwrite the existing OPAR. If provided, it will automatically overwrite the existing OPAR. This is disabled by default.
-conn_desc	The connection descriptor that will connect to the Management Repository that the plug-in metadata will be written to when the plug-in is imported into Enterprise Manager. Specify the connection descriptor using the following syntax: host:port:sid For example: myhost.us.example.com:25055:\$ORACLE_SID
-repos_user	The user to connect to the Management Repository.

If the command runs successfully, a *plugin_version.plugin_id.opar* archive will be created in the directory where you ran this command.

If the command fails, an appropriate error message will be displayed. The parameters passed to the commands will vary from user to user and across systems where the plug-in is being created.

Some common mistakes while trying to create the plug-in archive are:

- If the path to the staging directory is entered incorrectly, which will throw a *File not found* or an *Input not found* exception.
- The *empdk command not found* exception will be shown if you have not changed your current directory to the expanded EDK directory.
- If disk where you are trying to create the OPAR has inadequate memory, an Input/Output related exception may occur.

13.5 Importing and Deploying the Plug-in Archive into Enterprise Manager

Once you have the plug-in archive ready with your *.opar file, you must import the plug-in into Enterprise Manager. Importing ensures that the content that you have created and packaged in the plug-in is available within Enterprise Manager.

Note: You must first import the plug-in before it can be deployed onto Enterprise Manager.

13.5.1 Prerequisites for Importing the Plug-in

The following prerequisites are met before importing the plug-in.

13.5.1.1 Setting Up the Software Library

1. Create a folder in the system where Enterprise Manager is installed. For example, `/net/hostname/scratch/aime/swlib1`.
2. From the console, select **Enterprise**, then **Provisioning and Patching**, and then **Software Library**.
3. Click **Actions**, then **Administration**.
4. Click **Add**.
5. In the pop up window, enter a name and location. For example, `swlib1` and `/net/hostname/scratch/aime/swlib1`. This should be the folder that you created in step 1.
6. Wait for the processing to finish.

13.5.1.2 Setting Up the EM CLI Utility

You will use the Enterprise Manager Command Line Utility, or EM CLI, to import the plug-in archive for deployment to Enterprise Manager.

- A page is provided in the Cloud Control console with instructions on setting up EM CLI. Access the page at the following URL:

```
https://em_host:em_port/em/console/emcli/download
```

For example:

```
https://emserver.acme.com:7799/em/console/emcli/download
```

- After setting up EM CLI, synchronize the EM CLI client with an Oracle Management Service (OMS):

```
emcli sync
```

After synchronization, all verbs and associated command line help available to this OMS become available at the EM CLI client.

13.5.2 Importing the Plug-in Archive

Once packaged, the plug-in must be imported into Enterprise Manager Cloud Control using the `emcli import_update` command. You have two options depending on where EM CLI is installed:

- If EM CLI is on the same system as the system where you created the plug-in archive (*.opar file), then run the following command.

```
emcli import_update
-file="<path to *.opar file you created>"
-omslocal
```

The `-omslocal` flag indicates that the plug-in archive is on the same system where you are running this command and the path exists on this system.

For example:

```
emcli import_update -file=/tmp/sample_plugin.opar -omslocal
```

- If you are running EM CLI on a different system than the system where you created the plug-in archive (*.opar file), then run the following command:

```
emcli import_update
-file="path to the .opar file"
-host="host name of plug-in host"
-credential_name="credential for plug-in host"
-credential_owner="credential owner on the plug-in host"
```

where:

- `-file`: The absolute path to the *.opar file on the system where you created the archive.
- `-host`: The host name for the host target where the file is available.
- `-credential_name`: The name of the credentials on the remote system you are connecting to.
- `-credential_owner`: The owner of the credentials on the host system you are connecting to.

For example:

```
emcli import_update -file=/tmp/sample_plugin.opar
-host="host1.acme.com" -credential_name="myOracleCred"
-credential_owner="mypassword"
```

- As an alternative to the previous step, you can also run the following command:

```
emcli import_update
-file="path to *.opar file you created"
-host="hostname"
-credential_set_name="setname"
```

`-credential_set_name`: The set name of the preferred credential stored in the Management Repository for the host target. It can be one of the following:

- `HostCredsNormal`: The default unprivileged credential set.
- `HostCredsPriv`: The privileged credential set.

13.5.3 Deploying the Plug-in on Oracle Management Service (OMS)

A plug-in must be deployed on Oracle Management Service (OMS) before it can be used to monitor targets. Follow the steps below to deploy the plug-in on Enterprise Manager Cloud Control.

Note: Plug-ins must be deployed on Oracle Management Service before being deployed on Management Agents.

Plug-ins for specific target types are deployed automatically on Management Agents that will monitor those target types. For more information, see [Section 13.6, "Adding a Target Instance"](#).

To deploy a plug-in on the Oracle Management Server:

1. From the **Setup** menu, select **Extensibility**, then **Plug-ins**.

Enterprise Manager displays the list of plug-ins that have been downloaded and can be deployed on the Plug-ins page.

2. On the Plug-ins page, select the specific plug-in you want to deploy.
3. Click **Deploy On**, then select **Management Servers**.

Ensure that dependent plug-ins are deployed and that all existing Management Agents are compatible with the version of the specified plug-in. Enterprise Manager prompts for credentials if the Management Agent is not available.

4. On the **Deploy Plug-in** window, enter the required details. Note that you will require the Management Repository SYS user password to complete the deployment process.

From the Version list, select the Plug-in version. The **Target Type** information is displayed in the table. Enter the **Repository sys Password**, then click **Continue**.

5. Proceed through the steps in the Deploy Plug-in windows.
6. Click **Deploy** to deploy the selected plug-in on all Enterprise Manager servers.
Enterprise Manager displays a page that allows you to monitor the deployment status. Enterprise Manager deploys the selected plug-in on all Enterprise Manager Servers.

You can also monitor the deployment status by going to the Enterprise Manager Cloud Control console, then going to the Plug-ins page as described in step 1, selecting the plug-in and select the **Recent Deployment Activities** tab at the bottom of the page for the selected plug-ins. This bottom section also lets you see details of your plug-in, which includes the plug-in ID, version, vendor, and so on.

If any of the steps during plug-in deployment fails, the log file is available in \$ORACLE_HOME/cfgtoollogs/pluginca/*. Append these files when logging a support request for failure while deploying the plug-in. You can also use them to debug the problem.

13.5.4 Important Details Regarding Plug-in Deployment

- You can import multiple versions of the same plug-in. The version to deploy can be selected from a list if you are using Cloud Control to deploy the plug-in, or can be specified on the command line if using EM CLI.
- Only one version can be deployed on the Oracle Management Service (OMS) at any given time. If a later version has been deployed previously, it cannot be downgraded to an earlier version.
- Updating a plug-in to a new version does not remove the content of the older plug-in(s) that were imported.

- The Management Agent can have the same or earlier version of the plug-in that is deployed on the OMS. A version later than the version on the OMS is not allowed on the Management Agent.
- The plug-in on OMS and the Management Agent can be updated independently of each other as long as the version on the OMS is the latest version.
- Available updates are visible on the Plug-ins page. They can be downloaded from the Enterprise Manager store or imported using EM CLI as described in [Section 13.5.2, "Importing the Plug-in Archive"](#).

13.6 Adding a Target Instance

When the plug-in is deployed on OMS, it is ready to monitor target instances.

Note: In the current Cloud Control release, deployment of a plug-in to a Management Agent that will monitor targets is no longer required. Instead, the plug-in for a specific target type is automatically deployed with the Management Agent that will monitor targets of that type.

This is a significant change from previous releases, in which plug-ins had to be manually deployed to a Management Agent first. Then a target instance had to be added to the Management Agent manually.

You can add targets that the plug-in will monitor through Enterprise Manager Cloud Control by selecting **Add Targets** from the **Setup** menu. The process for adding targets - known in Cloud Control terminology as *target promotion* - varies depending on the option you choose.

You can also add a target instance using the EM CLI utility. Open a command prompt and run the following command:

```
emcli add_target
  -name="name"
  -type="type"
  -host="hostname"
  [-properties="pname1:pval1;pname2:pval2;..."]...
  [-separator=properties="sep_string"]
  [-subseparator=properties="subsep_string"]
  [-credentials="userpropname:username;pwdpropname:password;..."]
  [-input_file="parameter_tag:file_path"]
  [-display_name="display name"]
  [-groups="groupname1:grouptype1;groupname2:grouptype2;..."]...
  [-timezone_region="gmt offset"]
  [-monitor_mode="monitor mode"]
  [-instances="rac database instance target name1:target type1;..."]
]
```

For example:

```
emcli add_target
  -name="cluster_database"
  -type="rac_database"
  -host="myhost.us.example.com"
  -monitor_mode="1"
  -properties="ServiceName:service.us.example.com;ClusterName:newdb_
cluster"
  -instances="database_inst1:oracle_database;database_inst2:oracle_
database"
```

Use the `emcli help add_target help` command to see more options when adding the target instance.

If targets have been added before, they will be promoted and monitored by the plug-in after it is deployed.

13.7 Updating Deployed Metadata Files Using the Metadata Registration Service (MRS)

As part of the plug-in development process, you will package your plug-in as an archive and deploy it to an Enterprise Manager Cloud Control installation to test it. However, you will likely not want to re-deploy the plug-in each time you make changes to various metadata files.

Note: You must update the metadata version each time you update a metadata file.

The Metadata Registration Service (MRS) allows you to upload one or more updated metadata files to the Oracle Management Service and Management Agents your plug-in has been deployed to. The updated metadata files will be registered with Enterprise Manager, and will overwrite the existing metadata with your updates.

Note: For target types and default collections, some additional steps are required for using MRS, see [Section 13.7.1, "Target Types and Default Collections"](#).

This service is invoked through the `emctl register oms metadata` command. The syntax is as follows:

```
emctl register oms metadata -service Metadata Service Id (-file metadata file to register | -file_list file containing list of files to register)
                        [-core | -pluginId Plugin Id] [-sysman_pwd "sysman password"]
```

For example, the following command registers changes to target type metadata file:

```
emctl register oms metadata -service targetType -file /staging/demo_hostsample.xml
-pluginId acme.demo.hostsample -sysman_pwd myempassword
```

[Table 13–4](#) describes the usage of the command.

Table 13–4 *emctl Command Usage*

Option	Required Y/N	Description
-service	Y	Specify the type of metadata to register. Values are: <ul style="list-style-type: none"> ■ targetType: Specify for target type metadata. ■ default_collection: Specify for default collection metadata. ■ LiveSnapshotRegistration: Specify for configuration metadata registration ■ CredStoreMetadata: Specify for ■ jobTypes: Specify for ■ report: Specify for Report Metadata Registration ■ bipublisherreport: Specify for BI Publisher report metadata. ■ discovery: Specify for discovery metadata. ■ derivedAssocs: Specify for associations metadata. ■ gccompliance: Specify for compliance rules, compliance standards, and compliance framework metadata. ■ mpcui: Specify for management user interface metadata.
-file	N	The path and file name for a single metadata file to upload and register. Either -file or -file_list can be included.
-file_list	N	The path and file name for a file containing a list of metadata file paths (one on each line).
-core	N	Not valid for plug-in development.
-pluginId	N	The unique three-part identifier given to the deployed plug-in to update. See Section 2.2.1, "Defining the Plug-in ID" for details.
-sysman_pwd	Y	The Enterprise Manager user password.

13.7.1 Target Types and Default Collections

For target types and default collections, some additional steps are required for using MRS if there are existing targets of this target type.

If you have an existing target and you want to update the metadata files during the plug-in development process, follow these steps:

1. Register the new metadata files using the `emctl register oms metadata` command.

```
emctl register oms metadata -service targetType -file full
path/TargetTypeMetadata.xml -pluginId Plugin Id -sysman_pwd sysman
```

```
emctl register oms metadata -service storeTargetType -file full
path/TargetTypeMetadata.xml -pluginId Plugin Id -sysman_pwd sysman
```

```
emctl register oms metadata -service default_collection -file full
path/TargetTypeCollection.xml -pluginId Plugin Id -sysman_pwd sysman
```

```
emctl register oms metadata-service systemStencil -file full  
path/TargetTypeStencil.xml -pluginId Plugin Id -sysman_pwd sysman
```

2. Place the metadata XML files in the correct directories of the plug-in home directory (*PLUGIN_AGENT_HOME*) in the Management Agent as shown. The *PLUGIN_AGENT_HOME* directory is created when the plug-in is deployed to the Management Agent. The default location is *AGENT_BASE_DIR/plugins*.

```
$PLUGIN_AGENT_HOME/metadata/  
$PLUGIN_AGENT_HOME/default_collection
```

3. Restart the Management Agent.

```
AGENT_HOME/agent/bin/emctl stop agent  
AGENT_HOME/agent/bin/emctl start agent
```

In the preceding command, *AGENT_HOME* represents the Management Agent home directory.

Defining Software Library Metadata

The chapter introduces Software Library framework, and describes how you can define and register metadata that is used by plug-in integrators, to extend a Software Library to include extensions, and Out-of-box entities. After registering the custom extensions, you can use interfaces like Software Library console, EM CLI, Action Script API, and so on, to create, manage, and access Software Library entities. Primarily, the chapter contains the following topics:

- [Introduction to Software Library Framework](#)
- [Defining Software Library Metadata](#)
- [Organizing Software Library Metadata Files](#)
- [Adding the Software Library Metadata to Enterprise Manager](#)
- [Using Software Library Entities](#)

Note: For conceptual information about any of the topics covered in this chapter, see *Oracle Enterprise Manager Cloud Control Extensibility Programmer's Guide*.

14.1 Introduction to Software Library Framework

After defining and registering the new metadata, you can extend the Software Library framework to become aware of new types of software or scripts or configurations. You can use these entities in the existing or custom automation logic represented by a job type or deployment procedure.

As a plug-in developer, you will be able to define the following Software Library metadata:

- **Folder**

Software Library allows you to organize the different user-defined or out-of-box entities into logical folders for efficient management.

For more information, see [Section 14.2.1](#).

- **Type and Subtype**

A type and subtype together define the common traits of the entities it represents in terms of common and searchable metadata/configuration properties, their default values, file association requirements, etc. All entities in Software Library have an associated type and a subtype.

For more information, see [Section 14.2.2](#) and [Section 14.2.3](#).

- **Entity**

Entities are the primary artifacts stored in Software Library. They are identified by the type/subtype they have been created with and the folder they are present in.

For more information, see [Section 14.2.4](#).

14.2 Defining Software Library Metadata

This section contains the following topics:

- [Defining Folders](#)
- [Defining Types](#)
- [Defining Subtypes](#)
- [Defining Entities](#)

Important: To view an example which describes how to use the Software Library artifacts, follow these steps:

1. In Cloud Control, from **Setup** menu, select **Extensibility**, and click **Development Kit**.
2. On the Extensibility Development Kit page, in the Getting Started section, five samples are listed in the samples directory. The Sample Host Plugin 2 (Sample II) is the plug-in sample that covers information about Software Library artifacts.

Once you have downloaded the EDK kit to your local system, select `oracle.samples.xsh2` sample available in `/samples/plugins/` directory. If you further drill down to `/samples/plugins/oracle.samples.xsh2/plugin_dist/oms/metadata` directory, the `swlib` folder is displayed which contains the necessary examples.

Note: All the Software Library Metadata described in this section must be included in the XML file as described in [Section 14.3](#). Following which, you must register them using the information available in [Section 14.4](#).

14.2.1 Defining Folders

The following example describes how you can create two top level folders called `MPFolder1` and `MPFolder2`, with a subfolder named `Subfolder` under each of these folders:

```
<Folders>
  <Folder name="MPFolder1">
  </Folder>
  <Folder name="MPFolder2">
  </Folder>
  <Folder name="MPFolder1/subfolder">
  </Folder>
  <Folder name="MPFolder2/subfolder">
  </Folder>
</Folders>
```

After defining and registering the folders, when the plug-in is deployed, these folders will be displayed in Software Library console with a lock symbol, which means that the folders are Oracle-Owned, and can not be edited. You can customize them using the Create Like functionality available in Software Library. All the other folders that are created using Enterprise Manager console by users appear without the lock symbol, and can be edited by anyone who has been given the required accesses on the folder.

14.2.2 Defining Types

The following example describes how to create two Type artifacts called MPType1 and MPType2:

```
<Types>
<EntityType internalName="MPType1"/>
<EntityType internalName="MPType2"/>
</Types>
```

14.2.3 Defining Subtypes

The following example describes how you can create a subtype MPSubtype1 for the type MPType1:

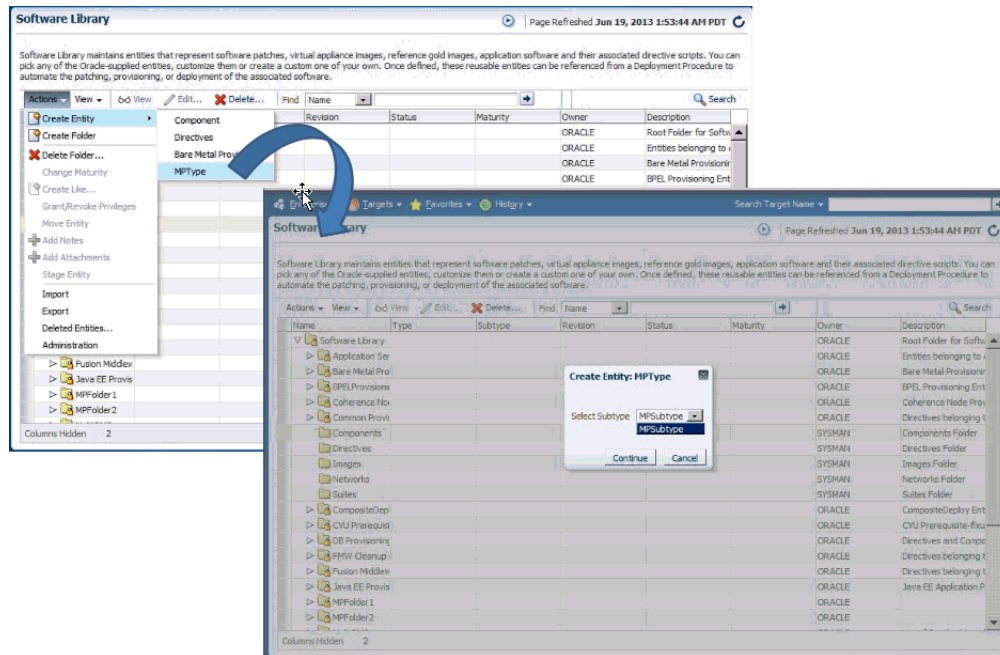
```
<EntitySubtype internalName="MPSubtype1" type="MPType1">
  <EntityProperties filename="MPSubtype-entPropDict.xml"/>
  <GenericUISpecification>
    <Create>
      <ContentDescriptor>
        <DisplayName default="Describe"/>
        <Description default="Describe"/>
        <Content
contentId="/WEB-INF/sdk/core/regions/swlib/sdkcore-regions-swlib-describe-task-flow.xml#sdkcore-regions-swlib-describe-task-flow" contentType="ADFRegion"/>
        </ContentDescriptor>
        <ContentDescriptor>
          <DisplayName default="Select Files"/>
          <Description default="Select Files"/>
          <Content
contentId="/WEB-INF/sdk/core/regions/swlib/sdkcore-regions-swlib-upload-task-flow.xml#sdkcore-regions-swlib-upload-task-flow" contentType="ADFRegion"/>
          </ContentDescriptor>
          <ContentDescriptor>
            <DisplayName default="Review"/>
            <Description default="Review"/>
            <Content
contentId="/WEB-INF/sdk/core/regions/swlib/sdkcore-regions-swlib-review-task-flow.xml#sdkcore-regions-swlib-review-task-flow" contentType="ADFRegion"/>
            </ContentDescriptor>
          </Create>
          <Edit>
            <ContentDescriptor>
              <DisplayName default="Describe"/>
              <Description default="Describe"/>
              <Content
contentId="/WEB-INF/sdk/core/regions/swlib/sdkcore-regions-swlib-describe-task-flow.xml#sdkcore-regions-swlib-describe-task-flow" contentType="ADFRegion"/>
              </ContentDescriptor>
              <ContentDescriptor>
                <DisplayName default="Select Files"/>
                <Description default="Select Files"/>
```

```

        <Content
contentId="/WEB-INF/sdk/core/regions/swlib/sdkcore-regions-swlib-upload-task-flow.xml#sdkcore-regions-swlib-upload-task-flow" contentType="ADFRegion"/>
        </ContentDescriptor>
    </Edit>
    <View>
        <ContentDescriptor>
            <DisplayName default="Describe"/>
            <Description default="Describe"/>
            <Content
contentId="/WEB-INF/sdk/core/regions/swlib/sdkcore-regions-swlib-describe-task-flow.xml#sdkcore-regions-swlib-describe-task-flow" contentType="ADFRegion"/>
        </ContentDescriptor>
        <ContentDescriptor>
            <DisplayName default="Select Files"/>
            <Description default="Select Files"/>
            <Content
contentId="/WEB-INF/sdk/core/regions/swlib/sdkcore-regions-swlib-upload-task-flow.xml#sdkcore-regions-swlib-upload-task-flow" contentType="ADFRegion"/>
        </ContentDescriptor>
    </View>
</GenericUISpecification>
</EntitySubtype>

```

Once a subtype defined is registered, entities of this subtype can be created (if creation is not blocked) using the Software Library console. To do so, from Enterprise menu, select **Provisioning and Patching**, then click **Software Library**. From the Software Library page, select a user-owned folder. From **Actions** menu, select **Create**, and then select **Type** and then the **SubType**. All the sub-types that belong to the type selected are displayed. For more information, see the following graphic file.



Ideally, you must see the subtypes you created as a part of this list, if not, there is an issue with registration. For more information on registering the sub-type, see [Section 14.4](#).

14.2.3.1 Entity Properties File

Following is the entity properties file `MPSubtype-entPropDict.xml`:

```
<?xml version="1.0"?>
<Dictionary xmlns="http://www.oracle.com/sysman/emgc/Properties">
  <TypeDefinitions>
    <DictionaryItem refID="MPAttr1" purpose="PURPOSE_TopLevel">
      <PropType typeCode="TYPE_String">
        <SimpleType className="java.lang.String" uiHintReadOnly="false"
uiHintHidden="false" secret="false" guid="false">
          <Constraints/>
        </SimpleType>
      </PropType>
    </DictionaryItem>

    <DictionaryItem refID="MPAttr2" purpose="PURPOSE_TopLevel">
      <PropType typeCode="TYPE_String">
        <SimpleType className="java.lang.String" uiHintReadOnly="false"
uiHintHidden="false" secret="false" guid="false">
          <Constraints/>
        </SimpleType>
      </PropType>
    </DictionaryItem>

    <DictionaryItem refID="MPAttr3" purpose="PURPOSE_TopLevel">
      <PropType typeCode="TYPE_String">
        <SimpleType className="java.lang.String" uiHintReadOnly="false"
uiHintHidden="false" secret="false" guid="false">
          <Constraints/>
        </SimpleType>
      </PropType>
    </DictionaryItem>
  </TypeDefinitions>
  <DynamicTypes/>
</Dictionary>
```

To use the file, ensure that you make the following changes:

- `MPAttr1`, `MPAttr2`, and `MPAttr3` are the names of the entity properties, remove the properties that are not required, and change the names of the properties that you retain appropriately.

Copy the `DictionaryItem` section from the above sample to create more entity properties. Note that, you can only update the name, you cannot change other aspects of the entity property.

- The entity properties for this sub-type are displayed as `Other Attributes` in Software Library console. Only, after the subtype metadata with entity property definitions are registered, while creating or editing entities, you can specify values for the entity properties. However, note that you can not define new entity properties.

Software Library

Describe Select Files Review

Create MPSubtype : Describe [Back] Step 1 of 3 [Next] [Save] [Cancel]

Parent Directory testfolder
Subtype MPSubtype

Enter name, description and other attributes that describe the entity. These attributes are shared by all revisions of this entity. Additionally, attach any documents and keep notes.

* Name MPTypeSubtypeEntity
Description

Other Attributes

Name	Value	Description
MPAttr3	value1	
MPAttr2	value2	
MPAttr1	value3	

Attachments

+ Add ✕ Remove

File Name	Size(KB)	Mime Type
No attachment has been added yet.		

Notes

New Note [] + Add

Note	Added By	Date
No note has been added yet.		

☒ **TIP** Notes once added cannot be deleted or edited.

14.2.4 Defining Entities

The following example describes how you can create an entity of the type `MPTType` and subtype `MPSubtype` in the folder `MPFolder1`:

```
<Entity name="MPTSEntity">
  <Type>MPTType</Type>
  <Directory>/MPFolder1</Directory>
  <Subtype>MPSubtype</Subtype>
  <Fileset>
    <FileEntry>
      <path>payload.zip</path>
      <sourcePath>payloadSrc.zip</sourcePath>
    </FileEntry>
  </Fileset>
  <ExternalID>0.1</ExternalID>
</Entity>
```

FileSet: A file `payloadSrc.zip` should be present in the same directory as of this XML. The entity will have a file entry `payload.zip`.

ExternalID: A decimal floating point number of the form `xxxxxxxx.x`, which means, a maximum 8 digits before decimal point, and one digit after decimal point is allowed. For example, 0.6, 23.9, and so on. When registering metadata again, the external IDs of the entities are compared. Entities whose ExternalID matches the ExternalID of the latest revision currently registered will not get registered again.

For example:

- If the entity `MPTSEntity` in `MPFolder1` is registered with ExternalID 0.1, the an entity with revision 0.1 is created. If ExternalID is not specified, then the default value is 0.1.

- If XML is updated and registered again without changing ExternalID, then this entity will not be registered, and a warning will be logged. Entity will continue to have 0.1 revision.
- If the ExternalID is changed to 0.2, and registered again, then the registration will be applied, and a revision 0.2 is created in addition to the earlier registered 0.1 revision. Consequently, multiple increments might need to be performed to ExternalID for it to be registered during the development and testing phase. However, before the release, ensure that you reset it back to the correct value and track it correctly.

14.3 Organizing Software Library Metadata Files

To organize the Software Library Metadata files, follow these steps:

1. Once you have the XMLs to define types, subtypes, folders, and entities ready, navigate to the following location:

```
$OMS_PLUGIN/metadata/swlib
```

2. Create a directory for your functional area (for example, `functionalArea`), in the following location:

```
$OMS_PLUGIN/metadata/swlib/functionalArea
```

3. Create a driving file (`swlib.xml`)

4. Edit `order.xml` file available at the following location: `$OMS_PLUGIN/metadata`, to add an entry for the newly created `functionalArea/swlib.xml` file:

Here is a sample example of how your directory structure should look:

```
swlib/
  order.xml
  functionalArea/swlib.xml
  functionalArea1/swlib.xml
```

Here is a sample example of the contents of `order.xml` file:

```
<order>
  <name>functionalArea/swlib.xml</name>
  <name>functionalArea1/swlib.xml</name>
</order>
```

Note: For each functional area in the plug-in, you must add an entry in the `order.xml` file as described in the example.

14.4 Adding the Software Library Metadata to Enterprise Manager

Adding the Software Library Metadata to Enterprise Manager is a two-step process as described in this section:

- [Step 1: Validating Metadata XML](#)
- [Step 2: Adding Metadata XML to OPAR](#)

14.4.1 Step 1: Validating Metadata XML

For the purpose of testing, use `emctl` to register the metadata as follows:

```
emctl register oms metadata -service swlib -file <Metadata Instance file>
-pluginId <Plugin Id> [-sysman_pwd "sysman password"]
```

Where:
Metadata Instance file is the path to the folder containing order.xml. For example, \$OMS_PLUGIN/metadata/swlib.

Note: If you have not added a Software Library location, then the `emctl register` command will not work, instead, you will see an error message as follows:

```
EM-04040: Metadata operation is skipped. Reason: Software Library
OMS shared storage is not configured, skipping metadata
registration. Check /u01/inst/em/EMGC_OMS1/sysman/log/emctl.log for
more details.
```

Also, if the metadata XML file is repeatedly registered for entities, then you must ensure that the external ID element is incremented each time. For example, if there are 11 entities defined in XML, four of which have same ExternalID, and seven have updated ExternalID, then you will see the following message:

```
Total 0 errors, 4 warnings. 7 entities imported.
Metadata registration successful
```

Note: OMS must be restarted for the registration to take effect, when a metadata registered using `emctl` contains types or subtypes. However, OMS need not be restarted for the registration of types and subtypes, when the Software Library metadata defined in a plug-in contains types or subtype, and is deployed.

Oracle recommends that you check the logs even if the registration is successful, as there may be some warnings. These warnings are mainly caused when the external ID is not modified, inturn causing entity registration to fail. To view the logs, navigate to the following location:

```
$INSTANCE_HOME/sysman/log/emctl.log
```

14.4.2 Step 2: Adding Metadata XML to OPAR

When a plug-in OPAR is deployed containing Software Library metadata organized in the way described above, Software Library metadata will be registered if Software Library is configured on the system. Software Library metadata XMLs are to be included in the OPAR like any other metadata. See Chapter 13, "Validating, Packaging, and Deploying the Plug-in" for more information.

Note 1: If Software Library is not configured at plug-in deployment time, the plug-in's Software Library metadata will get registered whenever Software Library is configured for the first time after the plugin deployment.

14.5 Using Software Library Entities

Software Library entities created by the plug-in may represent a patch/script/configuration or any other software relevant to the plug-in. To use these entities after they have been created using the options described in previous sections, consider one of the following approaches:

- [Using Job Types](#)
- [Using EMCLI Verbs](#)

Note: For information about how plug-in UI uses the Software Library search service, see [Section 8.11.9](#).

14.5.1 Using Job Types

Software Library makes use of the following job types:

- `SwlibStageEntities`
- `SwlibUploadFiles`

You can create your own job types, which inturn contain these jobtypes. For example, you can create a jobtype XML that contains the `SwlibStageEntities` jobtype as follows:

```
<?xml version="1.0"?>
<jobType name="StageWrap" version="1.0" singleTarget="true" targetTypes="host"
editable="true">
  <credentials>
    <credential usage="destHostCreds" authTargetType="host"
defaultCredentialSet="HostCredsNormal">
    </credential>
    <credential usage="destNfsHostCreds" authTargetType="host"
defaultCredentialSet="HostCredsPriv">
    </credential>
  </credentials>
  <paramInfo>
    <paramSource sourceType="user" paramNames="stageLocation, entityURN"
      required="true" evaluateAtSubmission="true" />
    <paramSource sourceType="user" paramNames="stageFileEntryPaths, operMode,
autoRetry" required="false"/>
    <paramSource sourceType="inline" paramNames="operMode">
      <sourceParam name="paramValues" value="mount"/>
      <sourceParam name="overwriteExistingFiles" value="yes"/>
    </paramSource>
  </paramInfo>
  <stepset ID="main" type="serial">
    <step ID="preStage" command="remoteOp">
      <credList>
        <cred usage="defaultHostCred" reference="destHostCreds"/>
      </credList>
      <paramList>
        <param name="remoteCommand">%job_default_shell%</param>
        <param name="args">ls, -R, %stageLocation%</param>
        <param name="targetName">%job_target_names%[1]</param>
        <param name="targetType">%job_target_types%[1]</param>
      </paramList>
    </step>
    <job ID="stage" type="SwlibStageEntities">
      <credList>
        <cred usage="destHostCreds" reference="destHostCreds"/>
        <cred usage="destNfsHostCreds" reference="destNfsHostCreds"/>
      </credList>
      <paramList>
        <param name="entityURN">%entityURN%</param>
        <param name="stageFileEntryPaths" valueOf="stageFileEntryPaths"/>
      </paramList>
    </job>
  </stepset>
</jobType>
```

```

        <param name="stageLocation">%stageLocation%</param>
        <param name="operMode">%operMode%</param>
        <param name="autoRetry">%autoRetry%</param>
        <param name="overwriteExistingFiles">%overwriteExistingFiles%</param>
    </paramList>
    <targetList allTargets="true" />
</job>
<step ID="postStage" command="remoteOp">
    <credList>
        <cred usage="defaultHostCred" reference="destHostCreds" />
    </credList>
    <paramList>
        <param name="remoteCommand">%job_default_shell%</param>
        <param name="args">ls, -R, %stageLocation%</param>
        <param name="targetName">%job_target_names%[1]</param>
        <param name="targetType">%job_target_types%[1]</param>
    </paramList>
</step>
</stepset>
<displayInfo
    useDefaultCreateUI="true"
    showParams="true">

    <jobTypeDisplayInfo>
        <nlsValue>Stage Wrap</nlsValue>
    </jobTypeDisplayInfo>
    <parameterDisplayInfo name="stageLocation" showInResults="true"
showInCreate="true">
        <parameterLabel>
            <nlsValue>Stage Location</nlsValue>
        </parameterLabel>
        <parameterHint>
            <nlsValue>Directory location on target where the files from the entity
will be transferred.</nlsValue>
        </parameterHint>
        <parameterTextBox lines="1" />
    </parameterDisplayInfo>
    <parameterDisplayInfo name="entityURN" showInResults="true"
showInCreate="true">
        <parameterLabel>
            <nlsValue>Entity URN</nlsValue>
        </parameterLabel>
        <parameterHint>
            <nlsValue>Internal identifier of the entity</nlsValue>
        </parameterHint>
        <parameterTextBox lines="1" />
    </parameterDisplayInfo>
</displayInfo>
</jobType>

```

Jobs of jobtype SwlibStageEntities and SwlibUploadFiles expect the following inputs:

- Stage Location: The directory path where the file of the entity will be transferred.
- Entity URN: This is the internal identifier of the entity, which can be obtained any one of the following methods:
 - In Enterprise Manager Cloud Control, from the **Enterprise** menu, select **Provisioning and Patching**, then click **Software Library**. On the Software Library home page, from **View** menu, select **Internal ID** to enable it. Copy

and supply the Internal ID value available to the job as the value of this parameter.

- You can use the emcli verb `list_swlib_entities` with the parameter `-show_entity_rev_id` to obtain the Internal ID. Copy and supply the Internal ID value available to the job as the value of this parameter.

14.5.2 Using EMCLI Verbs

You can use the EMCLI verbs provided by Software Library to add Software Library storage location, create folders, create entities, upload files for entities, modify entities, and so on. For more information about EMCLI verbs, see *Oracle Enterprise Manager Command Line Interface* and *Oracle Enterprise Manager Lifecycle Management Administrator's Guide*.

Defining Credentials

As part of the target type definition, you can define the types of credentials specific to the plug-in target type. Examples could be the username and password required by the plug-in to connect to a target instance to collect metric data, or to invoke a specific Enterprise Manager job.

The Enterprise Manager credential subsystem enables Enterprise Manager administrators to store credentials, in a secure manner, as preferences or operation credentials. The credentials can then be used to perform different system management activities, such as real-time monitoring, patching, provisioning, and other target administrative operations.

In this release, the credential subsystem supports the storing, accessing, and modifying of fixed number user name/password based credentials as preferred credentials, which other Enterprise Manager subsystems access to build automation solutions. The credential subsystem also supports sudo/powerbroker based impersonation support.

This chapter covers the following:

- [Introduction to Security Concepts](#)
- [Defining Credential Metadata](#)

15.1 Introduction to Security Concepts

The following sections describe the concepts associated with credential service integration:

- **Credential Types**

Credential type is the type of authentication supported by a target type. Various authentication schemes are supported, including native agent authentication and SSH. For more information, see [Section 15.1.1, "Understanding Credential Types"](#).

- **Named Credentials**

A named credential is a user's authentication information on a system and can be a user name/password, a public key-private key pair, or an X509v3 certificate. For more information, see [Section 15.1.2, "About Named Credentials"](#).

- **Authentication Target Type**

An authenticating target type is the target type that a credential can authenticate against. For more information, see [Section 15.1.3, "Authenticating Target Types"](#).

- **Credential Sets**

The credential set is a placeholder for a credential and can be used to decouple credentials from the system that uses a credential. For more information, see [Section 15.1.4, "Overview of Credential Sets"](#)

- **Credential Store**

The credential store is a logical store for all the named credentials of an Enterprise Manager administrator in the Enterprise Manager. For more information, see [Section 15.1.5, "Using the Credential Store"](#)

- **Credential Reference**

The credential reference is a way to refer to a credential. For more information, see [Section 15.1.6, "About the Credential Reference"](#)

15.1.1 Understanding Credential Types

Credential type is the type of authentication supported by a target type. For example, a host can support a user name/password based authentication, public key authentication, or kerberos authentication. Various authentication schemes are supported, including native agent authentication and SSH.

The native agent authentication scheme employs a user name/password structure, while the SSH Key authentication scheme user a user name/private key/public key structure.

15.1.2 About Named Credentials

A named credential is a users' authentication information on a system. A named credential can be a user name/password, a public key-private key pair, or an X509v3 certificate. An Enterprise Manager administrator can store these credentials as named entities in Enterprise Manager to use when performing operations like running jobs, patching, and other system management tasks. For example, you can store the user name and password that you want to use for patching as `MyPatchingCreds`. You can then later submit a patching job that uses `MyPatchingCreds` to patch the production databases.

Named Credentials can be created for the credential types in Enterprise Manager 12c. The most commonly used credential types for host and database target types are described in the following sections.

For more information on named credentials, see the *Configuring Security* chapter in the *Oracle Enterprise Manager Cloud Control Administrator's Guide*. This can be found at the following location:

http://docs.oracle.com/cd/E24628_01/doc.121/e24473/security.htm

Host Target Type

- **Host Credentials**

Users can create named credentials by providing the username and password for the host. Privilege delegation properties such as run privilege, runas, and profile can also be provided.

- **SSH Key Credentials**

Named credentials of type SSH Key credential can be created by providing the host username, SSH public key, and SSH private key. Privilege delegation properties such as run privilege, runas, and profile can be also be provided.

Database Target Type

- Database Credentials

Named credentials of this type can be created by providing the database username, password, and role.

- Database Kerberos Credentials

Named credentials of this type can be created by providing the Kerberos user name and Kerberos password. Database Kerberos credentials can not be used in this release for automation purposes. These can be used only for user interface operations, such as logging in to the database and viewing pages.

15.1.3 Authenticating Target Types

Authenticating target type is the target type that a credential can authenticate against. For example, a SQLScript job has a host credential DBHostCreds that is used to authenticate against the database host. Therefore, the target type for DBHostCreds is Database Instance and the authenticating target type is Host.

15.1.4 Overview of Credential Sets

The credential set is a placeholder for a credential. Credential sets can be used to decouple credentials from the system that uses a credential. For example, a patching job can be submitted to use the credential set "Normal Host Credentials" while being executed.

The "Normal Host Credentials" credential set can also be set to the actual named credential. The credential set to named credential mapping for the target can be changed without editing the system that uses the credential.

15.1.5 Using the Credential Store

The credential store is a logical store for all the named credentials of an Enterprise Manager administrator in the Enterprise Manager. The Enterprise Manager administrator's user name has a logical private credential store. Individual credentials can be identified by credential names. Enterprise Manager administrators can add, edit, and delete named credentials in the credential store.

15.1.6 About the Credential Reference

The credential reference is a way to refer to a credential. There are three ways credentials can be referenced:

- Credential Name

The credential is referenced using the name of the credential in the credential store.

- Credential Set

The credential is referenced using the credential set name and the target name. The lookup gets the credential associated with the credential set name and target name.

- Direct

The credential is specified by providing the values of the attributes. This reference does not refer to a credential in the credential store.

15.2 Defining Credential Metadata

Credential metadata is defined within the target type metadata file. See [Section 3.3, "Creating the Target Type Metadata File"](#) for details on this file.

All credential metadata for a target type is defined within a `CredentialInfo` element. This element in turn contains the following subelements:

- A `CredentialType` element that defines the type of credentials to be used to access target instances
- A `CredentialSet` element that instantiates an instance of `CredentialType`

The following shows a basic example defining the credentials required to authenticate with hosts running instances of the target: username and password.

Example 15–1 Credential Metadata

```
<TargetMetadata>

...
<CredentialInfo>
  <!-- The types of credentials: target host username/password -->
  <CredentialType NAME="HostCreds">
    <Display>
      <Label NLSID="CREDS_HOST_HOSTCREDS">Host Credentials</Label>
    </Display>
    <CredentialTypeColumn NAME="HostUserName" IS_KEY="TRUE">
      <Display>
        <Label NLSID="CREDS_HOST_USERNAME">UserName</Label>
      </Display>
    </CredentialTypeColumn>
    <CredentialTypeColumn NAME="HostPassword">
      <Display>
        <Label NLSID="CREDS_HOST_Password">Password</Label>
      </Display>
    </CredentialTypeColumn>
  </CredentialType>
  <!-- The CredentialSet that creates an instance of CredentialType -->
  <CredentialSet NAME="HostCredsNormal" CREDENTIAL_TYPE="HostCreds"
    USAGE="PREFERRED_CRED">
    <Display>
      <Label NLSID="CREDS_HOST_HOSTCREDS_NORMAL">Normal Host Credentials</Label>
    </Display>
    <CredentialSetColumn TYPE_COLUMN="HostUserName" SET_COLUMN="username">
      <Display>
        <Label NLSID="CREDS_NORMAL_USER">Normal Username</Label>
      </Display>
    </CredentialSetColumn>
    <CredentialSetColumn TYPE_COLUMN="HostPassword" SET_COLUMN="password">
      <Display>
        <Label NLSID="CREDS_NORMAL_PASSWORD">Normal Password</Label>
      </Display>
    </CredentialSetColumn>
  </CredentialSet>
</CredentialInfo>
...
</TargetMetadata>
```

15.2.1 Overview of Credential Elements

The key elements that define credentials are described in the following table:

Table 15–1 Key elements in a *plugin.xml* file

Element	Required (Y/N)	Description
CredentialInfo	Y	The root element for the credentials definition. Contains CredentialType and CredentialSet elements.
CredentialType	Y	Contains one or more CredentialTypeColumn elements, each defining a credential - such as "TargetUsername" or "TargetPassword" - to be used to access target instances.
CredentialSet	Y	<p>Instantiates an instance of the credential set defined in CredentialType. It includes the following attributes:</p> <ul style="list-style-type: none"> ■ CREDENTIAL_TYPE Identifies the CredentialType this CredentialSet is created from. ■ USAGE Values are MONITORING (default), which are used to directly connect to the target; PREFERRED_CRED, which are the user's preferred credentials; or SYSTEM, which are used by specialized applications like patching or cloning.
CredentialSetColumn	Y	<p>Subelement of CredentialType. Defines a single credential and maps that credential to its corresponding column in the CredentialType. It includes the following attributes:</p> <ul style="list-style-type: none"> ■ TYPE_COLUMN Specifies the CredentialTypeColumn that this CredentialSetColumn maps to. ■ SET_COLUMN Identifies the column definition in the CredentialSet

Converting an Existing Metadata Plug-in

Enterprise Manager Cloud Control includes a new metadata plug-in framework that is significantly enhanced from the framework used in previous releases. Existing metadata plug-ins, that is, plug-ins released for any version earlier than Cloud Control 12c are not compatible with the new framework.

This chapter provides instructions on converting existing plug-ins to the format required by the new framework.

Note: This conversion requirement applies to existing plug-ins only. If you are building a new plug-in for Enterprise Manager Cloud Control 12c, the issues discussed in this chapter will not affect you.

This chapter contains the following sections:

- [Introduction to Converting an Existing Metadata Plug-in](#)
- [Impact of the New Plug-in Framework on Existing Plug-ins](#)
- [Using the `convert_mp` Utility to Convert Plug-in Metadata](#)
- [Post-Conversion Steps](#)

16.1 Introduction to Converting an Existing Metadata Plug-in

As a plug-in developer, you are responsible for the following steps within the conversion process:

1. Familiarize yourself with the changes to the new plug-in framework.

For more information, see [Section 16.2, "Impact of the New Plug-in Framework on Existing Plug-ins"](#).

2. Convert the plug-in metadata to the new format by running the `convert_mp` utility.

For more information, see [Section 16.3, "Using the `convert_mp` Utility to Convert Plug-in Metadata"](#).

3. Manually modify the generated plug-in configuration files.

For more information, see [Section 16.4, "Post-Conversion Steps"](#).

4. Use the Extensibility Development Kit (EDK) to package the plug-in as a deployable archive.

For more information, see [Chapter 13, "Validating, Packaging, and Deploying the Plug-in"](#).

5. Validate, deploy, and test your plug-in against the Enterprise Manager Cloud Control 12 installation.

For more information, see [Chapter 13, "Validating, Packaging, and Deploying the Plug-in"](#).

16.2 Impact of the New Plug-in Framework on Existing Plug-ins

Plug-in developers should be aware of the following changes to the new plug-in framework and impact to the Enterprise Manager Cloud Control 12c upgrade process.

16.2.1 Plug-in Metadata Must Be Converted to the New Format

The plug-in metadata format used by Enterprise Manager Cloud Control 12c is incompatible with the format used in prior releases. As a result, the metadata for all existing plug-ins must be converted to the new format per the instructions in this document.

The Extensibility Development Kit (EDK) includes a `convert_mp` utility that converts existing plug-in metadata files to the new format, and generates additional metadata files required by the new framework.

16.2.2 Plug-ins Must Be Packaged as Oracle Plug-in Archive (OPAR) Files

Plug-ins must be packaged (or re-packaged) as Oracle Plug-in Archive (OPAR) files for deployment into Cloud Control 12c. The Metadata Plug-in Archive (MPA) files used to package previous plug-ins cannot be deployed into Cloud Control 12c.

16.2.3 In-place Upgrade of Existing Plug-ins Not Supported

An “in-place” upgrade of existing plug-ins is not supported as part of the Enterprise Manager Cloud Control 12c upgrade process. That is, there is no automated conversion or migration of plug-in artifacts already deployed to the existing installation to the upgraded installation.

Before beginning the upgrade process, all plug-ins included in the installation being upgraded must be converted to the new plug-in format as noted above.

16.2.4 Plug-ins Should Be Registered Before Upgrading

Once converted to the new format, plug-ins should be registered with the Cloud Control upgrade console prior to initiating the Enterprise Manager Cloud Control 12c upgrade process.

If the appropriate plug-in is not available and is not registered with the upgrade console, targets associated with the target type(s) specified in the plug-in metadata cannot be migrated from the existing installation to the upgraded installation.

16.3 Using the `convert_mp` Utility to Convert Plug-in Metadata

Existing plug-in metadata can be transformed to the new format using the `convert_mp` command-line utility, which is included with the Plug-in Developer Kit. The tool also generates two new metadata files required by the new format.

The tool can be run in two modes:

- Against a 10.2.x or 11.1 Management Repository containing the metadata for deployed plug-ins. See [16.3.1](#) , "[Converting Plug-In Metadata Stored In A Management Repository](#)" for instructions.
- On an existing metadata plug-in archive (MPA) file. See [16.3.2](#) , "[Converting a Metadata Plug-in Archive \(MPA\)](#)" for instructions.

The tool outputs new metadata files to the directory structure required to create a Plug-in Archive (OPAR) file, which can then be deployed into Enterprise Manager Cloud Control 12c. Note that the tool does not create the archive; this must be done as a separate step.

The tool converts the following existing metadata to the new format:

- Target type definitions
- Performance metrics and metric collection definitions
- Collection scripts and binaries
- Job type definitions
- Job scripts and binaries
- Report definitions
- Homepage chart customizations
- Related links to the homepage

The utility also creates the following new files, which contain metadata required by the new framework:

- plugin.xml

This file is used during plug-in deployment. It contains properties that identify the plug-in, such as name and version, and declares the set of target types that will be added to Enterprise Manager

- plugin_registry.xml

This file declares those components included in the plug-in that are to be deployed to the Management Agent.

Note: The plugin.xml and plugin_registry.xml files must be manually updated after the conversion process is complete.

For more information, see [Section 16.4.1](#), "[Modifying the plugin.xml and plugin_registry.xml Files](#)".

16.3.1 Converting Plug-In Metadata Stored In A Management Repository

The `convert_mp` command can be run against an Enterprise Manager release 10.2.x or 11.1 Management Repository containing metadata for deployed plug-ins. Note that only one plug-in can be converted on each invocation of the command.

Note: If the plug-in includes report definition metadata, it must be converted using this process. Report definitions cannot be converted by running the conversion utility against a metadata plug-in archive (MPA) file.

If there are no report definitions, then you can use either mode to convert plug-in metadata.

Before running the utility, you must import the plug-in to an existing Enterprise Manager Enterprise Manager 10.2.x or 11.1 installation and deploy the plug-in to a Management Agent. This causes the plug-in metadata to be written to the Management Repository. It is not necessary to create any target instances before running the tool.

The tool connects to the Management Repository using the specified connection string as the specified Cloud Control 10.2.x or 11.1 user. It then writes the converted files to a subdirectory named /opar within the specified output directory.

To use the convert_mp utility, navigate to the /bin directory within the EDK.

The basic usage syntax for the convert_mp utility is as follows:

```
convert_mp -conn_desc repository_connection_string -repos_user username
[-repos_password repos_password] [-mp_name plugin_name] [-mp_version plugin_
version]
-out_dir output_directory [-plugin_id converted_plug_id] [-force]
```

The following is a basic conversion example:

```
C:\edk\bin>convert_mp -conn_desc myhost.us.example.com:25055:$ORACLE_SID -repos_
user sysman -mp_name sample_plugin_01 -out_dir /tmp/plugins
```

Table 16–1 provides the options that can be used to convert metadata stored in a Management Repository.

Table 16–1 Options to Convert Metadata Stored in a Management Repository

Option	Required Y/N	Description
-conn_desc	Y	The connection string used to connect to the Management Repository. Specify the string in the form of host:port:sid.
-repos_user	Y	The user to connect to the 10.2.x or 11.1 Management Repository. Oracle recommends that you connect as SYSMAN. Note: You must not supply a Cloud Control 12.0 user name. The tool prompts you for the password if not supplied with -repos_password.
-repos_password	N	The 10.2.x or 11.1 Management Repository password. You will be prompted for the password if not specified. Note: Providing a password on the command line is insecure and should be avoided in a production environment.

Table 16–1 (Cont.) Options to Convert Metadata Stored in a Management Repository

Option	Required Y/N	Description
-mp_name	Y	Required if multiple plug-ins exist in the Management Repository. Because only one plug-in can be converted at a time, the name of the plug-in to convert must be supplied.
-mp_version	Y	Required if multiple plug-ins with the same name exist in the Management Repository.
-out_dir	Y	The directory to output the converted XML files to. A new subdirectory named /opar will be created in this directory
-plugin_id	N	An identifier that will be written to the plugin.xml and plugin_registry.xml files generated as part of the conversion process. If not included on the command line, the identifier must be specified in the plugin.xml and plugin-registry.xml files before packaging the plug-in.
-force	N	If set, an existing output directory for the plug-in will be deleted, and a new directory will be created. If not set, an error will occur if an output directory already exists.

16.3.2 Converting a Metadata Plug-in Archive (MPA)

The `convert_mp` utility can also convert the metadata files for one or more plug-ins packaged in a metadata plug-in archive (MPA). You can specify a single plug-in to convert, or convert all plug-ins within the archive simultaneously.

Note: If the plug-in includes report definition metadata, it must be converted using the Management Repository option. Report definitions cannot be converted by running the conversion utility against a metadata plug-in archive (MPA) file.

For more information about the Management Repository option, see [Section 16.3.1, "Converting Plug-In Metadata Stored In A Management Repository"](#).

Note that unlike the Management Repository option, the plug-in does not need to be deployed to Enterprise Manager Cloud Control before conversion.

The syntax is as follows:

```
convert_mp -mpa mpa_file_location [-mp_name plugin_name] [-mp_version plugin_version] -out_dir output_directory [-plugin_id converted_plugin_id] {-force}
```

The following is a basic conversion example. The tool will convert all plug-ins within `my_gc11_plugins.mpa` and output the converted files for each to a unique /opar directory within `/tmp/plugins/`.

```
convert_mp -mpa /tmp/my_gc11_plugin.mpa -out_dir /tmp/plugins
```

[Table 16–2](#) describes the options that can be used to convert an MPA file.

Table 16–2 Options to Convert an MPA File

Option	Required Y/N	Description
-mpa	Y	The path to the MPA file, relative to the EDK installation.
-mp_name	N	The name of a plug-in within the archive to convert. If not supplied, all plug-ins in the archive will be converted.
-mp_version	N	If multiple plug-ins have the same name, the plug-in version to convert. If not supplied, all versions of the specified mp_name will be converted.
-out_dir	Y	The directory to output the converted XML files to. A new subdirectory named /opar will be created in this directory. If multiple plug-ins are converted, a unique /opar directory will be created for each.
-plugin_id	N	An identifier that will be written to the plugin.xml and plugin_registry.xml generated as part of the conversion process. If not included on the command line, the identifier must be specified in the plugin.xml and plugin-registry.xml files before packaging the plug-in.
-force	N	If set, an existing output directory for the plug-in will be deleted, and a new directory will be created. If not set, an error will occur if an output directory already exists.

16.4 Post-Conversion Steps

After the transformation process is complete, several additional manual changes that should be considered before packaging the plug-in.

16.4.1 Modifying the plugin.xml and plugin_registry.xml Files

The plugin.xml file is generated during the conversion process. It is used during the plug-in deployment process.

The file contains properties that identify the plug-in, such as name and version, and declares the set of target types that will be added to Enterprise Manager using other metadata included in the plug-in.

If the -plugin_id option is not included on the command line when running the convert_mp command, a <PluginID> element must be specified in the plugin.xml and plugin_registry.xml files before packaging the plug-in.

For more information about the contents of the plugin.xml file and plugin_registry.xml file, see [Section 2.4, "Creating the plugin.xml File"](#) and [Section 2.5, "Creating the plugin_registry.xml File"](#).

Example 16–1 Modifying the plugin.xml File

```
<?xml version="1.0" encoding="UTF-8" ?>
<Plugin xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
        xsi:schemaLocation="http://www.oracle.com/EnterpriseGridControl/plugin_
metadata plugin_metadata.xsd"
        xmlns="http://www.oracle.com/EnterpriseGridControl/plugin_metadata">
<!-- If the -plugin_id option is not included on the command line when running
the convert_mp command, the PluginID element must be manually specified. -->
  <PluginId vendorId="oracle" productId="sysman" pluginTag="db2"/>
  <PluginVersion value="11.2.0.1.0"/>
  <PluginOMSOSAruId value="2000">
</PluginOMSOSAruId>
```

```

<TargetTypeList>
  <TargetType name="ibm_db2_database" isIncluded="TRUE">
  </TargetType>
</TargetTypeList>

<EMPlatforms>
  <CertifiedPlatform version="11.2.0.1.0"/>
</EMPlatforms>
<AgentSideCompatibility>
  <Version>11.2.0.1.0</Version>
</AgentSideCompatibility>
<PluginJ2EEArtifacts/>
</Plugin>

```

16.4.2 Converting Report Definitions Created With PL/SQL

Report definitions created using PL/SQL that include SQL statement references through the “oracle.sysman.eml.ip.render.elem.sqlStatement” type cannot be converted directly to the new report metadata format supported in Enterprise Manager Cloud Control 12c.

If the report definition includes SQL statements that retrieve the report data, the SQL must be extracted and registered as “namedSQL”.

An existing report definition might contain a SQL statement such as:

```

l_param_values(N) := MGMT_IP_PARAM_VALUE_
RECORD('oracle.sysman.eml.ip.render.elem.sqlStatement', 'select target_name,
target_type from mgmt$target');

```

It might also include the following:

```

l_element_guid := mgmt_ip.add_element_to_report_def (
  p_report_guid      => l_report_guid,
  p_element_name_nlsid => 'IPMSG_TABLE_FROM_SQL',
  p_element_type_nlsid => 'IPMSG_ANY_TARGET_TYPE',
  p_header_nlsid     => null,
  p_element_order    => 2,
  p_element_row      => 2,
  p_parameters       => l_param_values,
  p_targets          => null
);

```

Before you begin to convert your report to XML format, check [Section 16.4.2.1, "Named SQL Statements"](#).

16.4.2.1 Named SQL Statements

If you have a Named SQL statement registered for your report, you must convert it to XML.

For example, if your PL/SQL report definition includes a line similar to the following:

```

mgmt_ip.register_sql_statement (
  p_name => 'oracle.sysman.db.storage.reports.TablespaceUsage',
  p_sql_statement => 'select * from dual');

```

Convert to an XML file as follows:

```

<?xml version = '1.0' encoding = 'UTF-8'?>
<NamedSQLStatements xmlns= "http://www.oracle.com/DataCenter/NamedSQL">

```

```
<NamedSQL sqlName="oracle.sysman.db.storage.reports.TablespaceUsage"
sqlValue="select * from dual"/>
</NamedSQLStatements>
```

16.4.2.2 Converting Report Definitions to XML

The steps for converting report definitions are as follows.

1. Run the report definition PL/SQL to register the report against your Enterprise Manager 10.2.x or 11.1 installation.
2. Log in to Enterprise Manager as SYSMAN and from the Reports page, use **Create Like** to make a copy of your out-of-the-box report.

Note: You must change the name of your report temporarily using the **Edit Report Definition** page.

3. Export the report definition to an XML file using the `emcli export_report` command as shown below. Note that you must specify the report ID as the value for `-title`:

```
emcli export_report -title='MY_REPORT_ID' -output_file='./tmp/report.xml'
```

MY_REPORT_ID is the name of the copied file that you created in step 2.

4. Include the generated file in the `/oms/metadata/report` directory within the plug-in staging area. The file will be imported when the plug-in is deployed
5. Create an XML file that includes the named SQL statements referenced in the report, as shown below. Multiple statements can be included in the file. Ensure that the same statement identifier is used:

```
<?xml version = '1.0' encoding = 'UTF-8'?>
  <NamedSQLStatements xmlns="http://www.oracle.com/DataCenter/NamedSQL">
    <NamedSQL sqlName="MY_UNIQUE_SQL_STATEMENT_ID"
      sqlValue="'select target_name, target_type from mgmt$target"/>
  </NamedSQLStatements>
```

Note: Oracle supports one XML file per report definition. Named SQL statements can be grouped into one file at your discretion.

These steps must be repeated for all report definitions that include SQL statements.

6. Include the file in the `/oms/metadata/namedsql` directory within the plug-in staging area. The file will be imported when the plug-in is deployed.

16.4.3 Converting Job Type Definitions

Job type definitions created in an existing plug-in will be converted by the `convert_mp` utility. However, the XSD that describes the job type metadata is more strictly validated in Cloud Control 12c. As a result, validation errors may occur during plug-in deployment if any existing job type definitions include XML that is not valid according to the new XSD.

The best way to proactively find such issues is to run the `empdk validate_plugin` command against the plug-in. Each time you encounter an error, fix it, then run the command again. Continue running the command until it no longer reports any validation errors.

See [Section 13.3, "Validating the Plug-in"](#) for details on using the `empdk validate_plugin` command.

Refer to the job type XSD to determine how to restructure incorrectly formatted job type definitions.

16.4.4 Packaging the Plug-in

Once you have completely converted your plug-in to the new format, you will be ready to package the plug-in for deployment into Enterprise Manager Cloud Control 12c. See [Section 13.4, "Creating the Plug-in Archive"](#) for details.

Monitoring Using Web Services and JMX

You can extend Enterprise Manager to monitor Web services and JMX-instrumented applications for critical events, performance problems, error conditions, and statistics.

Enterprise Manager's ability to monitor WSDL and JMX-enabled targets enables you to consolidate monitoring and management operations. When added to the Enterprise Manager framework, Enterprise Manager functionality, such as notifications, jobs, and reporting, is automatically extended to these targets.

This chapter contains the following topics:

- [Overview](#)
- [Monitoring Using Web Services in Enterprise Manager](#)
- [Monitoring Using WS-Management in Enterprise Manager](#)
- [Monitoring JMX Applications Deployed on Oracle Application Servers \(OC4J\)](#)
- [Monitoring a Standalone JMX-instrumented Java Application or JVM Target](#)
- [Monitoring JMX Applications Deployed on Oracle WebLogic Application Servers](#)
- [Adding a Target to a Management Agent](#)
- [Monitoring Credential Setup](#)
- [Viewing Monitored Metrics](#)
- [Creating JMX Metric Extensions](#)
- [Surfacing Metrics from a Standalone JVM or Oracle Coherence](#)

17.1 Overview

Using Enterprise Manager to monitor targets that expose a Web services management interface, JMX-instrumented applications and servers, and standalone Java Virtual Machine (JVM) targets entails defining a new target type via metadata plug-ins.

Creating a metadata plug-in consists of four basic steps:

1. Generate the target metadata and default collection files to be added to the plug-in.
2. Create an Oracle Plug-in Archive containing the target definition files for one or more plug-ins. A single archive may contain more than one plug-in.
3. Import the plug-in into Enterprise Manager.
4. Deploy the plug-in to the appropriate Management Agents.

For more information about each of these steps, see [Chapter 13, "Validating, Packaging, and Deploying the Plug-in"](#).

Procedural information for the monitoring targets can be found in the following sections:

- [Section 17.2](#) discusses software components exposing an external interface that communicate across a network using a standard messaging protocol.
- [Section 17.4](#) discusses J2EE applications running on an OC4J that are instrumented using JMX MBeans.
- [Section 17.5](#) discusses standalone Java applications running on J2SE5.0 or higher that are instrumented using JMX MBeans.
- [Section 17.6](#) discusses JMX applications running on Oracle WebLogic Application Servers 9.x or above.

[Section 17.4](#) and [Section 17.5](#) explain how to generate metadata and default collection files for your custom JMX-enabled application by guiding you through the MBeans for which you are interested in collecting data, and helping you define the MBeans as metrics in Enterprise Manager. Even if your standalone Java application is not instrumented through JMX, you can still monitor the JVMs it is running on by directly creating the built-in JVM target instances as defined in [Section 17.7.3](#).

After the metadata and default collection files are created, you can follow the normal metadata plug-in mechanism to deploy your plug-in and create target instances of your Java application target type.

17.2 Monitoring Using Web Services in Enterprise Manager

Web services are loosely coupled software components that expose an external interface via the Web Service Definition Language (WSDL). These components communicate across a network using a standard messaging protocol called Simple Object Access Protocol (SOAP). The Management Agent's Web service Fetchlet (with ID WSF) supports SOAP communication.

Note: For more information about the Web services standard, see the World Wide Web Consortium (W3C) website:

<http://www.w3.org>

Prerequisites

- Management Agent version 12.1.0.0.0 or later installed on that host.
- Oracle Management Server (OMS) version 12.1.0.0.0 or later with which the Management Agent communicates.

17.2.1 Creating Metadata and Default Collection Files

Defining a target type to be monitored through a Web services interface includes creating the requisite target definition files, which are required to collect metrics from resources that support the WSDL interface:

- Target Metadata
- Default Collection

Enterprise Manager provides an easy-to-use Web services command-line tool that simplifies creating plug-ins by automatically generating these requisite files. Information retrieval is achieved through the Web services fetchlet that is integrated with the Management Agent.

The command-line tool works by parsing a specified WSDL file for all operations, and enables you to select one or more operations to be invoked. If multiple port types are specified in the WSDL file, the tool prompts you to select one of them. Operations are listed along with their parameters. A Web service operation can be one of four types:

- One Way
- Request Response
- Solicit Response
- Notification

The Request Response operation type is particularly useful: The selected operation could have primitive or complex parameters and results. The result of Web service invocation is displayed in a table (the tool prompts you to provide labels for the table columns). You can also filter result attributes by specifying an Xpath expression (see the `RowType` property in the generated target metadata, [Example 17-3](#)). Filter attributes can be useful for complex return types from which only few attributes are interesting.

The Web services command-line tool supports Web services with the following binding and encoding styles:

- DOC/literal
- DOC/Wrapped
- RPC/encoded

17.2.1.1 Web Services CLI Command-line Tool Syntax

The Web services CLI command-line tool syntax is as follows:

```
emctl wscli [-metadata | -help] [-options]
```

The command accepts the following options:

- `-wsdl=file | URL`: WSDL file or URL (mandatory)
- `-username=user ID`: user name if the WSDL is protected

The command-line tool requires a WSDL file name or URL to locate the WSDL for a Web service. For example, for a Calculator service Web service, a WSDL URL would be as follows:

```
http://localhost:44861/CalWS/CalculatorPort?WSDL
```

The command tool script requires access to the Enterprise Manager home directory (`EM_HOME`) to run. The tool defaults to `ORACLE_HOME` (ensure this environment variable is set properly before using this tool).

The tool parses specified WSDL for all the port types and binding (supported protocols such as HTTP get/post, SOAP) to list all the operations. If there are multiple port types in WSDL, you will first be prompted to choose a port type.

17.2.1.2 Web Services Command-line Tool Security

The command-line tool generates metadata required by Enterprise Manager for target monitoring purposes through the WSDL file. When you run this tool, you only require read permission on the WSDL file or URL and permission to save generated files to the appropriate directory.

17.2.1.3 Generating the Files

[Example 17-1](#) shows a sample WSDL file passed to the command-line tool to generate the target metadata and collection files.

Example 17-1 Sample WSDL File CalculatorService.wsdl

```
<?xml version="1.0" encoding="UTF-8"?>
<!-- Published by JAX-WS RI at http://jax-ws.dev.java.net. RI's version is Oracle
JAX-WS 2.1.5. -->
<wsdl:definitions xmlns:wsdl="http://schemas.xmlsoap.org/wsdl/"
xmlns:tns="http://tests.jaxws.oracle.com/"
xmlns:ns0="http://www.oracle.com/jaxws/tests"
xmlns:xsd="http://www.w3.org/2001/XMLSchema"
xmlns:mime="http://schemas.xmlsoap.org/wsdl/mime/"
xmlns:soap="http://schemas.xmlsoap.org/wsdl/soap/"
xmlns:soap12="http://schemas.xmlsoap.org/wsdl/soap12/" name="CalculatorService"
targetNamespace="http://tests.jaxws.oracle.com/">
  <wsdl:types>
    <xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema" version="1.0"
targetNamespace="http://www.oracle.com/jaxws/tests/types">
      <xs:complexType name="calculatorFaultInfo">
        <xs:sequence>
          <xs:element name="number" type="xs:int"/>
          <xs:element name="reason" type="xs:string" minOccurs="0"/>
        </xs:sequence>
      </xs:complexType>
    </xs:schema>
    <xs:schema xmlns:ns1="http://www.oracle.com/jaxws/tests/types"
xmlns:tns="http://www.oracle.com/jaxws/tests"
xmlns:xs="http://www.w3.org/2001/XMLSchema" version="1.0"
targetNamespace="http://www.oracle.com/jaxws/tests">
      <xs:import namespace="http://www.oracle.com/jaxws/tests/types"/>
      <xs:element name="CalculatorException" nillable="true"
type="tns:CalculatorException"/>
      <xs:element name="CalculatorWrapperException" nillable="true"
type="ns1:calculatorFaultInfo"/>
      <xs:complexType name="CalculatorException">
        <xs:sequence>
          <xs:element name="Message" type="xs:string"/>
          <xs:element name="Number" type="xs:int"/>
          <xs:element name="Reason" type="xs:string"/>
        </xs:sequence>
      </xs:complexType>
    </xs:schema>
    <schema xmlns="http://www.w3.org/2001/XMLSchema"
xmlns:xsd="http://www.w3.org/2001/XMLSchema"
xmlns:tns="http://tests.jaxws.oracle.com/"
targetNamespace="http://tests.jaxws.oracle.com/">
      <xsd:complexType name="add">
        <xsd:sequence>
          <xsd:element name="arg0" type="xsd:int"/>
          <xsd:element name="arg1" type="xsd:int"/>
        </xsd:sequence>
      </xsd:complexType>
    </schema>
  </wsdl:types>
</wsdl:definitions>
```

```

        </xsd:complexType>
        <xsd:element name="add" type="tns:add"/>
        <xsd:complexType name="addResponse">
            <xsd:sequence>
                <xsd:element name="return" type="xsd:int"/>
            </xsd:sequence>
        </xsd:complexType>
        <xsd:element name="addResponse" type="tns:addResponse"/>
        <xsd:complexType name="square">
            <xsd:sequence>
                <xsd:element name="arg0" type="xsd:int"/>
            </xsd:sequence>
        </xsd:complexType>
        <xsd:element name="square" type="tns:square"/>
        <xsd:complexType name="squareResponse">
            <xsd:sequence>
                <xsd:element name="arg0" type="xsd:int"/>
            </xsd:sequence>
        </xsd:complexType>
        <xsd:element name="squareResponse" type="tns:squareResponse"/>
        <xsd:complexType name="checkNumber">
            <xsd:sequence>
                <xsd:element name="arg0" type="xsd:int"/>
            </xsd:sequence>
        </xsd:complexType>
        <xsd:element name="checkNumber" type="tns:checkNumber"/>
        <xsd:complexType name="checkNumberResponse">
            <xsd:sequence>
                <xsd:element name="return" type="xsd:boolean"/>
            </xsd:sequence>
        </xsd:complexType>
        <xsd:element name="checkNumberResponse"
type="tns:checkNumberResponse"/>
    </schema>
</wsdl:types>
<wsdl:message name="addInput">
    <wsdl:part name="parameters" element="tns:add"/>
</wsdl:message>
<wsdl:message name="addOutput">
    <wsdl:part name="parameters" element="tns:addResponse"/>
</wsdl:message>
<wsdl:message name="squareInput">
    <wsdl:part name="parameters" element="tns:square"/>
</wsdl:message>
<wsdl:message name="squareOutput">
    <wsdl:part name="parameters" element="tns:squareResponse"/>
</wsdl:message>
<wsdl:message name="checkNumberInput">
    <wsdl:part name="parameters" element="tns:checkNumber"/>
</wsdl:message>
<wsdl:message name="checkNumberOutput">
    <wsdl:part name="parameters" element="tns:checkNumberResponse"/>
</wsdl:message>
<wsdl:message name="CalculatorWrapperException">
    <wsdl:part name="CalculatorWrapperException"
element="ns0:CalculatorWrapperException"/>
</wsdl:message>
<wsdl:message name="CalculatorException">
    <wsdl:part name="CalculatorException" element="ns0:CalculatorException"/>
</wsdl:message>

```

```
<wsdl:portType name="Calculator">
  <wsdl:operation name="add">
    <wsdl:input xmlns:ns1="http://www.w3.org/2006/05/addressing/wsdl"
message="tns:addInput" ns1:Action=""/>
    <wsdl:output xmlns:ns1="http://www.w3.org/2006/05/addressing/wsdl"
message="tns:addOutput" ns1:Action=""/>
  </wsdl:operation>
  <wsdl:operation name="square">
    <wsdl:input xmlns:ns1="http://www.w3.org/2006/05/addressing/wsdl"
message="tns:squareInput" ns1:Action=""/>
    <wsdl:output xmlns:ns1="http://www.w3.org/2006/05/addressing/wsdl"
message="tns:squareOutput" ns1:Action=""/>
  </wsdl:operation>
  <wsdl:operation name="checkNumber">
    <wsdl:input xmlns:ns1="http://www.w3.org/2006/05/addressing/wsdl"
message="tns:checkNumberInput" ns1:Action=""/>
    <wsdl:output xmlns:ns1="http://www.w3.org/2006/05/addressing/wsdl"
message="tns:checkNumberOutput" ns1:Action=""/>
    <wsdl:fault name="CalculatorWrapperException"
message="tns:CalculatorWrapperException"/>
    <wsdl:fault name="CalculatorException"
message="tns:CalculatorException"/>
  </wsdl:operation>
</wsdl:portType>
<wsdl:binding name="CalculatorSoapHttp" type="tns:Calculator">
  <soap:binding style="document"
transport="http://schemas.xmlsoap.org/soap/http"/>
  <wsdl:operation name="add">
    <soap:operation soapAction=""/>
    <wsdl:input>
      <soap:body use="literal"/>
    </wsdl:input>
    <wsdl:output>
      <soap:body use="literal"/>
    </wsdl:output>
  </wsdl:operation>
  <wsdl:operation name="square">
    <soap:operation soapAction=""/>
    <wsdl:input>
      <soap:body use="literal"/>
    </wsdl:input>
    <wsdl:output>
      <soap:body use="literal"/>
    </wsdl:output>
  </wsdl:operation>
  <wsdl:operation name="checkNumber">
    <soap:operation soapAction=""/>
    <wsdl:input>
      <soap:body use="literal"/>
    </wsdl:input>
    <wsdl:output>
      <soap:body use="literal"/>
    </wsdl:output>
    <wsdl:fault name="CalculatorWrapperException">
      <soap:fault name="CalculatorWrapperException" use="literal"
encodingStyle=""/>
    </wsdl:fault>
    <wsdl:fault name="CalculatorException">
      <soap:fault name="CalculatorException" use="literal"
encodingStyle=""/>
    </wsdl:fault>
  </wsdl:operation>
</wsdl:binding>
</wsdl:service>
```

```

        </wsdl:fault>
    </wsdl:operation>
</wsdl:binding>
<wsdl:service name="CalculatorService">
    <wsdl:port name="CalculatorPort" binding="tns:CalculatorSoapHttp">
        <soap:address
location="http://localhost:8888/CalWSBA/CalculatorPort"/>
    </wsdl:port>
</wsdl:service>
</wsdl:definitions>

```

[Example 17–2](#) uses the WSDL file shown in [Example 17–1](#), "Sample WSDL File CalculatorService.wsdl". First, the tool parses the WSDL for all port types and bindings (supported protocols such as HTTP get/post or SOAP) to list all the operations. If there are multiple port types in the WSDL, the tool first prompts you to select a port type.

To start the command-line tool:

1. Go to the \$AGENT_HOME/bin directory.
2. Run the following command:

```
$ emctl wscli -metadata -wsdl=/tmp/CalculatorWS.wsdl
```

Once invoked, the command-line tool automatically prompts you for the requisite information, as shown in [Example 17–2](#), "Sample Web Services Command-Line Tool Session". If you need to quit a command-line tool session, you can press Ctrl+C at any point to exit. Session information will not be saved.

Example 17–2 Sample Web Services Command-Line Tool Session

```
Oracle Enterprise Manager 12c Release 1 Cloud Control 12.1.0.1.0
Copyright (c) 1996, 2011 Oracle Corporation. All rights reserved.
```

```
OracleHome : /oracle/oms/agent
EMDROOT    : /oracle/oms/agent
```

```
Generate Metric Metadata for Web Service Monitoring
```

```
Reading WSDL Document at /tmp/CalculatorWS.wsdl...done.
```

```
==> Enter the metadata file name [/tmp/target/metadata/CalculatorService.xml] :
```

```
* Selected Service: CalculatorService
```

```
* Selected Port: CalculatorPort
```

```
All operations for the selected Port "CalculatorPort":
```

```
[1]  squareResponse square(int arg0)
[2]  checkNumberResponse checkNumber(int arg0)
[3]  addResponse add(int arg0, int arg1)
```

```
==> Enter the index [1-3] of operation to select: 1
```

```
* Selected Operation:
    squareResponse square(int arg0)
```

```
Define new metric group:
```

```
==> Enter the name for this metric group [square]:
```

```
Return value(s) for the selected operation:
[1] //ns0:squareResponse/arg0 <int>

==> Enter the index [1-1] of metric to display: 1
==> Enter the name for this metric [arg0]: SquareResult
==> Enter the label for this metric [SquareResult]:
==> Is this a key metric <y/n>? [n] :
==> Do you want to create threshold for this item <y/n>? [n] :

Setup operation Argument: square.arg0 <type:int>
==> Enter value [%square.arg00001%] :

==> Do you want to use jps-config-jse.xml <y/n>? [n] :

==> Do you want to add User/Password Credential <y/n>? [n] : y
==> Enter the name for User/Password credential set [UserCredentialSet01] :

==> Do you want to add SSL TrustStore Credential <y/n>? [n] :

==> Do you want to add SSL KeyStore Credential <y/n>? [n] :

==> Do you want to add KeyStore Credential <y/n>? [n] :

==> Do you want to add Encryption Key Credential <y/n>? [n] :

==> Do you want to add Signature Key Credential <y/n>? [n] :

==> Is this metric group for periodic collection <y/n>? [y] :
The following units are for collection frequency:
[1] Min
[2] Hr
[3] Day

==> Enter the index [1-3] of unit for this collection: 1
==> Enter the frequency of collection in Min: 30

==> Do you want to add another metric group <y/n>? [n] :

Files Generated:
- Target Metadata file: /tmp/target/metadata/CalculatorService.xml
- Target Collection file: /tmp/target/metadata/CalculatorServiceCollection.xml
```

The command-line tool generates the metadata required to monitor the `CalculatorService` target type as shown in [Example 17-3](#).

Example 17-3 CalculatorService Target Metadata File

```
<!DOCTYPE TargetMetadata SYSTEM "../dtds/TargetMetadata.dtd">
<TargetMetadata META_VER="1.0" TYPE="CalculatorService">
  <Display>
    <Label NLSID="NLSID_CALCULATOR_SERVICE">CalculatorService</Label>
    <ShortName NLSID="NLSID_CALCULATOR_SERVICE">CalculatorService</ShortName>
    <Description NLSID="NLSID_CALCULATOR_SERVICE">CalculatorService</Description>
  </Display>
  <Metric NAME="square" TYPE="TABLE">
    <Display>
      <Label NLSID="NLSID_SQUARE">square</Label>
    </Display>
  </Metric>
</TargetMetadata>
```



```

<TableDescriptor>
  <ColumnDescriptor IS_KEY="FALSE" NAME="SquareResult" TYPE="STRING">
    <Display>
      <Label NLSID="COL_SQUARE_RESULT">SquareResult</Label>
    </Display>
  </ColumnDescriptor>
</TableDescriptor>
<QueryDescriptor FETCHLET_ID="WSF">
  <Property NAME="ProxyHost" SCOPE="INSTANCE"
OPTIONAL="TRUE">ProxyHost</Property>
  <Property NAME="ProxyPort" SCOPE="INSTANCE"
OPTIONAL="TRUE">ProxyPort</Property>
  <Property NAME="SecurityPolicy" SCOPE="INSTANCE"
OPTIONAL="FALSE">square.SecurityPolicy</Property>
  <Property NAME="ServiceEndpoint" SCOPE="INSTANCE"
OPTIONAL="FALSE">square.ServiceEndpoint</Property>
  <Property NAME="ServiceName" SCOPE="GLOBAL"
OPTIONAL="FALSE">ns0:CalculatorService</Property>
  <Property NAME="PortName" SCOPE="GLOBAL"
OPTIONAL="FALSE">ns0:CalculatorPort</Property>
  <Property NAME="OperationName" SCOPE="GLOBAL"
OPTIONAL="FALSE">square</Property>
  <Property NAME="MessageType" SCOPE="GLOBAL" OPTIONAL="FALSE">SOAP</Property>
  <Property NAME="SOAPBindingStyle" SCOPE="GLOBAL"
OPTIONAL="FALSE">DOCUMENT</Property>
  <Property NAME="SOAPBindingUse" SCOPE="GLOBAL"
OPTIONAL="FALSE">LITERAL</Property>
  <Property NAME="ParameterStyle" SCOPE="GLOBAL"
OPTIONAL="FALSE">WRAPPED</Property>
  <Property NAME="SOAPVersion" SCOPE="GLOBAL" OPTIONAL="FALSE">SOAP_1_
1</Property>
  <Property NAME="Namespace" SCOPE="GLOBAL"
OPTIONAL="FALSE"><![CDATA[ [ns0="http://tests.jaxws.oracle.com/"] ]]></Property>
  <Property NAME="RowType" SCOPE="GLOBAL"
OPTIONAL="FALSE"> /ns0:squareResponse/arg0</Property>
  <Property NAME="ColType" SCOPE="GLOBAL"
OPTIONAL="FALSE">SquareResult:STRING</Property>
  <Property NAME="Payload" SCOPE="GLOBAL"
OPTIONAL="FALSE"><![CDATA[<soap:Envelope
xmlns:soap="http://schemas.xmlsoap.org/soap/envelope/">
  <soap:Body xmlns:ns1="http://tests.jaxws.oracle.com/">
    <ns1:square>
      <arg0>%square.arg00001%</arg0>
    </ns1:square>
  </soap:Body>
</soap:Envelope>]]></Property>
  <Property NAME="UserCredential" SCOPE="GLOBAL"
OPTIONAL="FALSE">UserCredentialSet01</Property>
  <CredentialRef
NAME="UserCredentialSet01">UserCredentialSet01</CredentialRef>
</QueryDescriptor>
</Metric>
<CredentialInfo>
  <CredentialType NAME="CSFKeyCredential">
    <Display>
      <Label NLSID="CRED_TYPE">CSF-Key Credential Type</Label>
    </Display>
  <CredentialTypeColumn NAME="CSFKey">
    <Display>
      <Label NLSID="CRED_C_S_F_KEY">Alias CSF Key</Label>
    </Display>
  </CredentialTypeColumn>
</CredentialInfo>

```

```
</Display>
</CredentialTypeColumn>
</CredentialType>
<CredentialType NAME="AliasCredential">
  <Display>
    <Label NLSID="CRED_TYPE">Alias Credential Type</Label>
  </Display>
  <CredentialTypeColumn NAME="Alias">
    <Display>
      <Label NLSID="CRED_ALIAS">Alias (i.e. username, encryption key,
signature key, etc)</Label>
    </Display>
  </CredentialTypeColumn>
  <CredentialTypeColumn NAME="Password">
    <Display>
      <Label NLSID="CRED_PASSWORD">Password for the alias</Label>
    </Display>
  </CredentialTypeColumn>
</CredentialType>
<CredentialSet NAME="UserCredentialSet01" USAGE="MONITORING">
  <AllowedCredType TYPE="CSFKeyCredential"/>
  <AllowedCredType TYPE="AliasCredential"/>
</CredentialSet>
</CredentialInfo>
<InstanceProperties>
  <InstanceProperty NAME="ProxyHost" CREDENTIAL="FALSE" OPTIONAL="TRUE">
    <Display>
      <Label NLSID="PROP_PROXY_HOST">Proxy Server Name</Label>
    </Display>
  </InstanceProperty>
  <InstanceProperty NAME="ProxyPort" CREDENTIAL="FALSE" OPTIONAL="TRUE">
    <Display>
      <Label NLSID="PROP_PROXY_PORT">Proxy Server Port</Label>
    </Display>
  </InstanceProperty>
  <InstanceProperty NAME="square.SecurityPolicy" CREDENTIAL="FALSE"
OPTIONAL="FALSE">
    <Display>
      <Label NLSID="PROP_SQUARE_SECURITY_POLICY">[square] Authentication/Web
Service Policy</Label>
    </Display>
  </InstanceProperty>
  <InstanceProperty NAME="square.ServiceEndpoint" CREDENTIAL="FALSE"
OPTIONAL="FALSE">
    <Display>
      <Label NLSID="PROP_SQUARE_SERVICE_ENDPOINT">[square] Web Service Endpoint
URL</Label>
    </Display>
  </InstanceProperty>
  <InstanceProperty NAME="square.arg00001" CREDENTIAL="FALSE" OPTIONAL="FALSE">
    <Display>
      <Label NLSID="PROP_SQUARE_ARG00001">[square] square.arg0</Label>
    </Display>
  </InstanceProperty>
</InstanceProperties>
</TargetMetadata>
```

The command-line tool also generates the requisite collection file as shown in [Example 17-4](#).

Example 17–4 CalculatorService Default Collection File

```
<!DOCTYPE TargetCollection SYSTEM "../dtds/TargetCollection.dtd">
<TargetCollection TYPE="CalculatorService">
  <CollectionItem NAME="square">
    <Schedule>
      <IntervalSchedule TIME_UNIT="Min" INTERVAL="30"/>
    </Schedule>
  </CollectionItem>
</TargetCollection>
```

After the tool generates the target metadata and collection files, you can create the Oracle Plug-in archive. For more information, see [Section 13.4, "Creating the Plug-in Archive"](#).

17.3 Monitoring Using WS-Management in Enterprise Manager

Beginning with Enterprise Manager 12c, WS-Management (WS-MAN)-compliant resources can be monitored using the fetchlet WSMangementFetchlet.

The fetchlet communicates with the WS-MAN resources using WS-Transfer protocol, which defines a number of management operations that the managed resources should support. However, in the current release, the fetchlet only supports the operation WS-Transfer GET.

Note: For more information about the monitor WS-Management standard, see the DMTF Web Services Management website:

<http://www.dmtf.org/standards/wsman>

Prerequisites

- Management Agent version 12.1.0.0.0 or greater installed on that host.
- Oracle Management Server (OMS) version 12.1.0.0.0 or greater with which the Management Agent communicates.

17.3.1 Creating Metadata and Default Collection Files

Enterprise Manager provides an easy-to-use WS-Management CLI command-line tool that simplifies creating new Management Plug-ins by automatically generating the requisite target metadata and default collection files. Information retrieval is achieved via the WSMangementFetchlet that is integrated with the Management Agent.

Resources, which support WS-Management interface, should describe their model-specific elements using XML Schema Definition (XSD) representation and expose the XSD as a public accessible link just like WSDL for Web services.

The command-line tool works by parsing a specified XSD file for the managed WS-MAN resource and then prompts you to select the interested resource properties to construct a monitoring metric.

17.3.1.1 WS-Management CLI Command-line Tool Syntax

The WS-Management CLI command-line tool syntax is as follows:

```
Usage: emctl wsmancli [-metadata | -help] [-options]
```

The command accepts the following options:

- `-schema=file | URL`: Resource XSD file or URL [mandatory]
- `-username=user ID`: User name if the schema is protected

The command-line tool requires a XSD file name or URL to locate the resource schema. For example, for a Traffic Light WS-Management service, a XSD URL would be as follows:

```
http://localhost:8888/TrafficLight?xsd
```

The command tool script requires access to the Enterprise Manager home directory (EM_HOME) to run. The tool defaults to ORACLE_HOME (ensure this environment variable is set properly before using this tool).

17.3.1.2 Command-line Tool Security

The command-line tool generates metadata required by Enterprise Manager for target monitoring purposes via the resource XSD. When you run this tool, you only need read permission on the XSD file or URL and permission to save generated files to the appropriate directory.

17.3.1.3 Generating Target Metadata and Collection Files

[Example 17-5](#) shows a sample XSD file passed to the command-line tool to generate the target metadata and collection files.

Example 17-5 Sample XSD File TrafficLight.xsd

```
<?xml version="1.0" encoding="UTF-8"?>
<xs:schema
targetNamespace="http://schemas.wiseman.dev.java.net/traffic/1/light.xsd"
elementFormDefault="qualified" blockDefault="#all"
xmlns:tl="http://schemas.wiseman.dev.java.net/traffic/1/light.xsd"
xmlns:xs="http://www.w3.org/2001/XMLSchema">
  <xs:complexType name="TrafficLightType">
    <xs:sequence>
      <xs:element name="name" type="xs:string"/>
      <xs:element name="color" type="xs:string"/>
      <xs:element name="x" type="xs:int"/>
      <xs:element name="y" type="xs:int"/>
    </xs:sequence>
  </xs:complexType>
  <xs:element name="trafficlight" type="tl:TrafficLightType"/>
</xs:schema>
```

To start the command-line tool:

1. Go to the \$AGENT_HOME/bin directory.
2. Execute the following command:

```
$ emctl wsmancli -metadata -schema= http://localhost:8080/Traffic?xsd
```

Once invoked, the command-line tool automatically prompts you for the requisite information, as shown in [Example 17-6, "Sample WS-Management CLI Command-Line Tool Session"](#). If you need to quit a command-line tool session, you can press Control+C at any point to exit. Session information will not be saved.

Example 17-6 Sample WS-Management CLI Command-Line Tool Session

```
Oracle Enterprise Manager 12c Release 1 Cloud Control 12.1.0.0.0
Copyright (c) 1996, 2011 Oracle Corporation. All rights reserved.
```

```

OracleHome : /oracle/oms/agent
EMDROOT    : /oracle/oms/agent

Generate Metric Metadata for WS-Management Resource Monitoring

Reading Resource XSD Document at http://localhost:8080/Traffic?xsd...done.

==> Enter the name for this target type: TrafficLight

==> Enter the metadata file name [/tmp/target/metadata/TrafficLight.xml] :

Define new metric group name:
==> Enter the name for this metric group: trafficLight

WS-Addressing namespaces:
[1] http://www.w3.org/2005/08/addressing
[2] http://schemas.xmlsoap.org/ws/2004/08/addressing

==> Enter the index [1-2] to select: 1

SOAP Envelope namespaces:
[1] http://www.w3.org/2003/05/soap-envelope
[2] http://schemas.xmlsoap.org/soap/envelope/

==> Enter the index [1-2] to select: 1

Resource properties:
[1] trafficlight:color
[2] trafficlight:name
[3] trafficlight:x
[4] trafficlight:y

==> Enter the index [1-4] of property to display: 2
==> Enter the name for this metric [name]:
==> Enter the label for this metric [name]:
==> Is this a key metric <y/n>? [n] : y
==> Do you want to add another metric <y/n>? [n] : y

Resource properties:
[1] trafficlight:color
[2] trafficlight:x
[3] trafficlight:y

==> Enter the index [1-3] of property to display: 1
==> Enter the name for this metric [color]:
==> Enter the label for this metric [color]:
==> Is this a key metric <y/n>? [n] :
==> Do you want to create threshold for this item <y/n>? [n] :
==> Do you want to add another metric <y/n>? [n] : y

Resource properties:
[1] trafficlight:x
[2] trafficlight:y

==> Enter the index [1-2] of property to display: 1
==> Enter the name for this metric [x]:
==> Enter the label for this metric [x]:
==> Is this a key metric <y/n>? [n] :

```

```

==> Do you want to create threshold for this item <y/n>? [n] :
==> Do you want to add another metric <y/n>? [n] : y

Resource properties:
[1]   trafficlight:y

==> Enter the index [1-1] of property to display: 1
==> Enter the name for this metric [y]:
==> Enter the label for this metric [y]:
==> Is this a key metric <y/n>? [n] :
==> Do you want to create threshold for this item <y/n>? [n] :

==> Enter comma-separated list of Selector elements: name

==> Do you want to add User/Password Credential <y/n>? [n] : y
==> Enter the name for User/Password credential set [UserCredentialSet01] :

==> Is this metric group for periodic collection <y/n>? [y] :
The following units are for collection frequency:
[1]   Min
[2]   Hr
[3]   Day

==> Enter the index [1-3] of unit for this collection: 1
==> Enter the frequency of collection in Min: 30

==> Do you want to add another metric group <y/n>? [n] :

Files Generated:
- Target Metadata file: /tmp/target/metadata/TrafficLight.xml
- Target Collection file: /tmp/target/metadata/TrafficLightCollection.xml

```

The command-line tool generates the metadata required to monitor the target type TrafficLight as shown in [Example 17-7](#).

Example 17-7 TrafficLight Target Metadata File

```

<!DOCTYPE TargetMetadata SYSTEM "../dtds/TargetMetadata.dtd">
<TargetMetadata META_VER="1.0" TYPE="TrafficLight">
  <Display>
    <Label NLSID="NLSID_TRAFFIC_LIGHT">TrafficLight</Label>
    <ShortName NLSID="NLSID_TRAFFIC_LIGHT">TrafficLight</ShortName>
    <Description NLSID="NLSID_TRAFFIC_LIGHT">TrafficLight</Description>
  </Display>
  <Metric NAME="trafficLight" TYPE="TABLE">
    <Display>
      <Label NLSID="NLSID_TRAFFIC_LIGHT">trafficLight</Label>
    </Display>
    <TableDescriptor>
      <ColumnDescriptor IS_KEY="TRUE" NAME="name" TYPE="STRING">
        <Display>
          <Label NLSID="COL_NAME">name</Label>
        </Display>
      </ColumnDescriptor>
      <ColumnDescriptor IS_KEY="FALSE" NAME="color" TYPE="STRING">
        <Display>
          <Label NLSID="COL_COLOR">color</Label>
        </Display>
      </ColumnDescriptor>
    </TableDescriptor>
  </Metric>
</TargetMetadata>

```

```

</ColumnDescriptor>
<ColumnDescriptor IS_KEY="FALSE" NAME="x" TYPE="STRING">
  <Display>
    <Label NLSID="COL_X">x</Label>
  </Display>
</ColumnDescriptor>
<ColumnDescriptor IS_KEY="FALSE" NAME="y" TYPE="STRING">
  <Display>
    <Label NLSID="COL_Y">y</Label>
  </Display>
</ColumnDescriptor>
</TableDescriptor>
<QueryDescriptor FETCHLET_ID="WSManagementFetchlet">
  <Property NAME="ProxyHost" SCOPE="INSTANCE"
OPTIONAL="TRUE">ProxyHost</Property>
  <Property NAME="ProxyPort" SCOPE="INSTANCE"
OPTIONAL="TRUE">ProxyPort</Property>
  <Property NAME="SecurityPolicy" SCOPE="INSTANCE"
OPTIONAL="TRUE">trafficLight.SecurityPolicy</Property>
  <Property NAME="ResourceURL" SCOPE="INSTANCE"
OPTIONAL="FALSE">trafficLight.ResourceURL</Property>
  <Property NAME="To" SCOPE="INSTANCE"
OPTIONAL="FALSE">trafficLight.To</Property>
  <Property NAME="OptionSet" SCOPE="INSTANCE"
OPTIONAL="TRUE">trafficLight.OptionSet</Property>
  <Property NAME="Locale" SCOPE="INSTANCE"
OPTIONAL="TRUE">trafficLight.Locale</Property>
  <Property NAME="MaxEnvelopeSize" SCOPE="INSTANCE"
OPTIONAL="TRUE">trafficLight.MaxEnvelopeSize</Property>
  <Property NAME="OperationTimeout" SCOPE="INSTANCE"
OPTIONAL="TRUE">trafficLight.OperationTimeout</Property>
  <Property NAME="Namespace" SCOPE="GLOBAL"
OPTIONAL="FALSE"><![CDATA[ [ns1="http://schemas.wiseman.dev.java.net/traffic/1/light
t.xsd"] [ns0="http://www.w3.org/2001/XMLSchema"] [wsa="http://www.w3.org/2005/08/add
ressing"] [env="http://www.w3.org/2003/05/soap-envelope"] ] ]></Property>
  <Property NAME="RowType" SCOPE="GLOBAL" OPTIONAL="FALSE">
//ns1:trafficlight/ns1:name,//ns1:trafficlight/ns1:color,//ns1:trafficlight/ns1:x,
//ns1:trafficlight/ns1:y</Property>
  <Property NAME="ColType" SCOPE="GLOBAL"
OPTIONAL="FALSE">name:STRING,color:STRING,x:STRING,y:STRING</Property>
  <Property NAME="ReplyTo" SCOPE="GLOBAL"
OPTIONAL="FALSE">http://www.w3.org/2005/08/addressing/role/anonymous</Property>
  <Property NAME="Action" SCOPE="GLOBAL"
OPTIONAL="FALSE">http://schemas.xmlsoap.org/ws/2004/09/transfer/Get</Property>
  <Property NAME="TransferOperation" SCOPE="GLOBAL"
OPTIONAL="FALSE">GET</Property>
  <Property NAME="SelectorSet" SCOPE="GLOBAL"
OPTIONAL="FALSE">[name,%trafficLight.name%]</Property>
  <Property NAME="UserCredential" SCOPE="GLOBAL"
OPTIONAL="FALSE">UserCredentialSet01</Property>
  <CredentialRef
NAME="UserCredentialSet01">UserCredentialSet01</CredentialRef>
</QueryDescriptor>
</Metric>
<CredentialInfo>
  <CredentialType NAME="CSFKeyCredential">
    <Display>
      <Label NLSID="CRED_TYPE">CSF-Key Credential Type</Label>
    </Display>
    <CredentialTypeColumn NAME="CSFKey">

```

```

        <Display>
            <Label NLSID="CRED_C_S_F_KEY">Alias CSF Key</Label>
        </Display>
    </CredentialTypeColumn>
</CredentialType>
<CredentialType NAME="AliasCredential">
    <Display>
        <Label NLSID="CRED_TYPE">Alias Credential Type</Label>
    </Display>
    <CredentialTypeColumn NAME="Alias">
        <Display>
            <Label NLSID="CRED_ALIAS">Alias (i.e. username, encryption key,
signature key, etc)</Label>
        </Display>
    </CredentialTypeColumn>
    <CredentialTypeColumn NAME="Password">
        <Display>
            <Label NLSID="CRED_PASSWORD">Password for the alias</Label>
        </Display>
    </CredentialTypeColumn>
</CredentialType>
<CredentialSet NAME="UserCredentialSet01" USAGE="MONITORING">
    <AllowedCredType TYPE="CSFKeyCredential"/>
    <AllowedCredType TYPE="AliasCredential"/>
</CredentialSet>
</CredentialInfo>
<InstanceProperties>
    <InstanceProperty NAME="ProxyHost" CREDENTIAL="FALSE" OPTIONAL="TRUE">
        <Display>
            <Label NLSID="PROP_PROXY_HOST">Proxy Server Name</Label>
        </Display>
    </InstanceProperty>
    <InstanceProperty NAME="ProxyPort" CREDENTIAL="FALSE" OPTIONAL="TRUE">
        <Display>
            <Label NLSID="PROP_PROXY_PORT">Proxy Server Port</Label>
        </Display>
    </InstanceProperty>
    <InstanceProperty NAME="trafficLight.SecurityPolicy"
        CREDENTIAL="FALSE" OPTIONAL="TRUE">
        <Display>
            <Label NLSID="PROP_TRAFFIC_LIGHT_SECURITY_POLICY">[trafficLight]
Authentication/Web Service Policy</Label>
        </Display>
    </InstanceProperty>
    <InstanceProperty NAME="trafficLight.ResourceURL" CREDENTIAL="FALSE"
OPTIONAL="FALSE">
        <Display>
            <Label NLSID="PROP_TRAFFIC_LIGHT_RESOURCE_U_R_L">[trafficLight] Resource
URL (wsman:ResourceURL)</Label>
        </Display>
    </InstanceProperty>
    <InstanceProperty NAME="trafficLight.To" CREDENTIAL="FALSE" OPTIONAL="FALSE">
        <Display>
            <Label NLSID="PROP_TRAFFIC_LIGHT_TO">[trafficLight] Network Address of the
service (wsa:To)</Label>
        </Display>
    </InstanceProperty>
    <InstanceProperty NAME="trafficLight.OptionSet" CREDENTIAL="FALSE"
OPTIONAL="TRUE">
        <Display>

```



```

        <Label NLSID="PROP_TRAFFIC_LIGHT_OPTION_SET">[trafficLight] Set of
wsman:Option. Format: [<OptionName1>; value:<value1>;,
type:<type1>;, mustComply:<true|false>;] [<OptionName2>;,
value:<value2>;, type:<type>;,
mustComply:<true|false>;][...]</Label>
    </Display>
</InstanceProperty>
<InstanceProperty NAME="trafficLight.Locale" CREDENTIAL="FALSE"
OPTIONAL="TRUE">
    <Display>
        <Label NLSID="PROP_TRAFFIC_LIGHT_LOCALE">[trafficLight] wsman:Locale (RFC
3066 language code). Format: e.g. en-US</Label>
    </Display>
</InstanceProperty>
<InstanceProperty NAME="trafficLight.MaxEnvelopeSize"
CREDENTIAL="FALSE" OPTIONAL="TRUE">
    <Display>
        <Label NLSID="PROP_TRAFFIC_LIGHT_MAX_ENVELOPE_SIZE">[trafficLight]
wsman:MaxEnvelopeSize in Octets. Format: e.g. 8192</Label>
    </Display>
</InstanceProperty>
<InstanceProperty NAME="trafficLight.OperationTimeout"
CREDENTIAL="FALSE" OPTIONAL="TRUE">
    <Display>
        <Label NLSID="PROP_TRAFFIC_LIGHT_OPERATION_TIMEOUT">[trafficLight]
wsman:OperationTimeout. Format: e.g. PT30S</Label>
    </Display>
</InstanceProperty>
<InstanceProperty NAME="trafficLight.name" CREDENTIAL="FALSE"
OPTIONAL="FALSE">
    <Display>
        <Label NLSID="PROP_TRAFFIC_LIGHT_NAME">[trafficLight] Value for the
Selector "name"</Label>
    </Display>
</InstanceProperty>
</InstanceProperties>
</TargetMetadata>

```

The command-line tool also generates the requisite collection file as shown in [Example 17–8, "TrafficLight Default Collection File"](#).

Example 17–8 TrafficLight Default Collection File

```

<!DOCTYPE TargetCollection SYSTEM "../dtds/TargetCollection.dtd">
<TargetCollection TYPE="TrafficLight">
    <CollectionItem NAME="trafficLight">
        <Schedule>
            <IntervalSchedule TIME_UNIT="Min" INTERVAL="30"/>
        </Schedule>
    </CollectionItem>
</TargetCollection>

```

After the command-line tool generates the target metadata and collection files, you can create the Metadata Plug-in archive. See [Section 13.4, "Creating the Plug-in Archive"](#).

17.4 Monitoring JMX Applications Deployed on Oracle Application Servers (OC4J)

The Java Management Extensions (JMX) framework improves manageability of your JMX-instrumented applications by enabling you to see what is happening inside. You gain insight into your applications and infrastructure through modular plug-ins called Managed Beans (MBeans). MBeans integrate with your application, components (such as Enterprise Java-Beans), or other resources to expose attributes (values) and operations.

The OJMX fetchlet, supplied with 10.2 Management Agents, enables you to monitor key metrics in your JMX-instrumented applications deployed on Oracle Application Server 10.1.3 or above. The fetchlet extends monitoring capabilities via JMX to the J2EE 1.4-compliant Oracle containers for J2EE (OC4J) servers themselves.

Monitoring JMX-instrumented applications/servers with Enterprise Manager entails defining a new target type that Enterprise Manager can monitor via Management Plug-ins. As with the Web services `wsccli` command-line tool, Enterprise Manager provides a `jmxcli` command-line tool to automate the generation of the target metadata and collection files.

Prerequisites

- Oracle Application Server 10.1.3 instance running on a specific host with a JMX-enabled application deployed on it that needs to be monitored as a target in Enterprise Manager.
- Management Agent version 10.2.0.2 or greater installed on that host.
- Oracle Management Server (OMS) version 10.2.0.2 or greater with which the Management Agent communicates.

Known Limitations

Currently, the `jmxcli` tool and OJMX fetchlet only allow you to browse and monitor MBeans (system and application-defined) that are available on the default MBeanServer on the target OC4J instance. The `jmxcli` tool primarily handles attributes and parameter and return values for operations that are OpenTypes. Examples: SimpleTypes, CompositeTypes, TabularTypes, and arrays of SimpleTypes.

17.4.1 Creating Metadata and Default Collection Files

As with Web services, the JMX command-line tool (`jmxcli`) simplifies creating the requisite target definition files: metadata and the default collection file. The tool is an offline configuration utility that connects you to an MBeanServer and enables you to browse available MBeans. It can also append metrics to an existing set of files during a subsequent invocation of the tool.

During a command-line tool session, you select specific MBeans and then choose the desired attributes/statistical values or operations Enterprise Manager needs to retrieve or invoke periodically on these MBeans to collect these values. The tool helps define packaging for these collected values as one or more Enterprise Manager metrics (with columns), and also enables you to specify a metric collection interval.

17.4.1.1 JMX Command-line Tool Syntax

The JMX command-line tool syntax is as follows for a JMX-enabled target on an OC4J: Note that usage has changed from earlier releases. The `cli` is now integrated with the `emctl` utility on the Management Agent.

```
cd Agent Instance Home/bin
emctl jmxcli TARGET_HOME
    [ -h hostname
      -p port
      -u username
      -c credential/password
      -w work directory
      -e true/false
      [-m MBeanName | -d jmx_domain | -s mBeanPattern]
    ]
```

TARGET_HOME is an Oracle Home directory 10.1.3 or later Oracle Application Server Container for J2EE (OC4J).

The `jmxcli` command accepts the following options:

- **-h** Host name of the OC4J. Default: "localhost"
- **-p** RMI/RMIS port of the OC4J. Default: "23791"
From the `ORACLE_HOME/opmn/bin` directory of your Application Server 10.1.3.0 or later instance, run `opmnctl status -l` to determine the RMI port for the OC4J for which MBeans were deployed.
- **-u** Valid user name for the OC4J. Default: "oc4jadmin"
- **-c** Password associated with the OC4J user specified by the `-u` option. Default: None. If you do not specify a password, you are prompted for the password.
- **-w** Directory where the metadata and default collection files created by the JMX command-line tool are placed. Default: Current directory. When invoking the command-line tool, you must have write permission on this directory to create subdirectories and add files. If the metadata and default collection files already exist within that directory, you have the option of appending to or overwriting the original files.
- **-e** Whether or not the RMIS connection to the OC4J is enabled (true or false). Default: false

You can also specify ONE of the following three parameters (`-m`, `-d` or `-s`) to retrieve a subset of MBeans available on the MBeanServer. By default, all MBeans on the MBeanServer are displayed for you to select from if none of these parameters are specified.

- **-m** MBean ObjectName of the required MBean that needs to be retrieved and examined. If this is an ObjectName pattern-matching multiple MBeans, you are shown a list of all MBeans that match the pattern, and you can select one at a time to work on.
- **-d** MBean domain of the required application whose MBeans need to be retrieved and examined. For example, you want to browse all MBeans for an application (myApp). MBeans for this application would be available in the JMX domain "myApp".
- **-s** MBean pattern-matching an existing set of MBeans from which the metrics are to be defined. The `-s` parameter allows bulk retrieval of JMX Attributes/Statistics from multiple MBeans of a similar type.

If you specify the `-s` parameter, the resulting metrics created during this `jmxcli` session appear as a table in the Enterprise Manager console with multiple rows — one row representing each MBean that matches the specified pattern, and with the MBean ObjectName as a key column. For example, if you specify `-s 'oc4j:j2eeType=Servlet,*'` the resulting metric will have multiple rows,

one for each servlet that matches the `ObjectName` pattern. Besides the `MBean ObjectName` column, other columns would be the attributes or fields from the return object of the operation, selected during the `jmxcli` session.

17.4.1.2 Generating the Files

To start the JMX command-line tool:

1. Go to the `$AGENT_HOME/bin` directory.
2. Run the following command:

```
emctl jmxcli Oracle Home of the target 10.1.3 or greater OC4J [OPTIONS]
```

Once invoked, the command-line interface automatically prompts you for the requisite information, as shown in [Example 17–10](#). If you need to quit a JMX command-line tool session, you can press `Control+C` at any point to exit. Session information will not be saved.

Example 17–9 Sample JMXCLI Invocation

```
./emctl jmxcli /scratch/shiphomes/oc4j/1013_SOA_M1/ -h localhost -p 12404 -m
'oc4j:J2EEApplication=orabpel,name=\"ServerBean\",*'
```

Example 17–10 Sample JMXCLI Session

```
oracleHome=/ade/sparmesw_10202_ssm/oracle
targetHome=/scratch/shiphomes/oc4j/1013_SOA_M1/
The Port is 12404
```

```
Connecting to server: localhost:12404
Connecting as user: oc4jadmin
Enter the password:
```

```
Obtained 1 MBeans matching pattern
oc4j:J2EEApplication=orabpel,name="ServerBean",*.
```

```
Enter the target type for this metric: [myJ2EEApp] myBPELApp
```

This is the target type for the new J2EE application as it should show up within Enterprise Manager.

```
Enter the target version: [1.0]
```

```
Enter the target metadata file: [./metadata/myBPELApp.xml]
```

This is the location of the metadata file that jmxcli generates. You must have write permission on the directories where the target metadata and default collection files are to be created.

```
Enter the default collections file: [./default_collection/myBPELApp.xml]
The file ./metadata/myBPELApp.xml already exists.
Do you want to overwrite the existing file, append to it, or quit <o/a/q>? [a] a
Appending to existing file: ./metadata/myBPELApp.xml.
The available targets are:
0: Identifies a deployed stateless session bean
   (oc4j:EJBModule="ejb_ob_
engine",J2EEApplication=orabpel,J2EEServer=standalone,j2eeType=StatelessSessionBea
n,name="ServerBean")
```

```
Enter the index of target/MBean you wish to monitor or press <Ctrl-C> to quit: 0
```

If multiple MBeans matched the -m <MBean pattern> specified when jmxcli was invoked, all MBean ObjectNames matching the pattern are listed during this part of the command-line session, at which point you can select one among the list. You can choose another MBean from the above list after creating metrics for the first one without exiting this jmxcli session.

If you want to append metrics from another MBean that does not match the above -m pattern, you must exit and start another jmxcli session with the MBean ObjectName/Pattern of the latter MBean, and create metrics from this MBean which will be appended to the original target metadata and default collection files from the previous jmxcli session. Using this method, you can append metrics created from multiple jmxcli sessions to the same target metadata and default collection files, if necessary.

Following metric source types are available for selected target(s):

- 0: JMX Attributes
- 1: JMX Operations
- 2: J2EE Statistics

Enter the index of your choice or press <Ctrl-C> to quit: 2

Statistics are:

- 0: CreateCount
- 1: ejbCreate()ClientActive
- 2: ejbCreate()ClientTime
- 3: ejbRemove()ClientActive
- 4: ejbRemove()ClientTime
- 5: MethodReadyCount
- 6: RemoveCount
- 7: setSessionContext(javax.ejb.SessionContext)ClientActive
- 8: setSessionContext(javax.ejb.SessionContext)ClientTime

Select one or more items as comma-separated indices: 0,6

JavaBean is : CreateCount

- 0: count
- 1: description
- 2: lastSampleTime
- 3: name
- 4: startTime
- 5: unit

This indicates that the Statistic call CreateCount is not a simple data type, but has a JavaBean pattern with the above listed properties, of which some may interest you.

Select one or more items as comma-separated indices: 0

JavaBean is : RemoveCount

- 0: count
- 1: description
- 2: lastSampleTime
- 3: name
- 4: startTime
- 5: unit

Select one or more items as comma-separated indices: 0

Number of possible columns in the resultant metric are 2.

Enter the name for this metric column at index=0 : [countOfCreateCount]
createCount

You can specify any meaningful name here. If you do not specify a name, the JMX command-line tool generates a default name that may not be appropriate in all cases.

Is this column a KEY Column <y/n>? [n]

In situations where multiple rows can be returned, as might be the case when the Attribute or return value of the Operation is TabularData, you need to specify one or more of your chosen metrics as "Key" columns.

```
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [createCount]
Enter the NLSID for column: [createCount]
Enter the UNIT for column "createCount": [count]
Do you want to create a threshold for this column <y/n>? [n] y
Creating threshold!!
Following operators are available for creating thresholds:
0: GT
1: EQ
2: LT
3: LE
4: GE
5: CONTAINS
6: NE
7: MATCH
```

If you want to create a threshold on this column, you can specify an operator and then a value that would trigger a CRITICAL or WARNING alert.

```
Enter the index of your choice or press <Ctrl-C> to quit: 0
Enter the CRITICAL threshold: [NotDefined] 100
Enter the WARNING threshold: [NotDefined] 85
Enter the number of occurrences that trigger threshold: [6] 3
```

This is the number of consecutive occurrences of above CRITICAL or WARNING values that would trigger an alert.

```
Enter the message to be used when threshold is triggered: [createCount is %value%
and has crossed warning (%warning_threshold%) or critical (%critical_threshold%)
threshold.]
Enter NLSID for the message used when threshold is triggered: [createCount_cond]
Enter the name for this metric column at index=1 : [countOfRemoveCount]
removeCount
Is this column a KEY Column <y/n>? [n]
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [removeCount]
Enter the NLSID for column: [removeCount]
Enter the UNIT for column "removeCount": [count]
Do you want to create a threshold for this column <y/n>? [n]
Enter the name of this metric: ServerBeanStats
Enter the label for this metric: [ServerBeanStats]

Do you want periodic collection for this metric <y/n>? [n] y
```

If the metric does not have to be collected periodically, as would be the case with real-time-only metrics, you can specify "no".

```
Enter the collection interval in seconds: 60
Periodic collection interval is: 60 seconds.
```

```
Do you want to create another metric <y/n>? [n] y
The available targets are:
0: Identifies a deployed stateless session bean
    (oc4j:EJBModule="ejb_ob_
engine",J2EEApplication=orabpel,J2EEServer=standalone,j2eeType=StatelessSessionBea
n,name="ServerBean")
```

Enter the index of target/MBean you wish to monitor or press <Ctrl-C> to quit: 0

If multiple MBeans match the MBean pattern for the -m option (specified when jmxcli was invoked) you can select a different MBean from the above list for the next iteration of this command-line session.

Following metric source types are available for selected target(s):

- 0: JMX Attributes
- 1: JMX Operations
- 2: J2EE Statistics

Enter the index of your choice or press <Ctrl-C> to quit: 0

Attributes are:

- 0: activeInstances Return Value: int
- 1: activeInstancesHighWaterMark Return Value: int
- 2: eventProvider Return Value: boolean
- 3: maxInstances Return Value: int
- 4: minInstances Return Value: int
- 5: ObjectName Return Value: javax.management.ObjectName
- 6: stateManageable Return Value: boolean
- 7: statisticsProvider Return Value: boolean
- 8: stats Return Value: javax.management.j2ee.statistics.Stats
- 9: transactionTimeout Return Value: int

Select one or more items as comma-separated indices: 0,3,4

Number of possible columns in the resultant metric are 3.

Enter the name for this metric column at index=0 : [activeInstances]

Is this column a KEY Column <y/n>? [n]

Is this column for SUMMARY_UI <y/n>? [n]

Enter the label for column: [activeInstances]

Enter the NLSID for column: [activeInstances]

Enter the UNIT for column "activeInstances": [millisec, kb etc..]

Do you want to create a threshold for this column <y/n>? [n]

Enter the name for this metric column at index=1 : [maxInstances]

Is this column a KEY Column <y/n>? [n]

Is this column for SUMMARY_UI <y/n>? [n]

Enter the label for column: [maxInstances]

Enter the NLSID for column: [maxInstances]

Enter the UNIT for column "maxInstances": [millisec, kb etc..]

Do you want to create a threshold for this column <y/n>? [n]

Enter the name for this metric column at index=2 : [minInstances]

Is this column a KEY Column <y/n>? [n]

Is this column for SUMMARY_UI <y/n>? [n]

Enter the label for column: [minInstances]

Enter the NLSID for column: [minInstances]

Enter the UNIT for column "minInstances": [millisec, kb etc..]

Do you want to create a threshold for this column <y/n>? [n]

Enter the name of this metric: ServerBeanCount

Enter the label for this metric: [ServerBeanCount]

Do you want periodic collection for this metric <y/n>? [n] y

Enter the collection interval in seconds: 300

Periodic collection interval is: 300 seconds.

Do you want to create another metric <y/n>? [n] n

Written the metadata xml file: ./metadata/myBPELApp.xml.

Updated the default collection file for myBPELApp at location ./default_

```
collection/myBPELApp.xml.
Exiting...
```

After the JMX command-line tool generates the target metadata and collection files, you can create the Oracle Plug-in archive. A sample of each generated file from the command-line tool session above is shown in [Example 17-11](#) and [Example 17-12](#).

Example 17-11 Generated Target Metadata File

```
<!DOCTYPE TargetMetadata SYSTEM "../dtds/TargetMetadata.dtd">
<TargetMetadata META_VER="1.0" TYPE="myBPELApp" CATEGORY
_PROPERTIES="VersionCategory">
  <Display>
    <Label NLSID="myBPELAppNLSID">myBPELApp</Label>
    <ShortName NLSID="myBPELAppShortName">myBPELApp</ShortName>
    <Description NLSID="myBPELAppDescription">myBPELApp</Description>
  </Display>

  <Metric NAME="ServerBeanStats" TYPE="TABLE">
    <Display>
      <Label NLSID="ServerBeanStats">ServerBeanStats</Label>
    </Display>
    <TableDescriptor>
      <ColumnDescriptor NAME="createCount" TYPE="NUMBER">
        <Display>
          <Label NLSID="createCount">createCount</Label>
          <Unit NLSID="count">count</Unit>
        </Display>
      </ColumnDescriptor>
      <ColumnDescriptor NAME="removeCount" TYPE="NUMBER">
        <Display>
          <Label NLSID="removeCount">removeCount</Label>
          <Unit NLSID="count">count</Unit>
        </Display>
      </ColumnDescriptor>
    </TableDescriptor>
    <QueryDescriptor FETCHLET_ID="OJMX">
      <Property NAME="machine" SCOPE="INSTANCE">HTTPMachine</Property>
      <Property NAME="OracleHome" SCOPE="INSTANCE">OracleHome</Property>
      <Property NAME="oc4jInstanceName" SCOPE="INSTANCE"
OPTIONAL="TRUE">OC4JInstanceName</Property>
      <Property NAME="jvmId" SCOPE="INSTANCE" OPTIONAL="TRUE">JVMId</Property>
      <Property NAME="mgmtWebSite" SCOPE="INSTANCE"
OPTIONAL="TRUE">MgmtWebSite</Property>
      <Property NAME="authuser" SCOPE="INSTANCE"
OPTIONAL="TRUE">authUser</Property>
      <Property NAME="authpwd" SCOPE="INSTANCE"
OPTIONAL="TRUE">authPasswd</Property>
      <Property NAME="metric" SCOPE="GLOBAL">ServerBeanStats</Property>
      <Property NAME="delimiter" SCOPE="GLOBAL">|</Property>
      <Property NAME="name" SCOPE="GLOBAL">getStatistics</Property>
      <Property NAME="signature"
SCOPE="GLOBAL">objectName,statNames,languageCode,countryCode</Property>
      <Property NAME="returnType" SCOPE="GLOBAL">arrayOfComplexObjectBean</Property>
      <Property NAME="dontAddDefaultRowKey" SCOPE="GLOBAL">true</Property>
      <Property NAME="columnOrder"
SCOPE="GLOBAL">/CreateCount/count,/RemoveCount/count</Property>
      <Property NAME="arguments" SCOPE="GLOBAL">
        <![CDATA[<arguments>
          <argument>
```



```

        <value>oc4j:EJBModule="ejb_ob_
engine",J2EEApplication=orabpel,J2EEServer=standalone,j2eeType=StatelessSessionBea
n,name="ServerBean"</value>
    </argument>
    <argument>
        <value>CreateCount</value>
        <value>RemoveCount</value>
    </argument>
    <argument>
        <value>en</value>
    </argument>
    <argument>
        <value>US</value>
    </argument>
</arguments>]]>
    </Property>
</QueryDescriptor>
</Metric>

<Metric NAME="ServerBeanCount" TYPE="TABLE">
    <Display>
        <Label NLSID="ServerBeanCount">ServerBeanCount</Label>
    </Display>
    <TableDescriptor>
        <ColumnDescriptor NAME="activeInstances" TYPE="NUMBER">
            <Display>
                <Label NLSID="activeInstances">activeInstances</Label>
            </Display>
        </ColumnDescriptor>
        <ColumnDescriptor NAME="maxInstances" TYPE="NUMBER">
            <Display>
                <Label NLSID="maxInstances">maxInstances</Label>
            </Display>
        </ColumnDescriptor>
        <ColumnDescriptor NAME="minInstances" TYPE="NUMBER">
            <Display>
                <Label NLSID="minInstances">minInstances</Label>
            </Display>
        </ColumnDescriptor>
    </TableDescriptor>
    <QueryDescriptor FETCHLET_ID="OJMX">
        <Property NAME="machine" SCOPE="INSTANCE">HTTPMachine</Property>
        <Property NAME="OracleHome" SCOPE="INSTANCE">OracleHome</Property>
        <Property NAME="oc4jInstanceName" SCOPE="INSTANCE"
OPTIONAL="TRUE">OC4JInstanceName</Property>
        <Property NAME="jvmId" SCOPE="INSTANCE" OPTIONAL="TRUE">JVMId</Property>
        <Property NAME="mgmtWebSite" SCOPE="INSTANCE"
OPTIONAL="TRUE">MgmtWebSite</Property>
        <Property NAME="authuser" SCOPE="INSTANCE"
OPTIONAL="TRUE">authUser</Property>
        <Property NAME="authpwd" SCOPE="INSTANCE"
OPTIONAL="TRUE">authPasswd</Property>
        <Property NAME="metric" SCOPE="GLOBAL">ServerBeanCount</Property>
        <Property NAME="delimiter" SCOPE="GLOBAL">|</Property>
        <Property NAME="name" SCOPE="GLOBAL">getAttributes</Property>
        <Property NAME="signature"
SCOPE="GLOBAL">objectName,attributeNames,languageCode,countryCode</Property>
        <Property NAME="returnType"
SCOPE="GLOBAL">arrayOfComplexObjectBean</Property>
        <Property NAME="dontAddDefaultRowKey" SCOPE="GLOBAL">true</Property>

```

```

        <Property NAME="columnOrder"
SCOPE="GLOBAL">/activeInstances,/maxInstances,/minInstances</Property>
<Property NAME="arguments" SCOPE="GLOBAL">
    <![CDATA[<arguments>
        <argument>         <value>oc4j:EJBModule="ejb_ob_
engine",J2EEApplication=orabpel,J2EEServer=standalone,j2eeType=StatelessSessionBea
n,name="ServerBean"</value>
        </argument>
        <argument>
            <value>activeInstances</value>
            <value>maxInstances</value>
            <value>minInstances</value>
        </argument>
        <argument>
            <value>en</value>
        </argument>
        <argument>
            <value>US</value>
        </argument>
    </arguments>]]>
    </Property>
</QueryDescriptor>
</Metric>

<Metric NAME="Response" TYPE="TABLE">
    <Display>
        <Label NLSID="Response">Response</Label>
    </Display>
    <TableDescriptor>
        <ColumnDescriptor NAME="Status" TYPE="NUMBER">
            <Display>
                <Label NLSID="Status">Status</Label>
            </Display>
        </ColumnDescriptor>
    </TableDescriptor>
    <QueryDescriptor FETCHLET_ID="OJMX">
        <Property NAME="machine" SCOPE="INSTANCE">HTTPMachine</Property>
        <Property NAME="OracleHome" SCOPE="INSTANCE">OracleHome</Property>
        <Property NAME="oc4jInstanceName" SCOPE="INSTANCE"
OPTIONAL="TRUE">OC4JInstanceName</Property>
        <Property NAME="jvmId" SCOPE="INSTANCE" OPTIONAL="TRUE">JVMId</Property>
        <Property NAME="mgmtWebSite" SCOPE="INSTANCE"
OPTIONAL="TRUE">MgmtWebSite</Property>
        <Property NAME="authuser" SCOPE="INSTANCE"
OPTIONAL="TRUE">authUser</Property>
        <Property NAME="authpwd" SCOPE="INSTANCE"
OPTIONAL="TRUE">authPasswd</Property>
        <Property NAME="metric" SCOPE="GLOBAL">Response</Property>
        <Property NAME="delimiter" SCOPE="GLOBAL">|</Property>
        <Property NAME="name" SCOPE="GLOBAL">getAttributes</Property>
        <Property NAME="signature"
SCOPE="GLOBAL">objectName,attributeNames,languageCode,countryCode</Property>
        <Property NAME="returnType"
SCOPE="GLOBAL">arrayOfComplexObjectBean</Property>
        <Property NAME="dontAddDefaultRowKey" SCOPE="GLOBAL">true</Property>
        <Property NAME="columnOrder" SCOPE="GLOBAL">/state</Property>
        <Property NAME="arguments" SCOPE="GLOBAL">
            <![CDATA[<arguments>
                <argument>

```

```

<value>oc4j:J2EEServer=standalone,j2eeType=J2EEApplication,name=orabpel</value>
  </argument>
  <argument>
    <value>state</value>
  </argument>
  <argument>
    <value>en</value>
  </argument>
  <argument>
    <value>US</value>
  </argument>
</arguments>]]>
  </Property>
</QueryDescriptor>
</Metric>

<InstanceProperties>
  <InstanceProperty NAME="HTTPMachine">
    <Display>
      <Label NLSID="dms_HTTPMachine_iprop">Machine name</Label>
    </Display>
  </InstanceProperty>
  <InstanceProperty NAME="OracleHome">
    <Display>
      <Label NLSID="dms_OracleHome_iprop">Oracle home path</Label>
    </Display>
  </InstanceProperty>
  <InstanceProperty NAME="OC4JInstanceName" OPTIONAL="TRUE"><Display><Label
NLSID="OC4JInstanceNameiprop">OC4JInstanceName</Label></Display>home</InstanceProp
erty>
    <InstanceProperty NAME="JVMIId" OPTIONAL="TRUE"><Display><Label NLSID="JVMIId_
iprop">JVMIId</Label></Display>1</InstanceProperty>
    <InstanceProperty NAME="MgmtWebSite" OPTIONAL="TRUE"><Display><Label
NLSID="MgmtWebSite_
iprop">MgmtWebSite</Label></Display>default-web-site</InstanceProperty>
    <InstanceProperty NAME="URI" OPTIONAL="TRUE"><Display><Label
NLSID="URI">URI</Label></Display>/JMXSoapAdapter/JMXSoapAdapter</InstanceProperty>
    <InstanceProperty NAME="authUser" OPTIONAL="TRUE">
      <Display>
        <Label NLSID="dms_authUser_iprop">Username for Basic
authorization</Label>
      </Display>
    </InstanceProperty>
    <InstanceProperty NAME="authPasswd" OPTIONAL="TRUE" CREDENTIAL="TRUE">
      <Display>
        <Label NLSID="dms_authPasswd_iprop">Password for Basic
authorization</Label>
      </Display>
    </InstanceProperty>
    <InstanceProperty NAME="Version" OPTIONAL="TRUE"><Display><Label
NLSID="oc4j_version_iprop">Version of
myBPPELApp</Label></Display>1.0</InstanceProperty>
  </InstanceProperties>
</TargetMetadata>

```

Example 17–12 Generated Metric Collection File

```

<!DOCTYPE TargetCollection SYSTEM "../dtds/TargetCollection.dtd">
<!-- This file is generated by Collector at 2011-04-28 12:11:55 -->

```

```

<TargetCollection TYPE="myBPELApp" INCLUDE_DEFAULT="TRUE">
  <CollectionItem NAME="ServerBeanStats" UPLOAD="YES">
    <Schedule>
      <IntervalSchedule INTERVAL="60" TIME_UNIT="Sec"/>
    </Schedule>
    <MetricColl NAME="ServerBeanStats">
      <Condition COLUMN_NAME="createCount" CRITICAL="100"
WARNING="85" OPERATOR="GT" OCCURRENCES="3" MESSAGE="createCount is %value% and has
crossed warning (%warning_threshold%) or critical (%critical_threshold%)
threshold." MESSAGE-NLSID="createCount_cond"/>
    </MetricColl>
  </CollectionItem>
  <CollectionItem NAME="ServerBeanCount" UPLOAD="YES">
    <Schedule>
      <IntervalSchedule INTERVAL="300" TIME_UNIT="Sec"/>
    </Schedule>
    <MetricColl NAME="ServerBeanCount">
    </MetricColl>
  </CollectionItem>

  <CollectionItem NAME="Response" UPLOAD="YES">
    <Schedule>
      <IntervalSchedule INTERVAL="30" TIME_UNIT="Sec"/>
    </Schedule>
    <MetricColl NAME="Response">
      <Condition COLUMN_NAME="Status" CRITICAL="1"
WARNING="NotDefined" OPERATOR="NE" OCCURRENCES="2" MESSAGE="Status is %value% and
has crossed warning (%warning_threshold%) or critical (%critical_threshold%)
threshold." MESSAGE-NLSID="Status_cond"/>
    </MetricColl>
  </CollectionItem>
</TargetCollection>

```

17.4.2 Displaying Target Status Information

For the status information of your targets to appear correctly within the Enterprise Manager console, you must define a metric, called Response, that has a column, named Status, with a critical threshold set. The status of target instances of this type appears in the console as "Up" (available) if the metric value is below the critical threshold. When the threshold is exceeded, the target status appears as "Down" in the console.

You can create the Response metric in another `jmxcli` session (append the metric to the metadata and collection files created in an earlier session). [Example 17–13](#) illustrates adding a Response metric to previously generated metadata and collection files from a new command-line session.

Example 17–13 Adding a Response Metric

```

./emctl jmxcli /scratch/shiphomes
//oc4j/1013_PRODUCTION/ -p 12403 -c welcome1 -m 'oc4j:j2eeType=J2EEApplication,n
ame=orabpel,*'

oracleHome=/ade/sparmesw_10202_ssm/oracle
targetHome=/scratch/shiphomes//oc4j/1013_PRODUCTION/
The Port is 12403

Connecting to server: localhost:12403
Connecting as user: oc4jadmin

```

Obtained 1 MBeans matching pattern oc4j:j2eeType=J2EEApplication,name=orabpel,*.

Enter the target type for this metric: [myJ2EEApp] myBPELApp

Enter the target version: [1.0]

Enter the target metadata file: [./metadata/myBPELApp.xml]

Enter the default collections file: [./default_collection/myBPELApp.xml]

The file ./metadata/myBPELApp.xml already exists.

Do you want to overwrite the existing file, append to it, or quit <o/a/q>? [a] a
Appending to existing file: ./metadata/myBPELApp.xml.

The available targets are:

0: Identifies a J2EE application EAR that has been deployed
(oc4j:J2EEServer=standalone,j2eeType=J2EEApplication,name=orabpel)

Enter the index of target/MBean you wish to monitor or press <Ctrl-C> to quit: 0

Following metric source types are available for selected target(s):

- 0: JMX Attributes
- 1: JMX Operations

Enter the index of your choice or press <Ctrl-C> to quit: 0

Attributes are:

0: allAccessibleGroups	Return Value: java.util.Set
1: allAccessibleUsers	Return Value: java.util.Set
2: applicationRootDirectoryPath	Return Value: java.lang.String
3: archivePath	Return Value: java.lang.String
4: childApplicationNames	Return Value: [Ljava.lang.String;
5: childApplications	Return Value: [Ljava.xml.management.ObjectName;
6: dataSourcesDescriptor	Return Value: java.lang.String
7: deploymentDescriptor	Return Value: java.lang.String
8: ejbClassLoaderPath	Return Value: java.lang.String
9: eventProvider	Return Value: boolean
10: groups	Return Value: java.util.Set
11: iiopStubs	Return Value: [B
12: metricRulesDescriptor	Return Value: java.lang.String
13: Modules	Return Value: [Ljava.xml.management.ObjectName;
14: objectName	Return Value: java.lang.String
15: ohsRouting	Return Value: boolean
16: parentApplication	Return Value: java.xml.management.ObjectName
17: parentApplicationName	Return Value: java.lang.String
18: properties	Return Value: java.util.Properties
19: proprietaryDeploymentDescriptor	Return Value: java.lang.String
20: proxyInterfaceSQLObjects	Return Value: [Ljava.lang.String;
21: routingId	Return Value: java.lang.String
22: Server	Return Value: java.xml.management.ObjectName
23: sharedLibraryImports	Return Value:
[Loracle.oc4j.admin.management.shared.SharedLibraryImport;	
24: startTime	Return Value: long
25: state	Return Value: int
26: stateManageable	Return Value: boolean
27: statisticsProvider	Return Value: boolean
28: syntheticWebModules	Return Value:
oracle.oc4j.admin.management.shared.WebModule	
29: users	Return Value: java.util.Set
30: webSite	Return Value: java.lang.String
31: webSiteBindings	Return Value: java.util.Map

```
Select one or more items as comma-separated indices: 25

Number of possible columns in the resultant metric are 1.

Enter the name for this metric column at index=0 : [state] Status
Is this column a KEY Column <y/n>? [n]
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [Status]
Enter the NLSID for column: [Status]
Enter the UNIT for column "Status": [millisec, kb etc.. ]
Do you want to create a threshold for this column <y/n>? [n] y
Creating threshold!!
Following operators are available for creating thresholds:
0: GT
1: EQ
2: LT
3: LE
4: GE
5: CONTAINS
6: NE
7: MATCH
Enter the index of your choice or press <Ctrl-C> to quit: 6
Enter the CRITICAL threshold: [NotDefined] 1
Enter the WARNING threshold: [NotDefined]
Enter the number of occurrences that trigger threshold: [6] 2
Enter the message to be used when threshold is triggered: [Status is %value% and
has crossed warning (%warning_threshold%) or critical (%critical_threshold%)
threshold.]
Enter NLSID for the message used when threshold is triggered: [Status_cond]

Enter the name of this metric: Response
Enter the label for this metric: [Response]

Do you want periodic collection for this metric <y/n>? [n] y
Enter the collection interval in seconds: 30
Periodic collection interval is: 30 seconds.

Do you want to create another metric <y/n>? [n] n
Written the metadata xml file: ./metadata/myBPELApp.xml.
Updated the default collection file for myBPELApp at location ./default_collecti
on/myBPELApp.xml.
Exiting...

Please note that the Response metric collected in this jmxcli session would be
appended to the metadata and default_collection file created in an earlier session
of the tool. (User can chose to overwrite the earlier file as well if they
specific the "o" option to the following prompt)

Do you want to overwrite the existing file, append to it, or quit <o/a/q>? [a] a
```

17.5 Monitoring a Standalone JMX-instrumented Java Application or JVM Target

Note: If your Java application is not JMX-instrumented, but you want to monitor the J2SE 5.0 JVM on which it is running, go directly to [Section 17.7.3, "Configuring a Standalone Java Application or JVM Target"](#) to create target instances of type JVM. This enables you to monitor these JVMs in Enterprise Manager, preferably from an Enterprise Manager Agent installed on the same host as your JVM. However, the prerequisites and known limitations discussed below still apply.

Enterprise Manager provides an out-of-box JVM target type. This enables you to add and configure metrics from standalone J2SE1.5 JVMs that are enabled for remote management in Enterprise Manager version 10.2.0.3 or later.

If your standalone Java application exposes data through JMX MBeans as for a J2EE application deployed on an Oracle Container for J2EE, you can use the JMX command-line tool to define such an application as an Enterprise Manager target type and generate a metadata and default collection file for this target type. You can monitor your standalone application targets from an Enterprise Manager Agent, preferably installed on the same host as your JVM. Multiple JVMs running on that host can be monitored by the same Enterprise Manager Agent.

You can collect metrics from user-defined MBeans on a standalone (J2SE1.5-based) JVM and place them into Enterprise Manager using the JMX fetchlet. The fetchlet is designed for a standalone Sun J2SE1.5-based (or later) JVM containing user-defined MBeans that use JMX OpenTypes as arguments and return values.

Prerequisites

- Java virtual machine J2SE 1.5 or higher instance running on a specific host. This JVM could be running a JMX-enabled application that exposes metrics via MBeans that need to be monitored as a target in Enterprise Manager. If the application does not expose MBeans, the JVM itself could be monitored using the built-in JVM target type provided in Enterprise Manager. See [Section 17.7.3, "Configuring a Standalone Java Application or JVM Target"](#) for more information.
- JMX agent enabled for local access. Set this system property when you start the JVM or Java application:

```
com.sun.management.jmxremote
```

- Monitoring and management from remote systems enabled. Set this system property when you start the JVM:

```
com.sun.management.jmxremot.port=portNum
```

For additional information about enabling the JVM for remote management, see the following document:

```
http://java.sun.com/j2se/1.5.0/docs/guide/management/agent.html#remote
```

- Management Agent version 10.2.0.3 or later installed on that host.
- Oracle Management Server (OMS) version 10.2.0.3 or greater with which the Management Agent communicates.

Known Limitations

Currently, the `jmxcli` tool only allows you to browse and monitor MBeans (platform and application-defined) that are available on the default platform MBeanServer on the target JVM instance. The tool does not support monitoring a custom MBeanServer on the target JVM instance. The `jmxcli` tool primarily handles attributes as well as parameter and return values for operations that are OpenTypes, such as SimpleTypes, CompositeTypes, TabularTypes, and arrays of SimpleTypes.

17.5.1 Generating Metadata and Default Collection Files

As with Web services and the J2EE application on OC4J, the command-line tool (`jmxcli`) simplifies creating the requisite target definition files: metadata and the default collection file for a standalone JMX-instrumented Java application. The tool is an offline configuration utility that connects you to an MBeanServer on a J2SE1.5 or higher JVM and enables you to browse available MBeans. It can also append metrics to an existing set of files during a subsequent invocation of the tool.

During a command-line tool session, you select specific MBeans and then choose the desired attributes/statistical values or operations Enterprise Manager needs to retrieve or invoke periodically on these MBeans to collect these values. The tool helps define packaging for these collected values as one or more Enterprise Manager metrics (with columns), and also enables you to specify a metric collection interval.

17.5.1.1 JMX Command-line Tool Syntax

The JMX command-line tool syntax is as follows:

```
cd Agent Instance dir/bin
emctl jmxcli -t JVM
[ -l JMXServiceURL
  -h hostname
  -p port
  -u username
  -c credential/password
  -w work directory
  -e true/false
  [-m MBeanName | -d jmx_domain | -s mBeanPattern]
]
```

The `jmxcli` command accepts the following options:

- **-t JVM** Indicates that the MBeanServer is on a standalone JVM
- **-l JMXServiceURL** of the target JVM
- **-h** Hostname of the JVM. Default: "localhost" if the `-l` option is not specified
- **-p** RMI/RMIS port of the JVM. Default: "23791" if the `-l` option is not specified. From the `ORACLE_HOME/opmn/bin` directory of your Application Server 10.1.3.0 or later instance, run `opmnctl status -l` to determine the RMI port for the OC4J for which MBeans were deployed.
- **-u** Valid user name for the JVM. Default: None
- **-c** Password for the above user. Default: None. The password is only used to retrieve data and is not stored anywhere.
- **-w** Work directory where the metadata and default collection files are created. Default: Current directory. When invoking the command-line tool, you must have write permission on this directory to create subdirectories and add files. If the

metadata and default collection files already exist within that directory, you have the option of appending to or overwriting the original files.

- **-e** True for enabling the SSL connection to the JVM. Default: false

You can also specify ONE of the following three parameters (**-m**, **-d** or **-s**) to retrieve a subset of MBeans available on the MBeanServer. By default, all MBeans on the MBeanServer are displayed for you to select from if none of these parameters are specified.

- **-m** MBean ObjectName of the required MBean that needs to be retrieved and examined. If this is an ObjectName pattern-matching multiple MBeans, you are shown a list of all MBeans that match the pattern, and you can select one at a time to work on.
- **-d** MBean domain of the required application whose MBeans need to be retrieved and examined. For example, you want to browse all MBeans for an application (myApp). MBeans for this application would be available in the JMX domain "myApp".
- **-s** MBean pattern matching an set of similar MBeans from which the metrics are to be defined. The **-s** parameter allows bulk retrieval of JMX Attributes/Statistics from multiple MBeans of a similar type.

If you specify the **-s** parameter, the resulting metrics created during this `jmxcli` session appear as a table in the Enterprise Manager console with multiple rows — one row representing each MBean that matches the specified pattern, and with the MBean ObjectName as a key column. For example, if you specify **-s 'oc4j:j2eeType=Servlet,*'** the resulting metric will have multiple rows, one for each servlet that matches the ObjectName pattern. Besides the MBean ObjectName column, other columns would be the attributes or fields from the return object of the operation, selected during the `jmxcli` session.

17.5.1.2 Generating the Files

The following steps explain how to prepare for and then use the JMX command-line tool to generate the files.

1. Bring up the standalone JVM instance with the MBeans. The following example shows an invocation of the JVM:

```
JDK15/bin/java -Dcom.sun.management.jmxremote
-Dcom.sun.management.jmxremote.port=6789
-Dcom.sun.management.jmxremote
-Dcom.sun.management.jmxremote.authenticate=false
-Dcom.sun.management.jmxremote.ssl=false MyJMXEnabledApp $*
```

The `jmxcli` tool connects to the port number above as a JSR-160 client.

2. Go to the `$ORACLE_HOME/bin` directory of the 10.2.0.3 or later version of the Enterprise Manager Agent.
3. Set the environment variable as follows:

```
setenv USER_JARS /myAppHome/myJar1.jar;/myAppHome/myJar2.jar
```

This step is needed if custom classes are being returned in attributes and/or operations in any of the MBeans registered with the target MBeanServer. The Enterprise Manager Agent (fetchlet) can only effectively monitor attributes and/or operations that return JMX OpenTypes, but it could also handle Java Bean properties (through getters and setters) on any custom classes.

Note: If the application-defined MBeans are returning custom classes, you need to also set up the corresponding user jar file in the CLASSPATH of the Enterprise Manager Agent monitoring this application. To do this, manually insert the location of this jar into the \$ORACLE_HOME/sysman/config/classpath.lst file on the Enterprise Manager Agent, then restart the Enterprise Manager Agent.

4. Run the following command:

```
./emctl jmxcli -t JVM -h localhost -p 6789 u user -c password
```

where:

- **-t** JVM indicates that the MBeanServer is running on a standard JVM
- **-h** Host name where the JVM is running
- **-p** Port number that enables the JVM for JSR-160 remote access

You can also specify an **-l** *JMXServiceURL* option instead of **-h** *host* and **-p** *port* options.

You can invoke *jmxcli* with a **-w** *work directory* option to create the metadata and default collection files in the specified work directory. If you do not specify **-w** when you start *jmxcli*, it defaults to the current directory, which is the directory where you start *jmxcli*.

Once invoked, the command-line interface automatically prompts you for the requisite information, as shown in [Example 17-14](#). For most of the prompts, you can just press enter to use defaults. If you need to quit a JMX command-line tool session, you can press Control+C at any point to exit. Session information will not be saved.

When the session concludes after you exit, the result will be a *myJ2EEApp.xml* file (or whatever target type you specified) as *metadata/myJ2EEApp.xml*, and a *default_collection/myJ2EEApp.xml* file if you specified periodic collection.

Sample JMXCLI Invocations

The following sample enables you to browse all MBeans on a remote MBeanServer:

```
./emctl jmxcli -t JVM -p 6789 (the host defaults to "localhost")
```

The following sample invokes the command-line interface and filters MBeans based on the MBeanPattern specified as the argument for the **-m** option:

```
./emctl jmxcli -t JVM -p 6789 -m "java.lang:*
```

Example 17-14 Sample JMXCLI Session

```
oracleHome=/ade/sparmesw_emas_ml/oracle
userJars=
```

```
Connecting to server: localhost:6789
Connecting without authentication. For specifying username and password use
the -u and -c options.
```

```
Obtained 14 MBeans matching pattern java.lang:*
```

```
Enter the target type for this metric: [myJ2EEApp] myJavaApp
```

```
Enter the target version: [1.0]
```

```
Enter the target metadata file: [./metadata/myJavaApp.xml]

Enter the default collections file: [./default_collection/myJavaApp.xml]

Enter a label for this target type: [myJavaApp]

Enter a description for this target type: [myJavaApp]
The available targets are:
0: sun.management.CompilationImpl
    (java.lang:type=Compilation)

1: sun.management.MemoryManagerImpl
    (java.lang:name=CodeCacheManager,type=MemoryManager)

2: sun.management.GarbageCollectorImpl
    (java.lang:name=Copy,type=GarbageCollector)

3: sun.management.MemoryPoolImpl
    (java.lang:name=Eden Space,type=MemoryPool)

4: sun.management.RuntimeImpl
    (java.lang:type=Runtime)

5: sun.management.ClassLoadingImpl
    (java.lang:type=ClassLoading)

6: sun.management.MemoryPoolImpl
    (java.lang:name=Survivor Space,type=MemoryPool)

7: sun.management.ThreadImpl
    (java.lang:type=Threading)

8: sun.management.GarbageCollectorImpl
    (java.lang:name=MarkSweepCompact,type=GarbageCollector)

9: com.sun.management.UnixOperatingSystem
    (java.lang:type=OperatingSystem)

10: sun.management.MemoryImpl
    (java.lang:type=Memory)

11: sun.management.MemoryPoolImpl
    (java.lang:name=Code Cache,type=MemoryPool)

12: sun.management.MemoryPoolImpl
    (java.lang:name=Tenured Gen,type=MemoryPool)

13: sun.management.MemoryPoolImpl
    (java.lang:name=Perm Gen,type=MemoryPool)

Enter the index of target/MBean you wish to monitor or press <Ctrl-C> to quit: 4

Following metric source types are available for selected target(s):
0: JMX Attributes

Enter the index of your choice or press <Ctrl-C> to quit: 0

Attributes are:
```

```
0: BootClassPath      Return Value: java.lang.String
1: BootClassPathSupported Return Value: boolean
2: ClassPath      Return Value: java.lang.String
3: InputArguments      Return Value: [Ljava.lang.String;
4: LibraryPath      Return Value: java.lang.String
5: ManagementSpecVersion Return Value: java.lang.String
6: Name      Return Value: java.lang.String
7: SpecName      Return Value: java.lang.String
8: SpecVendor      Return Value: java.lang.String
9: SpecVersion      Return Value: java.lang.String
10: StartTime      Return Value: long
11: SystemProperties      Return Value:
javax.management.openmbean.TabularData
12: Uptime      Return Value: long
13: VmName      Return Value: java.lang.String
14: VmVendor      Return Value: java.lang.String
15: VmVersion      Return Value: java.lang.String
Select one or more items as comma-separated indices: 6,7,8

Number of possible columns in the resultant metric are 3.

Enter the name for this metric column at index=0 : [Name]
Is this column a KEY Column <y/n>? [n] y
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [Name]
Enter the NLSID for column: [Name]
Enter the UNIT for column "Name": [millisec, kb etc.. ]

Enter the name for this metric column at index=1 : [SpecName]
Is this column a KEY Column <y/n>? [n]
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [SpecName]
Enter the NLSID for column: [SpecName]
Enter the UNIT for column "SpecName": [millisec, kb etc.. ]
Do you want to create a threshold for this column <y/n>? [n]

Enter the name for this metric column at index=2 : [SpecVendor]
Is this column a KEY Column <y/n>? [n]
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [SpecVendor]
Enter the NLSID for column: [SpecVendor]
Enter the UNIT for column "SpecVendor": [millisec, kb etc.. ]
Do you want to create a threshold for this column <y/n>? [n]

Enter the name of this metric: RuntimeMetric
Enter the label for this metric: [RuntimeMetric]

Do you want periodic collection for this metric <y/n>? [n] y
Enter the collection interval in seconds: 300
Periodic collection interval is: 300 seconds.

Do you want to create another metric <y/n>? [n]
Written the metadata xml file: ./metadata/myJavaApp.xml.
Creating new file: ./default_collection/myJavaApp.xml.
Updated the default collection file for myJavaApp at location
./default_collection/myJavaApp.xml.
Exiting...
```

17.5.2 Using the Metadata and Default Collection Files

Look at the `currentDir/metadata` and `currentDir/default_collection` directories to see the `myTarget.xml` files (for the target type you specified earlier).

You can use these files as follows:

- Convert the files to an Oracle Plug-in Archive (OPAR) and push them from OMS to the Agent and target instances created from OMS. See [Chapter 13, "Validating, Packaging, and Deploying the Plug-in"](#) and [Section 17.7.3, "Configuring a Standalone Java Application or JVM Target"](#).
- Edit the files, extract the metric definitions and `QueryDescriptors`, move them to other metadata and default collection files, and post-process them by creating `ExecutionDescriptors` as needed.

If you want the status information of your targets to appear correctly in the Enterprise Manager console, you need to define a Response metric. See [Section 17.4.2, "Displaying Target Status Information"](#) for more information.

17.6 Monitoring JMX Applications Deployed on Oracle WebLogic Application Servers

The JMX fetchlet, supplied with 11.1 Management Agents, enables you to monitor key metrics in your JMX-instrumented applications deployed on Oracle WebLogic Application Server 9.x or later.

Monitoring JMX-instrumented applications with Enterprise Manager entails defining a new target type that Enterprise Manager can monitor via Management Plug-ins. As with the Web services `wscli` command-line tool (and as was possible for Oracle Application Servers (OC4J)), Enterprise Manager provides an `jmxcli` command-line tool to automate the generation of the target metadata and collection files for custom JMX instrumented applications on weblogic servers..

Prerequisites

- Oracle WebLogic Server 9.x instance running on a specific host with a JMX-enabled application deployed on it that needs to be monitored as a target in Enterprise Manager.
- Management Agent version 11.1 or greater installed (preferably) on that host.
- Oracle Management Server (OMS) version 10.2.0.4 or greater with which the Management Agent communicates.
- The `jmxcli` tool primarily handles attributes and parameter and return values for operations that are `OpenTypes`. Examples: `SimpleTypes`, `CompositeTypes`, `TabularTypes`, and arrays of `SimpleTypes`.

17.6.1 Creating Metadata and Default Collection Files using `jmxcli`

As with Web services, the JMX command-line tool (`jmxcli`) simplifies creating the requisite target definition files: metadata and the default collection file. The tool is an offline configuration utility that connects you to an `MBeanServer` and enables you to browse available `MBeans`. It can also append metrics to an existing file during a subsequent invocation of the tool. During a command-line tool session, you select specific `MBeans` and then choose the desired JMX attributes/statistical values the Enterprise Manager needs to retrieve or JMX operations that need to be invoked periodically on these `MBeans` to collect these values. The tool helps define packaging for these collected values as one or more Enterprise Manager metrics (with columns),

and also enables you to specify a metric collection interval and thresholds for the columns.

17.6.1.1 JMX Command-line Tool Syntax

The JMX command-line tool syntax is as follows for a JMX-enabled target on an Oracle WebLogic Application Server:

```
./emctl jmxcli -t WebLogic [help|options]
  where options are:
      [ -l JMX ServiceURL
        -u username
        -c credential/password
        -w work directory
        [-m MBeanName | -d jmx_domain | -s mBeanPattern]
      ]
```

The `jmxcli` command accepts the following options:

- `l` - JMXServiceURL to the WebLogic managed server hosting the custom MBeans in the form

service:jmx:t3://host:t3port/jndi/weblogic.management.mbeanservers.runtime
- `u` - Valid user name for the WebLogic domain. Default: "weblogic"
- `c` - Password associated with the user specified by the `-u` option. Default: None. If you do not specify a password, you are prompted for the password.
- `w` - Directory where the metadata and default collection files created by the JMX command-line tool are placed. Default: Current directory.

When invoking the command-line tool, you must have write permission on this directory to create subdirectories and add files. If the metadata and default collection files already exist within that directory, you have the option of appending to or overwriting the original files.

You can also specify ONE of the following three parameters (`-m`, `-d` or `-s`) to retrieve a subset of MBeans available on the MBeanServer. By default, all MBeans on the MBeanServer are displayed for you to select from if none of these parameters are specified.

- `m` - MBean ObjectName of the required MBean that needs to be retrieved and examined. If this is an ObjectName pattern-matching multiple MBeans, you are shown a list of all MBeans that match the pattern, and you can select one at a time to work on.
- `d` - MBean domain of the required application whose MBeans need to be retrieved and examined. For example, you want to browse all MBeans for an application (myApp). MBeans for this application would be available in the JMX domain "myApp".
- `s` - MBean pattern-matching an existing set of MBeans from which the metrics are to be defined. The `-s` parameter allows retrieval of JMX Attributes/Statistics from multiple MBeans of a similar type into one Metric.

If you specify the `-s` parameter, the resulting metrics created during this `jmxcli` session appear as a table in the Enterprise Manager console with multiple rows - one row representing each MBean that matches the specified pattern (with the MBean ObjectName as a key column if no other key columns are defined). For example, if you specify `-s 'com.bea:Type=ServletRuntime,*'` the resulting metric will have multiple rows, one for each servlet that matches the ObjectName pattern.

Besides the MBean ObjectName key column, other columns would be the attributes or fields from the return object of the operation, selected during the `jmxcli` session.

17.6.1.2 Generating the Files

To start the JMX command-line tool:

1. Go to the `$AGENT_HOME/bin` directory.
2. Run the following command:

```
./emctl jmxcli -t WebLogic [OPTIONS]
```

Once invoked, the command-line interface automatically prompts you for the requisite information, as shown in the following example. If you need to quit a JMX command-line tool session, you can press Control+C at any point to exit. Session information will not be saved.

The following example illustrates a sample `jmxcli` session:

Example 17–15 `jmxcli` Session

```
$ ./emctl jmxcli -t WebLogic -l
"service:jmx:t3://stbct14:22048/jndi/weblogic.management.mbeanservers.runtime" -u
weblogic -c welcome1 -s "*:type=soainfra_bpel_requests,*"
```

NOTE 1: The `-s` option above will result in a metric with as many rows as the number of MBeans which match the ObjectName pattern specified, every time the metric is collected by the agent. If you need to always collect from a specific Mbean then use the `-m <ObjectName>` option instead of the `-s <Mbean pattern>` used in above example.

NOTE 2: If you need to use `t3s` to connect to the weblogic server then the following `env` variable needs to be set before invoking the `jmxcli`

```
setenv USER_JAVA_PROPS=-Dweblogic.security.TrustKeyStore=CustomTrust
-Dweblogic.security.CustomTrustKeyStoreFileName=$ORACLE_
HOME/sysman/config/montrust/AgentTrust.jks
-Dweblogic.security.SSL.enforceConstraints=off
-Dweblogic.security.SSL.ignoreHostnameVerification=true
-Djavax.net.ssl.trustStore=$ORACLE_HOME/sysman/config/montrust/AgentTrust.jks
(or set USER_JAVA_PROP= ... equivalent on win32)
setenv USER_JARS <; separated list of jars needed in classpath if custom
authentication modules are involved in SSL connection>
```

A semi-colon is used as a delimiter for the list of jar files.

Example: `setenv USER_JARS "jar1;jar2;jar3"`

In some cases, if MBeans return custom WebLogic objects in their MBeanInfo, the `weblogic.jar` may need to be set to the above `env` variable before invoking the `jmxcli`.

Example: `setenv USER_JARS $BEA_HOME/server/lib/weblogic.jar`

```
oracleHome=/ade/sparmesw_egcli/oracle/work/middleware/oms
userJars=
Connecting to server:
  service:jmx:t3://stbct14:22048/jndi/weblogic.management.mbeanservers.runtime
Connecting as user: weblogic
Obtained 3 MBeans matching pattern *:type=soainfra_bpel_requests,*.
```

```
Enter the target type for this metric: [myJ2EEApp] myCustomWLApp
Enter the target version: [1.0]
Enter the target metadata file: [./metadata/myCustomWLApp.xml]
```

```
Enter the default collections file: [./default_collection/myCustomWlApp.xml]
Enter a label for this target type: [myCustomWlApp]
Enter a description for this target type: [myCustomWlApp]
The available targets are:
0: DMS metric mbean
    (oracle.dms:name=/soainfra/engines/bpel/requests/engine,type=soainfra_
    bpel_requests)
1: DMS metric mbean
    (oracle.dms:name=/soainfra/engines/bpel/requests/system,type=soainfra_
    bpel_requests)
2: DMS metric mbean
    (oracle.dms:name=/soainfra/engines/bpel/requests/invoke,type=soainfra_
    bpel_requests)
Following metric source types are available for selected target(s):
0: JMX Attributes
Enter the index of your choice or press <Ctrl-C> to quit: 0
Attributes are:
0: active_count    Return Value: java.lang.Integer
1: active_maxValue Return Value: java.lang.Integer
2: active_minValue Return Value: java.lang.Integer
3: active_value    Return Value: java.lang.Integer
4: Name            Return Value: java.lang.String
5: Parent          Return Value: java.lang.String
6: scheduled_count Return Value: java.lang.Integer
7: scheduled_maxValue Return Value: java.lang.Integer
8: scheduled_minValue Return Value: java.lang.Integer
9: scheduled_value Return Value: java.lang.Integer
10: threadCount_count Return Value: java.lang.Integer
11: threadCount_maxValue Return Value: java.lang.Integer
12: threadCount_minValue Return Value: java.lang.Integer
13: threadCount_value Return Value: java.lang.Integer

Select one or more items as comma separated indices: 4,0,1,2
Number of possible columns in the resultant metric are 4.

Enter the name for this metric column at index=0 : [Name]
Is this column a KEY Column <y/n>? [n] y

Specifying "y" signifies that the value of this column is unique in case multiple rows are
returned.

Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [Name]
Enter the NLSID for column: [Name]
Enter the UNIT for column "Name": [millisec, kb etc.. ]

Enter the name for this metric column at index=1 : [active_count]
Is this column a KEY Column <y/n>? [n]
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [active_count]
Enter the NLSID for column: [active_count]
Enter the UNIT for column "active_count": [millisec, kb etc.. ]
Do you want to create a threshold for this column <y/n>? [n] y
Creating threshold!!
Following operators are available for creating thresholds:
0: GT
1: EQ
2: LT
3: LE
4: GE
5: CONTAINS
```


6: NE
7: MATCH

```

Enter the index of your choice or press <Ctrl-C> to quit: 0
Enter the CRITICAL threshold: [NotDefined] 50
Enter the WARNING threshold: [NotDefined] 45
Enter the number of occurrences that trigger threshold: [6] 3
Enter the message to be used when threshold is triggered: [active_count is %value%
and has crossed warning (%warning_threshold%) or critical (%critical_threshold%)
threshold.]
Enter NLSID for the message used when threshold is triggered: [active_count_cond]
Enter the name for this metric column at index=2 : [active_maxValue]
Is this column a KEY Column <y/n>? [n]
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [active_maxValue]
Enter the NLSID for column: [active_maxValue]
Enter the UNIT for column "active_maxValue": [millisec, kb etc.. ]
Do you want to create a threshold for this column <y/n>? [n]

Enter the name for this metric column at index=3 : [active_minValue]
Is this column a KEY Column <y/n>? [n]
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [active_minValue]
Enter the NLSID for column: [active_minValue]
Enter the UNIT for column "active_minValue": [millisec, kb etc.. ]
Do you want to create a threshold for this column <y/n>? [n]

Enter the name of this metric: bpm_requests
Enter the label for this metric: [bpm_requests]

Do you want periodic collection for this metric <y/n>? [n] y
Enter the collection interval in seconds: 300
Periodic collection interval is: 300 seconds.

Do you want to create another metric <y/n>? [n]
Written the metadata xml file: ./metadata/myCustomWLApp.xml.
Creating new file: ./default_collection/myCustomWLApp.xml.
Updated the default collection file for myCustomWLApp at location ./default_
collection/myCustomWLApp.xml.
Exiting...
```

Example 17-16 Sample jmxcli Invocation (using -m and defining multiple metrics from multiple Mbeans in one jmxcli session)

```

$ ./emctl jmxcli -t WebLogic -l
"service:jmx:t3://stbct14:22048/jndi/weblogic.management.mbeanservers.runtime" -u
weblogic -c welcome1 -m
"com.bea:ApplicationRuntime=soa-infra,WebAppComponentRuntime=soa_server1_/b2b,*"

oracleHome=/ade/sparmesw_egcli/oracle/work/middleware/oms
userJars=
Connecting to server: service:jmx:t3://stbct14:22048/jndi/weblogic.management.m
beanservers.runtime
Connecting as user: weblogic
Obtained 8 MBeans matching pattern
com.bea:ApplicationRuntime=soa-infra,WebAppComponentRuntime=soa_server1_/b2b,*.

Enter the target type for this metric: [myJ2EEApp] myCustomWLApp
```

Enter the target version: [1.0]

Enter the target metadata file: [./metadata/myCustomWlApp.xml]

Enter the default collections file: [./default_collection/myCustomWlApp.xml]
The file ./metadata/myCustomWlApp.xml already exists.

Do you want to overwrite the existing file, append to it, or quit <o/a/q>? [a]

Note: Because the file already exists, it will be appended.

Appending to existing file: ./metadata/myCustomWlApp.xml.

The available targets are:

0: (com.bea:ApplicationRuntime=soa-infra,Name=JspServlet,ServerRuntime=soa_server1,Type=ServletRuntime,WebAppComponentRuntime=soa_server1_/b2b)
1: (com.bea:ApplicationRuntime=soa-infra,Name=transportServlet,ServerRuntime=soa_server1,Type=ServletRuntime,WebAppComponentRuntime=soa_server1_/b2b)
2: (com.bea:ApplicationRuntime=soa-infra,Name=transportServletV,ServerRuntime=soa_server1,Type=ServletRuntime,WebAppComponentRuntime=soa_server1_/b2b)
3: (com.bea:ApplicationRuntime=soa-infra,Name=b2b_starter_wls,ServerRuntime=soa_server1,Type=ServletRuntime,WebAppComponentRuntime=soa_server1_/b2b)
4: (com.bea:ApplicationRuntime=soa-infra,Name=soa_server1_soa_server1_/b2b,ServerRuntime=soa_server1,Type=PageFlowsRuntime,WebAppComponentRuntime=soa_server1_/b2b)
5 (com.bea:ApplicationRuntime=soa-infra,Name=WebServiceServlet,ServerRuntime=soa_server1,Type=ServletRuntime,WebAppComponentRuntime=soa_server1_/b2b)
6: (com.bea:ApplicationRuntime=soa-infra,Name=RedirectUIServlet,ServerRuntime=soa_server1,Type=ServletRuntime,WebAppComponentRuntime=soa_server1_/b2b)
7: (com.bea:ApplicationRuntime=soa-infra,Name=FileServlet,ServerRuntime=soa_server1,Type=ServletRuntime,WebAppComponentRuntime=soa_server1_/b2b)

Enter the index of target/MBean you wish to monitor or press <Ctrl-C> to quit: 4
Following metric source types are available for selected target(s):
0: JMX Attributes
1: JMX Operations

Enter the index of your choice or press <Ctrl-C> to quit: 0

Attributes are:

0: AppName	Return Value: java.lang.String
1: ContextPath	Return Value: java.lang.String
2: HttpServerName	Return Value: java.lang.String
3: Name	Return Value: java.lang.String
4: PageFlows	Return Value: [Ljavaax.management.ObjectName;
5: Parent	Return Value: javax.management.ObjectName
6: ServerName	Return Value: java.lang.String
7: Type	Return Value: java.lang.String

Select one or more items as comma separated indices: 3,0,1

Number of possible columns in the resultant metric are 3.

Enter the name for this metric column at index=0 : [Name]

Is this column a KEY Column <y/n>? [n] y

Is this column for SUMMARY_UI <y/n>? [n]

Enter the label for column: [Name]

Enter the NLSID for column: [Name]

Enter the UNIT for column "Name": [millisec, kb etc..]

Enter the name for this metric column at index=1 : [AppName]

Is this column a KEY Column <y/n>? [n]

Is this column for SUMMARY_UI <y/n>? [n]

Enter the label for column: [AppName]

Enter the NLSID for column: [AppName]

```

Enter the UNIT for column "AppName": [millisec, kb etc.. ]
Do you want to create a threshold for this column <y/n>? [n]

Enter the name for this metric column at index=2 : [ContextPath]
Is this column a KEY Column <y/n>? [n]
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [ContextPath]
Enter the NLSID for column: [ContextPath]
Enter the UNIT for column "ContextPath": [millisec, kb etc.. ]
Do you want to create a threshold for this column <y/n>? [n]
Enter the name of this metric: PageFlowsRuntime
Enter the label for this metric: [PageFlowsRuntime]

Do you want periodic collection for this metric <y/n>? [n] y
Enter the collection interval in seconds: 3600
Periodic collection interval is: 3600 seconds.

Do you want to create another metric <y/n>? [n] y

```

This indicates more metrics need to be created in this `jmxcli` session. This process will repeat until you answer "n" to the question.

```

Do you want to create another metric <y/n>? [n]
Written the metadata xml file: ./metadata/myCustomWLApp.xml.
Updated the default collection file for myCustomWLApp at location ./default_collection/myCustomWLApp.xml.
Exiting...

```

After the JMX command-line tool generates the target metadata and collection files, you can create the Oracle Plug-in archive (OPAR).

17.6.1.3 Displaying Target Status Information

For the status information of your targets to appear correctly within the Enterprise Manager console, you must define a metric, called Response, that has a column, named Status, with a critical threshold set. The status of target instances of this type appears in the console as "Up" (available) if the metric value is below the critical threshold. When the threshold is exceeded, the target status appears as "Down" in the console.

You can create the Response metric in another `jmxcli` session (append the metric to the metadata and collection files created in an earlier session).

Example 17-17 Adding a Response Metric

```
setenv USER_JARS $T_WORK/middleware/wlserver_10.3/server/lib/weblogic.jar
```

This is required as some MBeans return WebLogic-specific classes which the JMX client (`jmxcli`) needs in its classpath.

```

$ ./emctl jmxcli -t WebLogic -l
"service:jmx:t3://stbct14:22048/jndi/weblogic.management.mbeanservers.runtime" -u
weblogic -c welcome1 -m com.bea.Type=ApplicationRuntime,Name=soa-infra,*"

```

For J2EE applications deployed on WebLogic it may be appropriate to make the `ActiveVersionState` JMX attribute of the `ApplicationRuntime` Mbean corresponding to the application deployment as the Status column. However, any other attribute of any other relevant Mbean to the application could also be used.

```

oracleHome=/ade/sparmesw_egcli/oracle/work/middleware/oms
userJars=

```

```
Connecting to server:
service:jmx:t3://stbct14:22048/jndi/weblogic.management.mbeanservers.runtime
Connecting as user: weblogic
Obtained 1 MBeans matching pattern
  com.bea.Type=ApplicationRuntime,Name=soa-infra,*
Enter the target type for this metric: [myJ2EEApp] myCustomWLApp
Enter the target version: [1.0]
Enter the target metadata file: [./metadata/myCustomWLApp.xml]
Enter the default collections file: [./default_collection/myCustomWLApp.xml]
The file ./metadata/myCustomWLApp.xml already exists.

Do you want to overwrite the existing file, append to it, or quit <o/a/q>? [a]
Appending to existing file: ./metadata/myCustomWLApp.xml.
The available targets are:
0: (com.bea.Name=soa-infra,ServerRuntime=soa_server1,Type=ApplicationRuntime)
Enter the index of target/MBean you wish to monitor or press <Ctrl-C> to quit: 0
Following metric source types are available for selected target(s):
  0: JMX Attributes
  1: JMX Operations
Enter the index of your choice or press <Ctrl-C> to quit: 0
Attributes are:
  0: ActiveVersionState      Return Value: java.lang.Integer
  1: ApplicationName         Return Value: java.lang.String
  2: ApplicationVersion      Return Value: java.lang.String
  3: ClassRedefinitionRuntime Return Value: javax.management.ObjectName
  4: ComponentRuntimes       Return Value: [Ljavaax.management.ObjectName;
  5: EAR                     Return Value: java.lang.Boolean
  6: HealthState             Return Value: weblogic.health.HealthState
  7: KodoPersistenceUnitRuntimes Return Value:
[Ljavax.management.ObjectName;
  8: LibraryRuntimes         Return Value: [Ljavaax.management.ObjectName;
  9: MaxThreadsConstraintRuntimes Return Value:
[Ljavax.management.ObjectName;
 10: MinThreadsConstraintRuntimes Return Value:
[Ljavax.management.ObjectName;
 11: Name                    Return Value: java.lang.String
 12: OptionalPackageRuntimes Return Value:
[Ljavax.management.ObjectName;
 13: Parent                  Return Value: javax.management.ObjectName
 14: QueryCacheRuntimes      Return Value: [Ljavaax.management.ObjectName;
 15: RequestClassRuntimes    Return Value:
[Ljavax.management.ObjectName;
 16: Type                    Return Value: java.lang.String
 17: WorkManagerRuntimes     Return Value: [Ljavaax.management.ObjectName;
 18: WseeRuntimes            Return Value: [Ljavaax.management.ObjectName;

Select one or more items as comma separated indices: 0

Number of possible columns in the resultant metric are 1.

Enter the name for this metric column at index=0 : [ActiveVersionState] Status

Note: The column name must be "Status".

Is this column a KEY Column <y/n>? [n]
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [Status]
Enter the NLSID for column: [Status]
Enter the UNIT for column "Status": [millisec, kb etc.. ]
Do you want to create a threshold for this column <y/n>? [n] y
Creating threshold!!
```

Following operators are available for creating thresholds:

```
0: GT
1: EQ
2: LT
3: LE
4: GE
5: CONTAINS
6: NE
7: MATCH
```

Enter the index of your choice or press <Ctrl-C> to quit: 6

Enter the CRITICAL threshold: [NotDefined] 2

Status of target is marked down if a CRITICAL THRESHOLD is triggered on the Status column of the Response Metric. In this case if value != ACTIVATED (such as: != 2)

Enter the WARNING threshold: [NotDefined]

Enter the number of occurrences that trigger threshold: [6] 1

Enter the message to be used when threshold is triggered: [Status is %value% and has crossed warning (%warning_threshold%) or critical (%critical_threshold%) threshold.]

Enter NLSID for the message used when threshold is triggered: [Status_cond]

Enter the name of this metric: Response

Note: The metric name must be "Response".

Enter the label for this metric: [Response]

Do you want periodic collection for this metric <y/n>? [n] y

Enter the collection interval in seconds: 30

Periodic collection interval is: 30 seconds.

Do you want to create another metric <y/n>? [n]

Written the metadata xml file: ./metadata/myCustomWLApp.xml.

Updated the default collection file for myCustomWLApp at location ./default_collection/myCustomWLApp.xml.

Exiting...

17.6.2 Using the Metadata and Default Collection Files

Look at the `currentDir/metadata` and `currentDir/default_collection` directories to see the `myTarget.xml` files (for the target type you specified earlier).

You can use these files as follows:

- Convert the files to an Oracle Plug-in Archive (OPAR). See [Section 13.4, "Creating the Plug-in Archive"](#).
- Move the OPAR to the OMS. See [Section 13.5, "Importing and Deploying the Plug-in Archive into Enterprise Manager"](#).
- Push the OPAR to the Agents. See [Section 13.5, "Importing and Deploying the Plug-in Archive into Enterprise Manager"](#).
- Create custom target instances. See [Section 17.7.4, "Adding a Target Instance for a Custom J2EE Application on WebLogic"](#)

If you want the status information of your targets to appear correctly in the Enterprise Manager console, you need to define a Response metric. See [Section 17.6.1.3, "Displaying Target Status Information"](#) for more information.

17.7 Adding a Target to a Management Agent

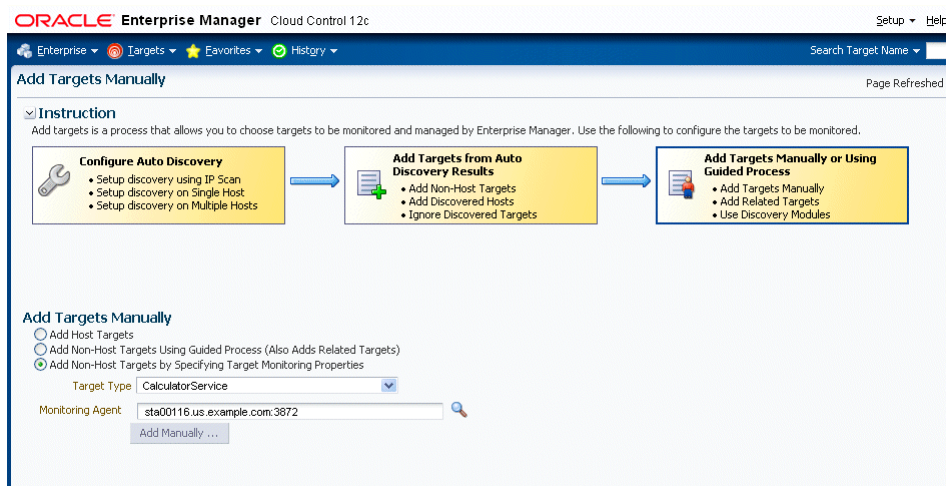
Once the plug-in has been deployed to the OMS, you are ready to add targets defined by your metadata plug-in to different monitoring Management Agents.

For illustrative purposes, the following steps show how to add the sample CalculatorService and TrafficLight as targets.

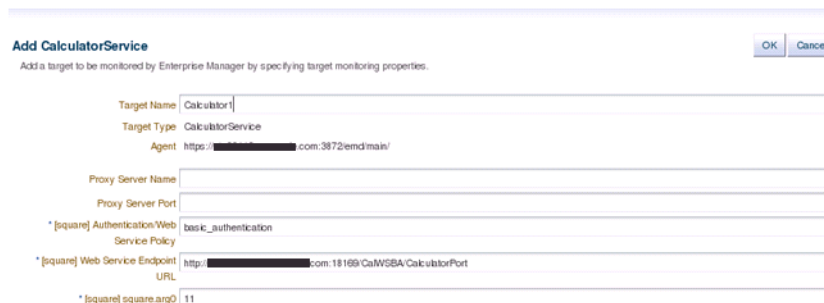
17.7.1 Adding a Web Services Target - CalculatorService

To add the CalculatorService target, perform the following steps:

1. From the **Setup** menu, select **Add Target** and then **Add Targets Manually**.
2. Select **Add Non-Host Targets by Specifying Target Monitoring Properties**.
3. From the **Target Type** menu, select **CalculatorService**.
4. From the **Monitoring Agent** menu, select the required monitoring Management Agent.



5. Click **Add Manually** to proceed.
6. Enter the property values of the target to be monitored.



- Click **OK** to complete the process. The confirmation window displays information on the newly added target.

17.7.2 Adding a WS-Management Target - TrafficLight

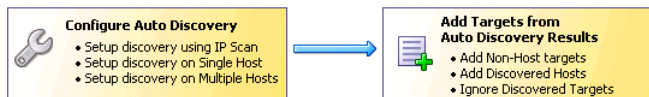
To add the sample TrafficLight target, perform the following steps:

- From the **Setup** menu, choose **Add Target** and then **Add Targets Manually**.
- Select **Add Non-Host Targets by Specifying Target Monitoring Properties**.
- Select **TrafficLight** from the **Target Type** drop-down menu.
- Select the desired agent from the **Monitoring Agent** drop-down menu.

Add Targets Manually

Instruction

Add targets is a process that allows you to choose targets to be monitored and managed by Enterprise Manager. Use



Add Targets Manually

- ☐ Add Non-Host Targets Using Guided Process (Also Adds Related Targets)
☒ Add Non-Host Targets by Specifying Target Monitoring Properties
☐ Add Host Targets

Target Type TrafficLight

Monitoring Agent com:3872

Add Manually ...

- Click **Add Manually** to proceed.
- Enter the property values of the target to be monitored.

- Click **OK** to complete the process. The confirmation window displays information on the newly added target.

17.7.3 Configuring a Standalone Java Application or JVM Target

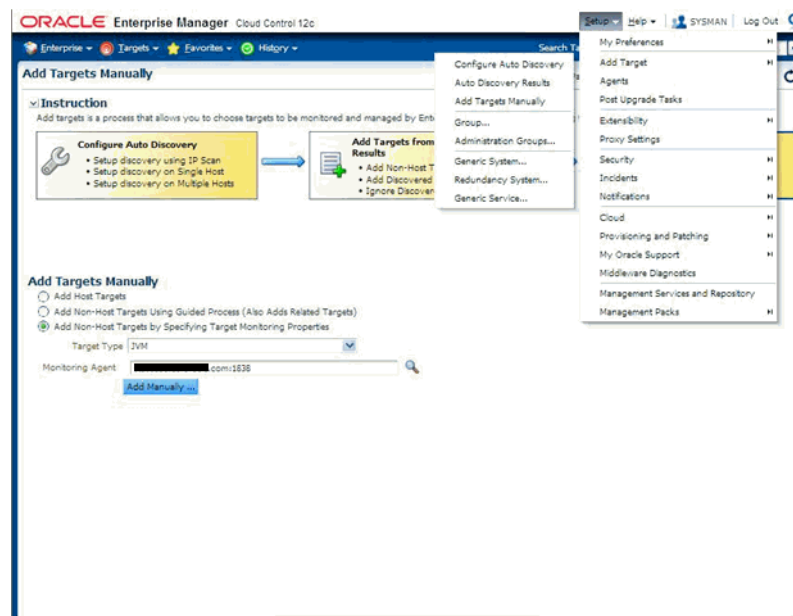
If you deployed a plug-in that defines a standalone Java application or you want to use the built-in JVM target type, you can begin configuring your JVM or JMX-enabled

Java application targets so that metrics for these targets can be collected in Enterprise Manager Cloud Control.

On the system running the JVM, install an Enterprise Manager Agent version 10.2.0.3 or later. Although recommended, this is not necessary for JVM and standalone Java application targets.

To add the JVM target instance, perform the following steps:

1. From the **Setup** menu of Enterprise Manager console (top right), select **Add Target** and then **Add Targets Manually**.
2. Select **Add Non-Host Targets by Specifying Target Monitoring Properties**.
3. From the **Target Type** menu, select **JVM**.
4. From the **Monitoring Agent** list, select the required Management Agent (preferably a Management Agent local to the JVM being monitored).



5. Enter the instance properties for this JVM or Java application instance that the Management Agent needs to monitor, then click **OK**.

ORACLE Enterprise Manager Cloud Control 12c Help

Add JVM
Add a target to be monitored by Enterprise Manager by specifying target monitoring properties.

Target Name: myJVMOnHost1
Target Type: JVM
Agent: https://[redacted]oracle.com:1838/emd/main/

* Machine name: adc2180736
* Admin Port number: 6789
User Name:
JVM Admin User Password:
Communication Protocol: rmi
Service Name: jmxrmi
SSLTrust Store:
SSLTrust Store Password:
Custom lookup provider Class:

OK Cancel

Table 17–1 provides definitions for the instance properties.

Table 17–1 JVM Instance Properties

Property	Definition
Name	Target name for this JVM instance.
MachineName	Host name where this JVM is running.
Admin Port Number	Port number a JSR-160 client can use (such as jconsole when using the “remote” option) to connect to the JVM. (This is the port specified for the <code>-Dcom.sun.management.jmxremote.port</code> property when the JVM is started up to enable remote management.)
User Name	Required if JVM started with: <code>Dcom.sun.management.jmxremote.authenticate=true</code> with a password and access file.
JVM Admin User Password	See the preceding User Name property.
Communication Protocol	Establishes a connection to the MBeanServer on the target JVM. This corresponds to the properties of the JMX ServiceURL needed to establish the JMX connection to the target MBeanServer. The default of <code>rmi</code> should be retained.
Service Name	Establishes a connection to the MBeanServer on the target JVM. This corresponds to the properties of the JMX ServiceURL needed to establish the JMX connection to the target MBeanServer. The default of <code>jmxrmi</code> should be kept.
SSL Trust Store	Location of the SSL Trust Store, which is needed if the target JVM has SSL enabled with <code>-Dcom.sun.management.jmxremote.ssl=true</code> on its startup.

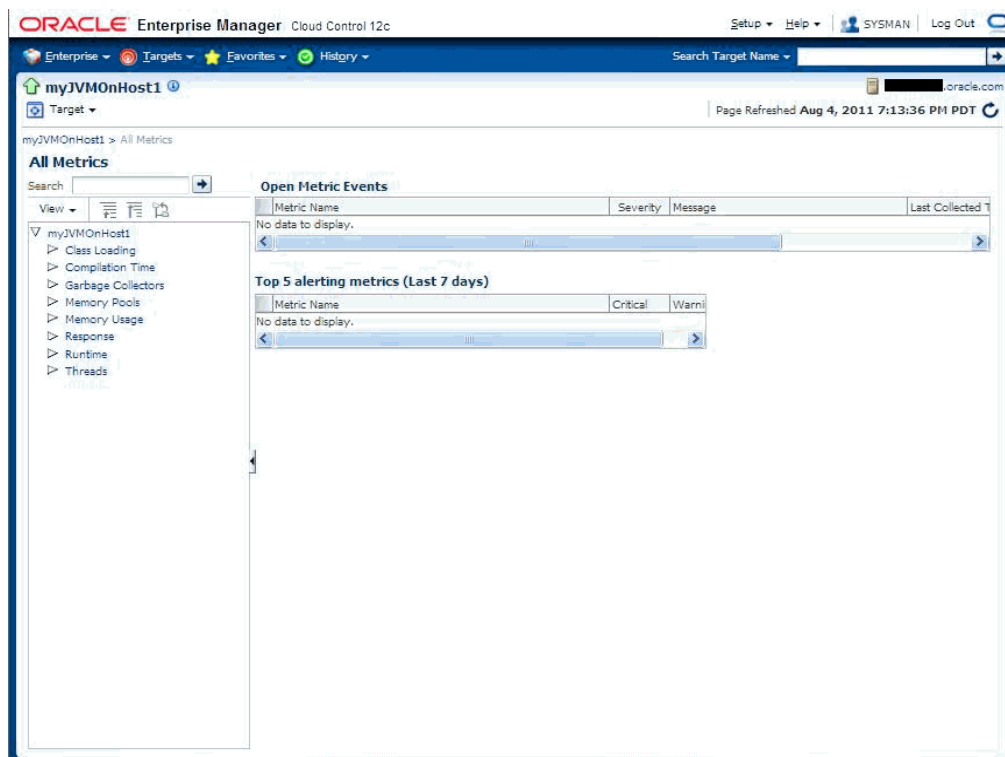
Table 17–1 (Cont.) JVM Instance Properties

Property	Definition
SSL Trust Store Password	Password needed to access the SSL Trust Store path.
Custom Lookup Provider Class	Full package name of a user-implemented Java lookup class that can be integrated into the Enterprise Manager client and be used to perform a custom lookup of the MBeanServer through LDAP or other lookup protocols.

- Navigate to the All Metrics page of the added JVM (Java application) target to see the metrics collected from the JVM (Java application) to Enterprise Manager. These metrics are exposed by the platform MBeans, which is available on JDK1.5 or above, or from application-defined MBeans for your Java application.

To navigate to JVM target home page from the **Targets** menu, choose **All Targets** and then select your JVM target instance.

To navigate to the **All Metrics** page, from the **Target** menu, select **Monitoring** and then **All Metrics** from the JVM target's home page menu.



The following graphic shows the collected metric details.

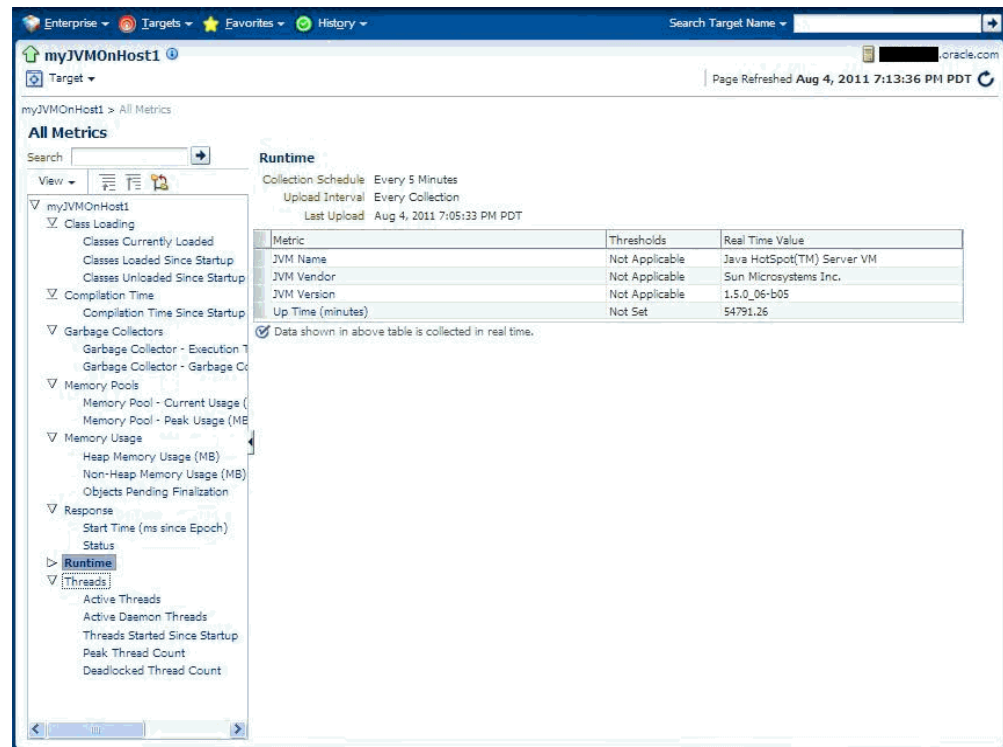


Table 17–2 Properties the Fetchlet Uses

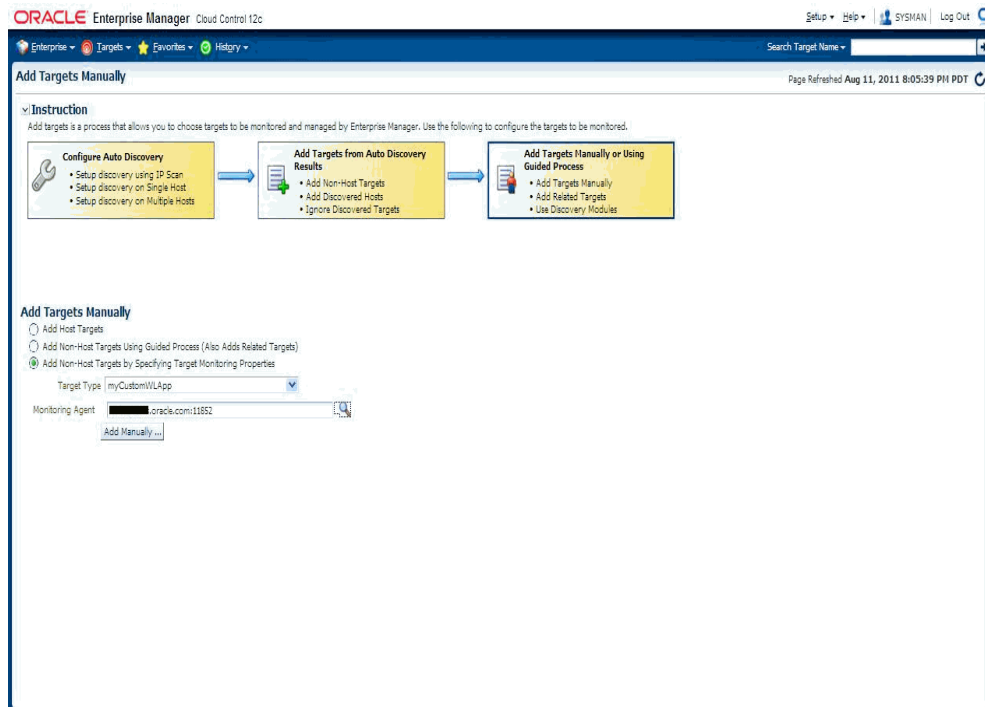
Property	Default	Description
MachineName	localhost	MBean server host machine name.
Port	8888	Port on which the MBean server is listening for connections.
Username	null	User name if required for a connection.
Password	null	Password if required for a connection.
Protocol	rmi	Protocol used for the connection.
Service	jmxrmi	Service used for the connection.
SSLTrustStore	null	Path to the SSLTrustStore.
SSLTrustStorePassword	null	Password needed to access the SSLTrustStore path.

17.7.4 Adding a Target Instance for a Custom J2EE Application on WebLogic

You have a custom J2EE application on WebLogic from which you need to collect custom metrics into Enterprise Manager that are exposed via JMX Mbeans. Once you have defined and deployed a plug-in that defines your custom target type, you can begin configuring your JMX-enabled J2EE application target instances on the various Management Agents to where you deployed the plug-ins. This enables Enterprise Manager to collect metrics for these target instances.

1. From the **Setup** menu, select **Add Target** and then **Add Target Manually**. Select the **Add non-Host targets by specifying Target Monitoring Properties** option.
2. Select your custom target type created earlier and deployed to the OMS

3. Select the monitoring Management Agent where you want to create an instance of this target type (this should preferably be an emagent local to the target)



4. Click **Add Manually**.
5. Enter the requisite target properties, as shown in the following graphic, then click **OK**. The newly added target appears in the "All Targets" list.

The screenshot shows the 'Add myCustomWLApp' dialog box. It contains the following fields and values: 'Target Name' is 'myCustomWLApp1', 'Target Type' is 'myCustomWLApp', 'Agent' is 'https://oracle.com:11552/and/main/', 'Machine Name' is '[redacted]', 'Admin Port number' is '27021', 'User Name' is 'weblogic', 'J2EE Admin User Password' is '*****', 'Communication Protocol' is 'G', 'Service Name' is 'weblogic.management.mbeanservers.runtime', 'Metric Source' is 'WebLogic', and 'Custom lookup provider Class' is empty. The 'OK' button is in the top right corner.

Table 17–3 Target Properties

Property	Definition
Name	Unique name for this target instance.
MachineName	Host name/IP Address of the system running the 9.x version or later of the Oracle WebLogic Application Server.

Table 17-3 (Cont.) Target Properties

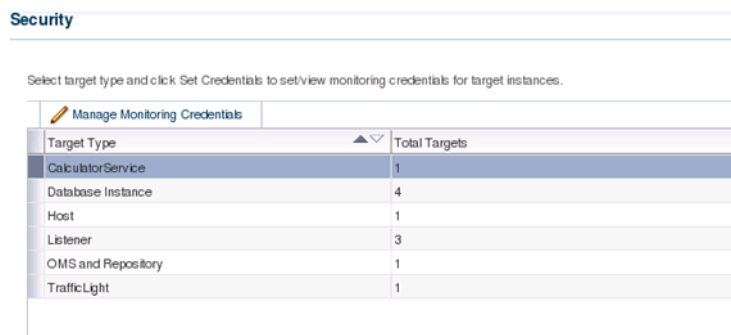
Property	Definition
Username	User Name used to establish the JMX connection to the WebLogic server. This could be either an administrator or monitor user.
JVM Admin User Password	Password for preceding user.
Communication Protocol	t3 (default) or t3s.
Service Name	<i>weblogic.management.mbeanservers.runtime</i> (or other MbeanServer where the application registers its Mbeans).
Metric Source	WebLogic

The metrics created can be viewed by navigating to the target instance home page and navigating to the **All Metrics** page (from the **Target** menu, choose **Monitoring** and then **All Metrics**).

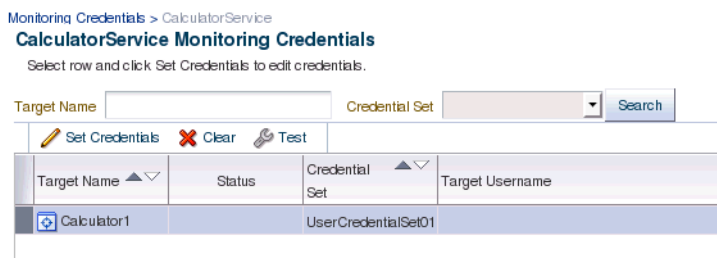
17.8 Monitoring Credential Setup

Some target types require monitoring credentials to be set for target instances. In the demo plug-ins, both CalculatorService and TrafficLight require monitoring credentials. The following steps demonstrate how to set up the credentials:

1. From the **Setup** menu, select **Security** and then **Monitoring Credentials**.
2. Select **CalculatorService** and then click **Manage Monitoring Credentials**.



3. Select **Calculator1** and then click **Set Credentials**.



4. Select **AliasCredential** from **Credential Type**. Enter values for **Alias** and **Password**.

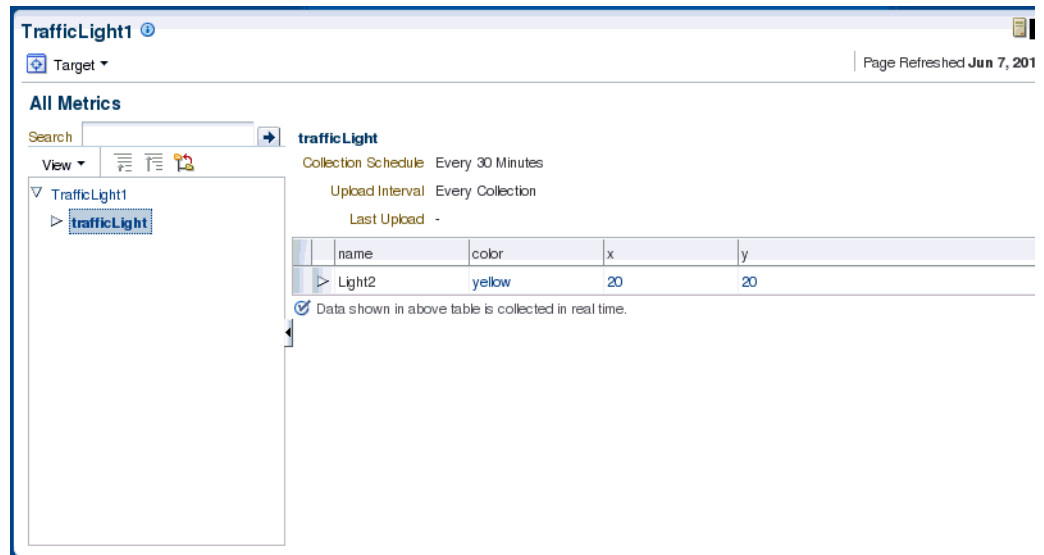
5. Click **Save** to finish.
6. Repeat the above steps for the target **TrafficLight1**.

17.9 Viewing Monitored Metrics

With a target instance added to the Management Agent for monitoring, you can now view metrics defined for your target type. As before, the sample targets are used to illustrate the procedure.

1. From the **All Targets** page, click the target you added in the previous step. Enterprise Manager takes you to that target's home page.
2. From the **Target** menu, select **Monitoring** and then **All Metrics**. The **All Metrics** page appears for the monitored target. An expandable tree list for each metric enables you to drill down to view specific metric parameters, as shown below:

Metric	Thresholds	Real Time Value
SquareResult	Not Applicable	121



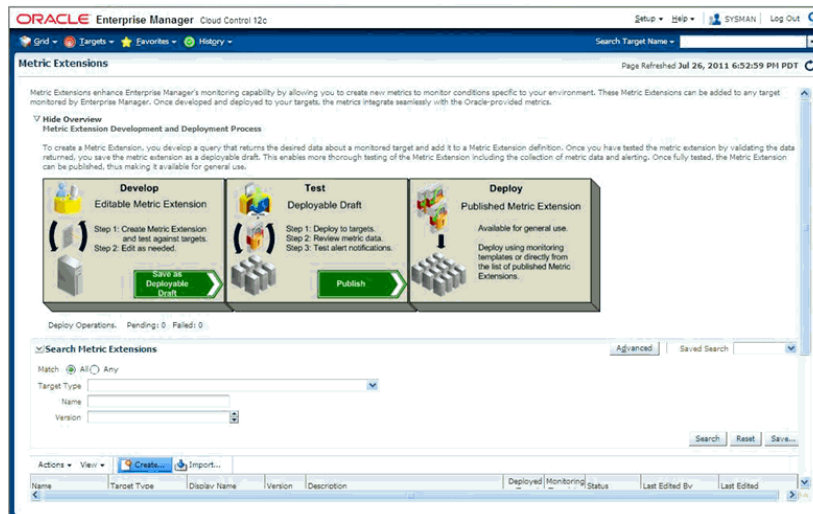
17.10 Creating JMX Metric Extensions

If you wish to collect metrics from your custom J2EE application deployed on Oracle Fusion middleware and exposed via JMX attributes into Enterprise Manager 12c, you can use either the Enterprise Manager console or the `jmxcli` command line tool. The latter also supports defining Metric Extensions from JMX operations and supports the creation of a Metric Extension Archive (MEA) which then must be imported into the OMS via the console and then tested and deployed to the desired J2EE application target instances representing your custom application.

Note: While you can select attributes that are not open types using the Mbean browser, the JMX metric extension UI supports only open type attributes. An error will occur if the UI is used to create metrics by selecting attributes which are not open types.

17.10.1 Using the Enterprise Manager Console

1. From the **Enterprise** menu, select **Monitoring** and the **Metric Extensions**. The Metric Extension page displays.
2. Click **Create** to create a Metric Extension.



3. Select "Application Deployment" target type (or any other appropriate Enterprise Manager target type for which this metric needs to be defined) and specify a meaningful name for your metric extension. Keep in mind that you might eventually end up creating additional metric extensions on the "Application Deployment" target type both for this application and for other custom applications so it is desirable to capture both the metric name and the application name in the metric extension name, whenever possible.

Also select JMX for the Adapter.

Note the "Collection Schedule" section below the "General Properties" section. This is where you define how often this metric is to be collected, or if this is realtime-only metric (in which case the **Disabled** button should be selected.).

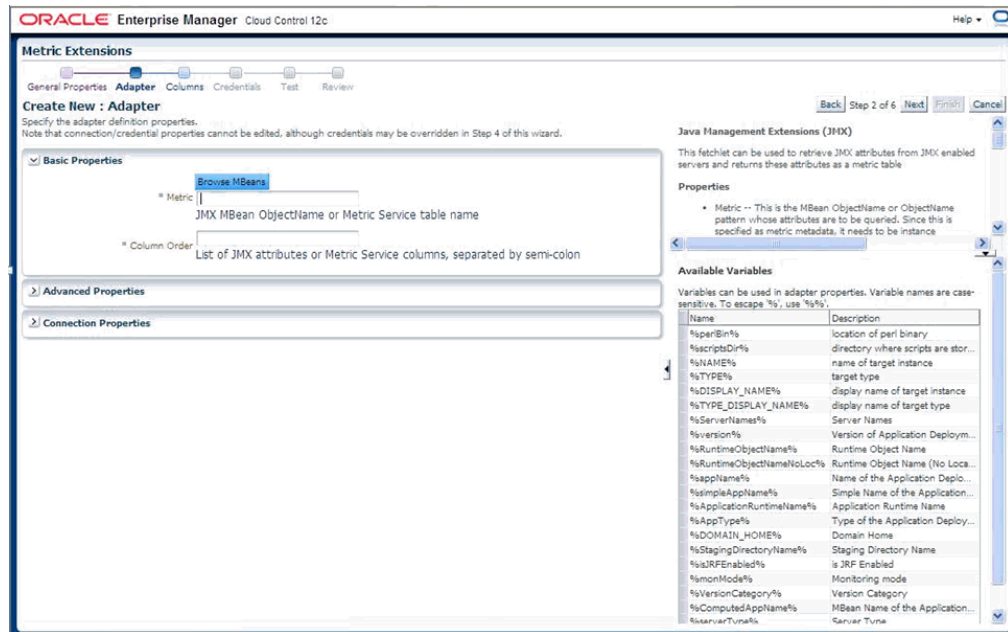
If "Alerting and Historical Trending" is selected, you can also select an **Upload Interval**, which indicates which samples (whose frequency is specified in the "Repeat Every" field) are uploaded to the Enterprise Manager repository for historical trending. For example if Collection frequency is specified as 15 minutes and the Upload Interval is 3, then every 3rd sample will be uploaded into the repository (every 45 minutes) and will be available for historical trending. However "alerts" that are possibly triggered due to threshold violations will be available for every collection (15 minutes).

4. Click **Next** and specify the required properties needed for a JMX-based metric. These are defined in the **Basic Properties** section and are:

- Metric: The Mbean ObjectName or Pattern and
- Column Order: A semi-colon separated list of JMX attributes for above Mbean (if a metric needs to be defined using a JMX operation, use the `jmxcli` as shown in a following section)

Note that the Mbean ObjectName or pattern defined previously must not have any server-specific key properties defined. These properties may be replaced with a wildcard ("*").

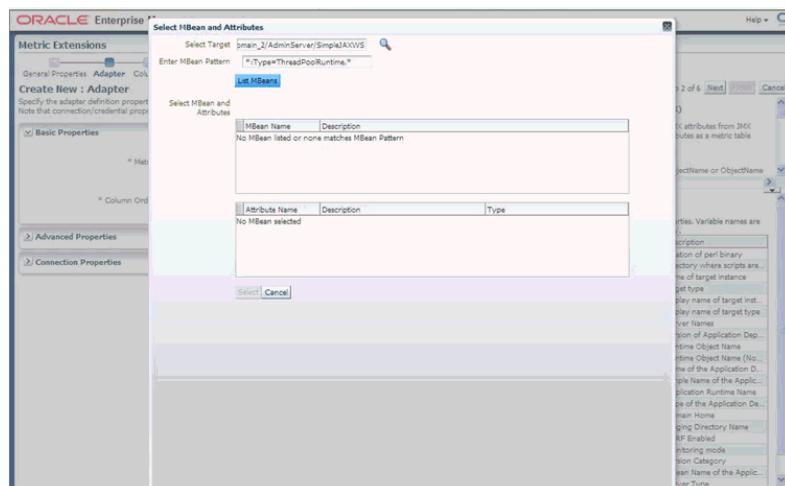
For example, if an Mbean object name is `com.bea:Type=foo,Location=Server1,Name=abc` then it may be appropriate to define this as `com.bea:Type=foo,*` in the "Metric" property described above. Also, if the Mbean ObjectName is a pattern, please be aware that multiple Mbeans could be returned making this metric a "table" with multiple rows (each row representing the JMX attributes of an Mbean matching the ObjectName pattern). In this case we need to define at least one or more columns as Key columns so that each row is unique in the resultant metric.



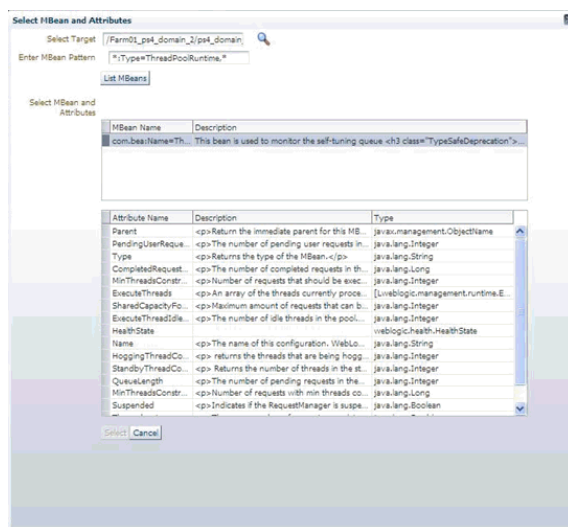
For the preceding step, there is a **Browse Mbeans** button that makes it easier to configure these two properties by allowing you to browse an MbeanServer and selecting an Mbean and its JMX attributes that need to be represented by this metric being defined in the metric extension.

If you click **Browse Mbeans**, you must perform the following in sequence.

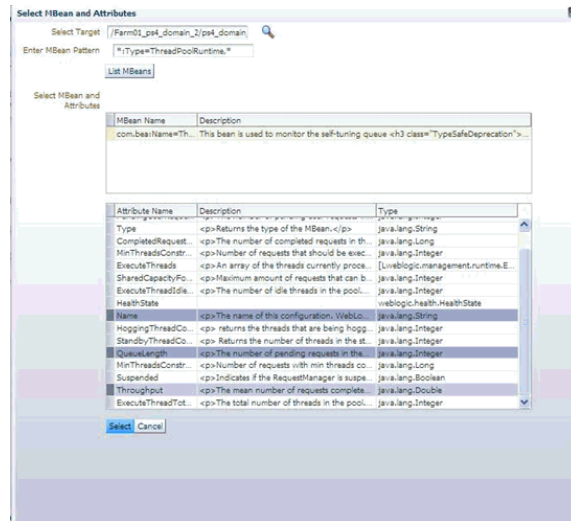
- Select the Target: Select an instance of the target type that you need to use to define this metric. This target instance is used to help configure the metric and does not have to be the target instance on which the metric is eventually defined.
- Enter the Mbean Pattern: Here, you enter an Mbean Object Name or pattern for the Mbean you are interested in monitoring
- Click **List Mbeans**: This will be displayed in the table under "Select Mbean and Attributes", the Mbeans that match the Mbean pattern or the text "No Mbean listed or none matches Mbean Pattern" if there is no match. You can iteratively update the previous "Enter Mbean pattern" field and click **List Mbeans** to refine the list of Mbeans displayed.



- Select an Mbean of interest: This will automatically populate the table below with the JMX attributes for the selected Mbean.

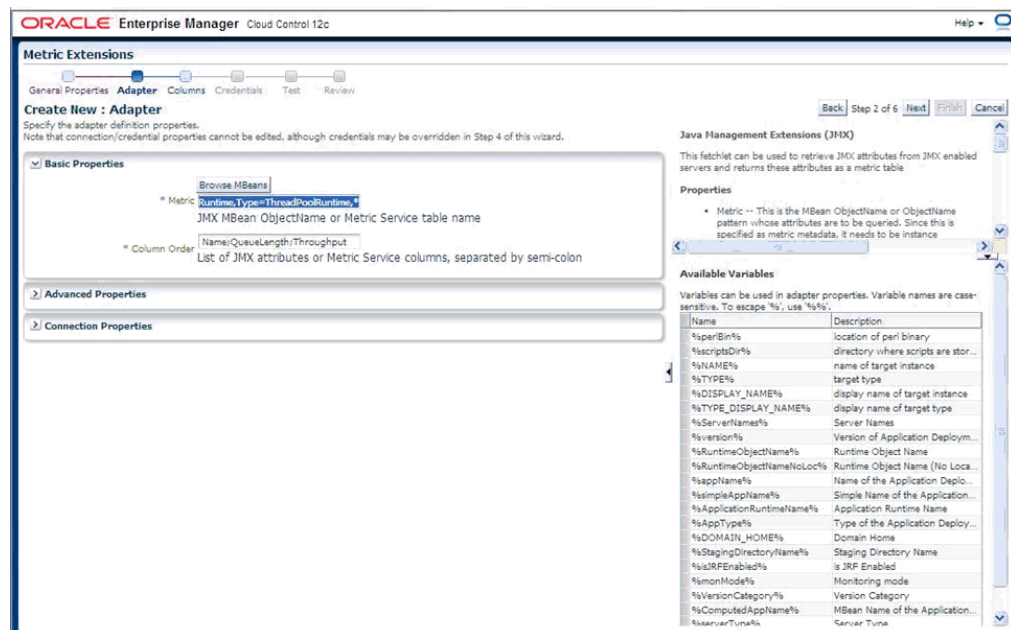


You can multi-select (using Control + click) multiple attributes and then click **Select** to accept the selections.



- You must now specify the required parameters "Metric" and "Column Order" needed to define a JMX based metric extension.

Note that the Mbean name populated in the "Metric" field should not have any instance specific information in its key properties (like Location=Server1 or ServerName=foo) if this metric extension can be applied to multiple servers besides the one that was selected/used to configure the metric extension using the "Browse Mbean" wizard above. These instance-specific key properties could be replaced with a wildcard "*" as appropriate to make this a valid Mbean ObjectName pattern.



Explanation of Specifiable Properties

Required Properties:

- metric -- This is the MBean ObjectName or ObjectName pattern whose attributes are to be queried. Since this is specified as metric metadata, it needs

to be instance agnostic so instance specific key-properties if any (like servername), on the MBean ObjectName may need to be replaced with wildcards.

- `columnOrder` -- This is a semi colon separated list of JMX attributes in the order they need to be presented in the metric

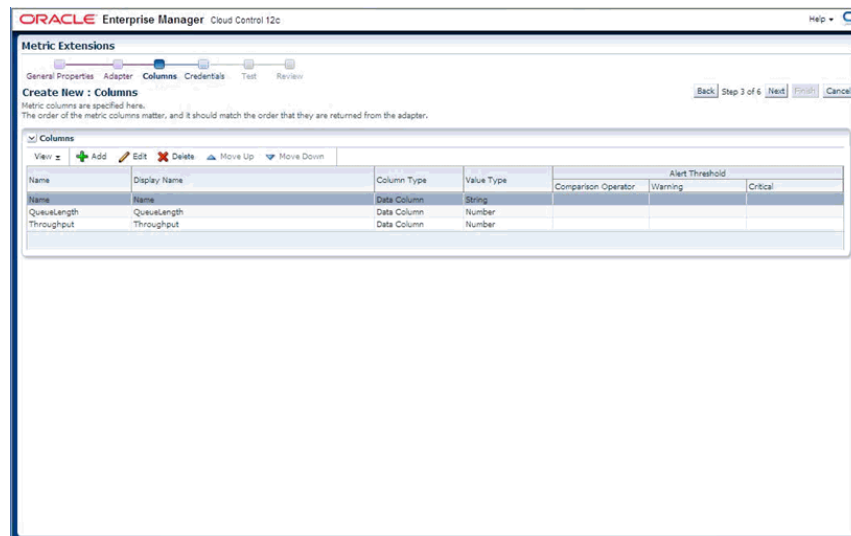
Advance Properties:

- `identityCol` -- This is an MBean key property that needs to be surfaced as a column when it is not available as a JMX attribute. For example: `com.myCompany:Name=myName,Dept=deptName,prop1=prop1Val,prop2=prop2Val` In above case setting `identityCol` as `Name;Dept` (note that separator is a semi-colon) will result in two additional key columns representing `Name` and `Dept` besides the columns representing the JMX attributes specified in the `columnOrder` property above.
- `autoRowId` -- This is the prefix used for an automatically generated row in case the MBean ObjectName pattern specified in metric property matches multiple MBeans and none of the JMX attributes specified in the `columnOrder` are unique for each. The `autoRowId` value specified here will be used as a prefix for the additional key column created. For example, if the metric is defined as `com.myCompany:Type=CustomerOrder,*` `columnOrder` is `CustomerName;OrderNumber;DateShipped` (and assuming `customerName;OrderNumber;DateShipped` may not be unique if an order is shipped in two parts).

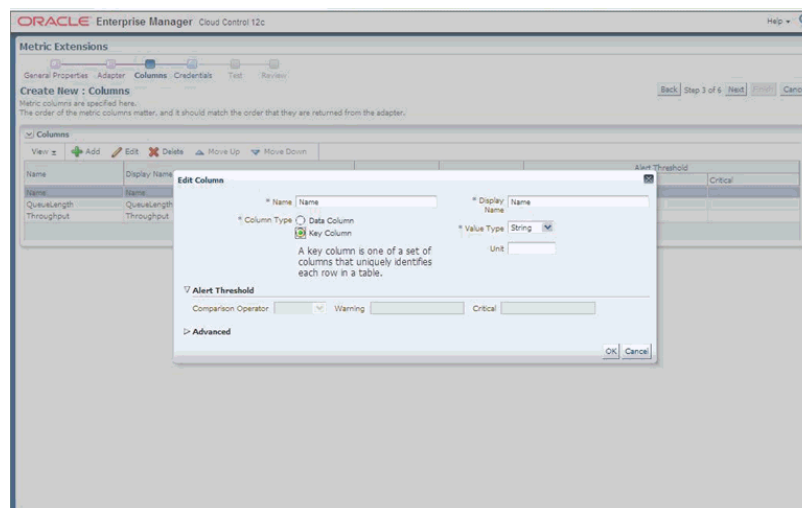
Setting `autoRowId` as `"ShipItem-"` will populate an additional key column for the metric for each row with `ShipItem-0`, `ShipItem-1`.

- `MetricService` -- True/False indicates whether `MetricService` is enabled on the target WebLogic domain. This would be unchecked or false in most cases for user-defined metrics except when metrics that are exposed via the Oracle DMS MBean needs to be collected. If set to true, then the basic property `"metric"` above should represent the `MetricService` table name and the basic property `"columnOrder"` will represent a semicolon separated list of column names for aforementioned `MetricService` table.
6. Specify the Columns for this metric (if you have used the "Browse Mbeans" step earlier, then these columns are automatically pre-filled for you). You may need to edit these pre-created columns by the "Browse Mbean" wizard to specify columns that are "Key columns". This done in the event an Mbean pattern is specified in the previous step for the "Metric" property, and multiple Mbeans could match this Mbean pattern for any of the target instances to which this metric extension will be applied to.

If the order of the columns are changed (using Move Up - Move Down buttons) then the corresponding order of the semi-colon separated columns in the "Column Order" property in the previous step also needs to be updated accordingly (using **Back**).

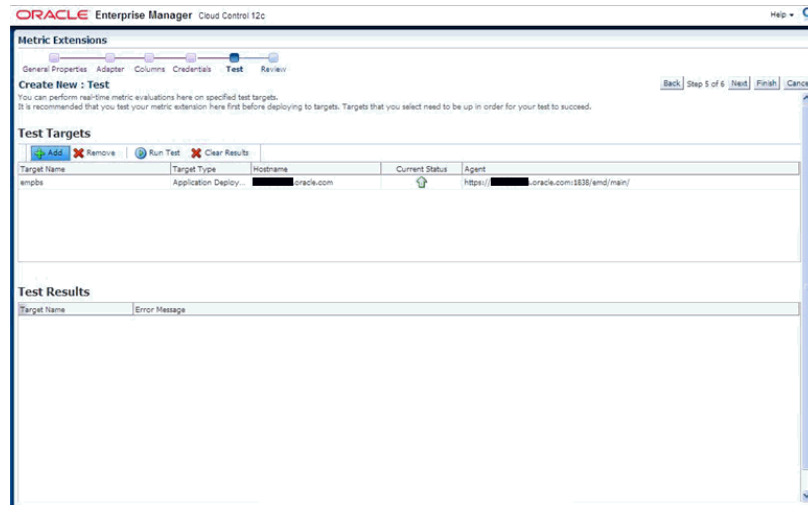


If needed, edit the columns as desired to make them a Key Column as shown in the following graphic.



Once columns are labeled and edited, click **Next**. We are now ready to test the metric extension

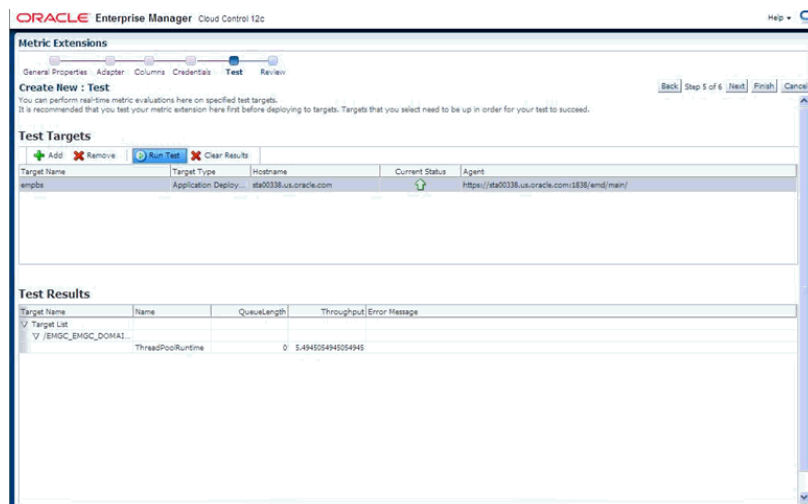
7. Click **Add** to select a target instance on which to test this metric extension. This could preferably be a different target instance than the one used to define the metric extension (if the **Browse Mbeans** button was used to help in defining the metric extension earlier).



Now select a target instance in the Test Targets table and then click **Run Test** above that table.

The metric values are displayed in the Test Results table (if there are errors ,then those are also shown).

If errors are present, click **Back** and fix the errors and re-run the test.



- Once satisfied with the Test, click **Next** to view a summary of the metric extension and then click the **Finish** to define the metric extension.

Oracle Enterprise Manager Cloud Control 12c

Metric Extensions

General Properties Adapter Columns Credentials Test Review

Create New: Review

Creating version 1 of metric extension ME\$myCRMAppME1 for target type j2ee_application

General Properties

Target Type: Application Deployment
 Name: ME\$myCRMAppME1
 Display Name: myCRMAppME1
 Adapter: Java Management Extensions (JMX)
 Description:
 Version Comment:

Collection Schedule

Data Collection: Enabled
 Collection Frequency: By Minutes
 Repeat Every: 15 Minutes
 Use of Metric Data: Alerting and Historical Trending
 Upload Interval: 1

Adapter Properties

Metric: com.bea.name.ThreadPoolRuntime.Type.ThreadPoolRuntime.*
 Column Order: Name;QueueLength;Throughput

Advanced Properties

Custom Files

Columns

Name	Display Name	Column Type	Value Type	Compute Expression	Comparison Operator	Warning	Critical
Name	Name	Key Column	String				
QueueLength	QueueLength	Data Column	Number				
Throughput	Throughput	Data Column	Number				

Credentials

Testing Status

- Before deploying the metric extension to selected target instances the metric extension needs to be saved as a "Deployable draft". This will let the metric extension designer deploy the metric extension to selected target instances and verify the metric collection but will prevent other administrators from deploying this metric extension until after it has been tested and the designer is satisfied.

Oracle Enterprise Manager Cloud Control 12c

Metric Extensions

Confirmation

Metric Extension: myCRMAppME1 v1 (ME\$myCRMAppME1) successfully created.

Editable Metric Extension

Step 1: Create Metric Extension and test targets.

Deployable Draft

Step 1: Deploy to targets.
 Step 2: Review metric data.
 Step 3: Test alert notifications.

Published Metric Extension

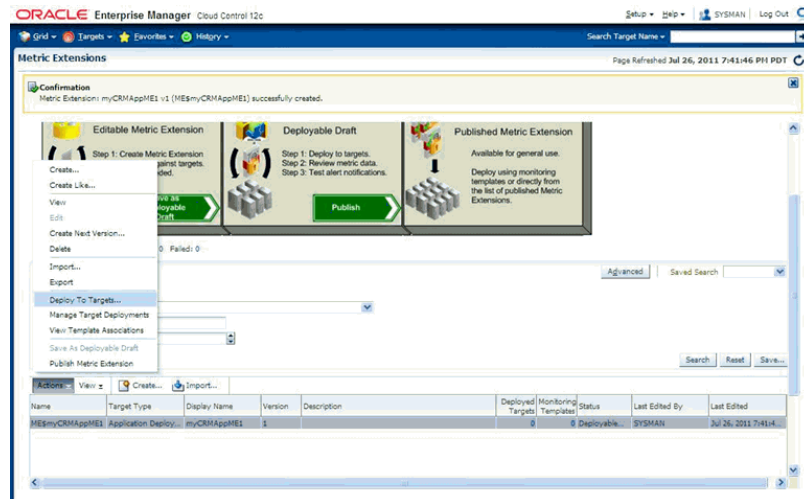
Available for general use.
 Deploy using monitoring templates or directly from the list of published Metric Extensions.

Actions

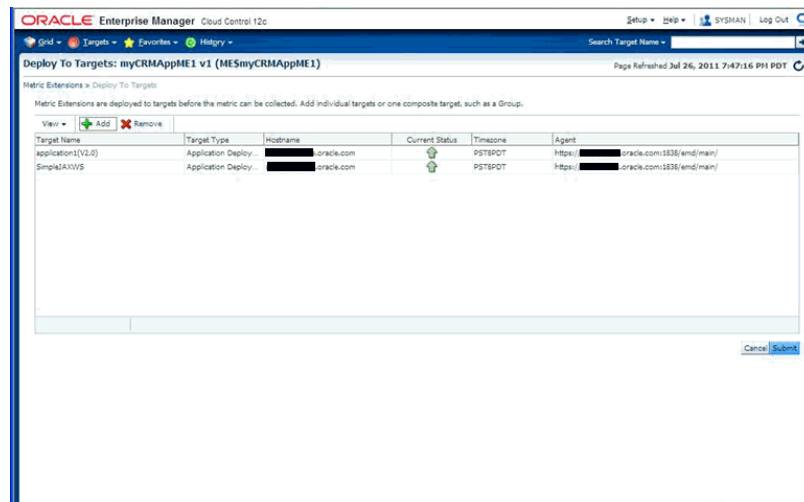
Create...
 Create Like...
 View...
 Edit...
 Create Next Version...
 Delete...
 Import...
 Export...
 Deploy To Targets...
 Manage Target Deployments...
 View Template Associations...
 Save As Deployable Draft...
 Publish Metric Extension...

Name	Target Type	Display Name	Version	Description	Deployed Targets	Monitoring Template	Status	Last Edited By	Last Edited
ME\$myCRMAppME1	Application Deployment	myCRMAppME1	1		0	0	Editable	SYSMAN	Jul 26, 2011 7:43:14

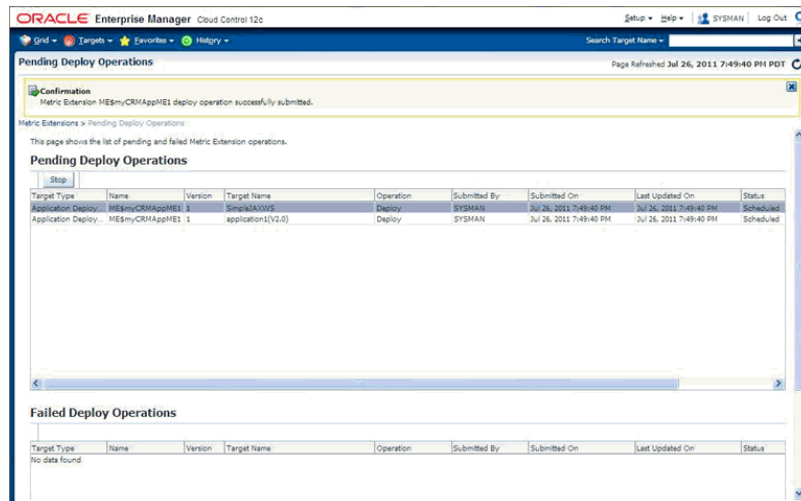
- Select the metric extension just created and saved as a deployable draft. From the **Actions** menu, choose **Deploy to Targets**.



11. Select the target instances that this metric extension needs to be tested on and click **Submit**. For example, if the metric extension is defined on an "Application Deployment" target type and represents a metric from a Custom Mbean registered by a custom JEE application, the instances of that custom application could be selected. This will schedule a job to asynchronously deploy the metric extension to the Management Agents monitoring the selected targets.

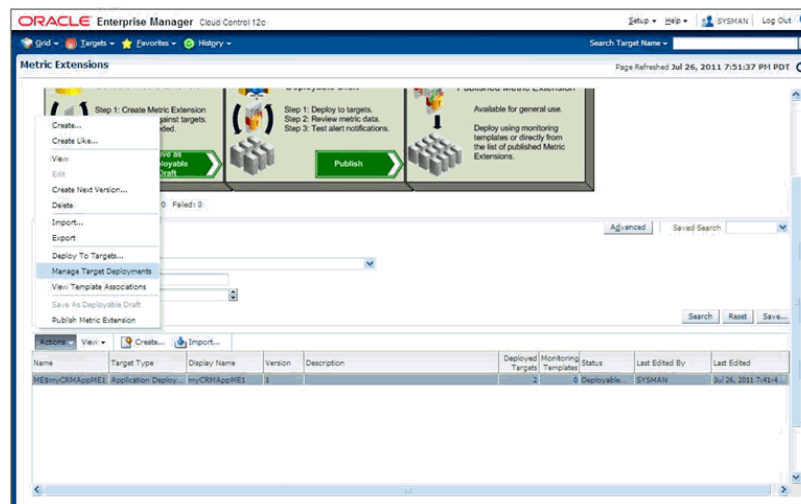


12. Monitor the status of the Pending deploy operation of the metric extension to selected targets by refreshing this page periodically to monitor the Status column and Failed Deploy Operations table for any possible errors during deployment.

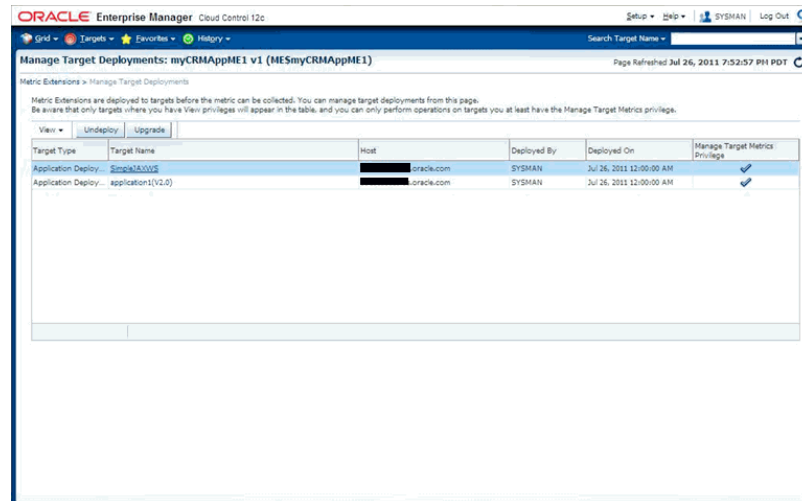


13. From the **Enterprise** menu, select **Monitoring** and then **Metric Extension**. On the Metric Extension home page, your metric extension appears as a row in the table with a column "Deployed Targets" representing the count of the number of targets this metric extension is deployed to.

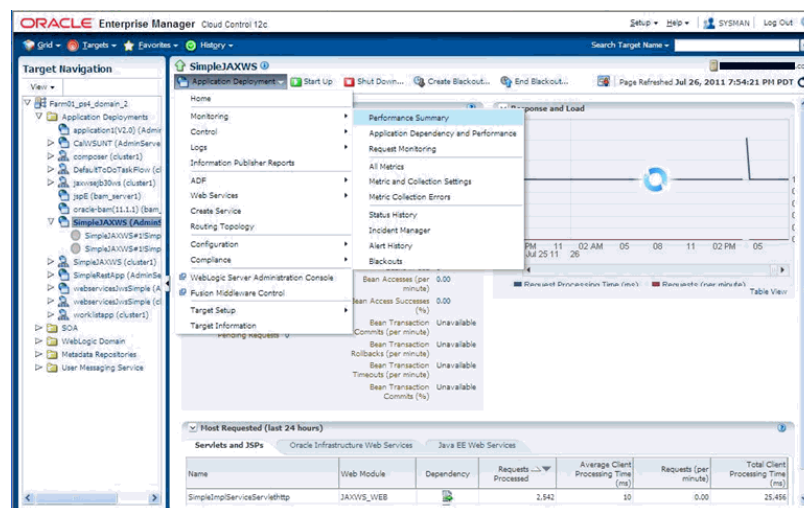
From the **Actions** menu, choose **Manage Target Deployments** from the table after selecting the desired metric extension. This will list the target instances this metric extension is deployed to.



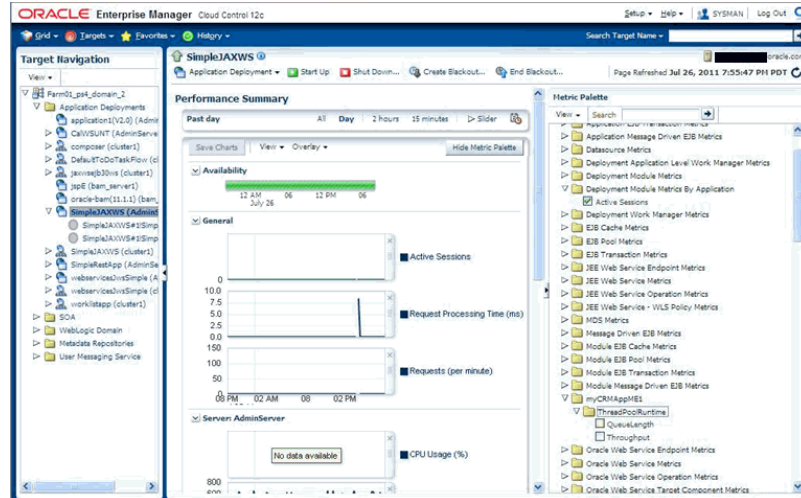
14. Click on the value in the "Target Name" column for the target instance you want to verify the metric extension on. This takes you to the home page of the target.



15. For middleware targets, navigate to the *Target Type/Monitoring/ Performance Summary* (or in general to the *Target Type /Monitoring /All Metrics*) page.



16. From the Performance Summary page, the newly created metric will be visible on the Metric Palette and can be selected and charted on the page.



17. Once satisfied with testing the metric extension on one or more target instances, the metric extension can be published from the Metric Extension page (from the Actions menu, choose Publish Metric Extensions) and then deployed to remaining target instances.

17.10.2 Using the JMXCLI to create a Metric Extension Archive

If you do not wish to use the Enterprise Manager console (or do not want to surface an Enterprise Manager metric exposed via a JMX operation), you can use the command line tool JMXCLI to create a Metric Extension Archive. This can then be imported into the OMS, edited, tested, published and then deployed to required instances of the target type on which it is defined. The following illustrates the use of jmxcli in creating a Metric Extension archive.

1. `cd Agent_Instance_Home/bin`
2. `setenv USER_JARS $T_WORK/middleware/wlserver_10.3/server/lib/weblogic.jar` (this should not be necessary if your Mbeans just return JMX Open Types and not any custom classes).
3. `emctl jmxcli -t WebLogic -MEXT -l "service:jmx:t3://sta00338:7018/jndi/weblogic.management.mbeanservers.runtime" -u weblogic -c welcome1 -m "*:Type=ThreadPoolRuntime,*" -w /scratch/TEMP/`

Options:

-l : JMX serviceURL to connect to the WebLogic server. Replace the host:port above with what is appropriate for your instance

-u : WebLogic user having access to required MBeans

-c : Password for the WebLogic user

-m : Mbean ObjectName or pattern.

-w : Temporary work directory where the Metric Extension Archive (which can later be imported into the OMS console) is created.

Oracle Enterprise Manager 12c Release 1 Cloud Control 12.1.0.0.0
Copyright (c) 1996, 2011 Oracle Corporation. All rights reserved.
Using Plugin Root /ade/sparnesw_
egc802/oracle/emagent/gcagent/plugins/oracle.sysman.emas.agent.plugin_

```

12.1.0.0.0
Connecting to server:
service:jmx:t3://sta00338:7018/jndi/weblogic.management.mbeanservers.runtime
Connecting as user: weblogic
Obtained 1 MBeans matching pattern *:Type=ThreadPoolRuntime,*.
Enter an existing target type for this Metric Extension: [j2ee_application]
Enter the name of the Metric Extension: [myMEXT] myAppME_1
Enter the Metric Extension version: [1.0]
Enter the Metric Extension metadata file location: [./metadata/ME#24#myAppME_
1.xml]
Enter the Metric Extension collection file location:
[./collection/ME#24#myAppME_1.xml]
Enter a label for this Metric Extension: [myAppME_1]
Enter a description for this Metric Extension: [myAppME_1]
The available targets are:
0: This bean is used to monitor the self-tuning queue <h3
class="TypeSafeDeprecation">Deprecation of MBeanHome and Type-Safe
Interfaces</h3>
<p class="TypeSafeDeprecation">This is a type-safe interface for a WebLogic
Server MBean,
which you can import into your client classes and access through
<code>weblogic.management.MBeanHome</code>.
As of 9.0, the <code>MBeanHome</code> interface and all type-safe interfaces
for WebLogic Server MBeans are deprecated.
Instead, client classes that interact with WebLogic Server MBeans should use
standard JMX design patterns in which clients use the
<code>javax.management.MBeanServerConnection</code> interface to discover
MBeans, attributes, and attribute types at runtime.
For more information, see "Developing Manageable Applications with JMX" on <a
href="http://www.oracle.com/technology/products/weblogic/index.html"
shape="rect">http://www.oracle.com/technology/products/weblogic/index.html</a>.
</p>
(com.bea:Name=ThreadPoolRuntime,ServerRuntime=EMGC_
ADMINSERVER,Type=ThreadPoolRuntime)
Enter the index of target/MBean you wish to monitor or press <Ctrl-C> to quit:
0
Following metric source types are available for selected target(s):
0: JMX Attributes
1: JMX Operations
Enter the index of your choice or press <Ctrl-C> to quit: 0
Attributes are:
0: CompletedRequestCount Return Value: java.lang.Long
1: ExecuteThreadIdleCount Return Value: java.lang.Integer
2: ExecuteThreads Return Value:
[Lweblogic.management.runtime.ExecuteThread;
3: ExecuteThreadTotalCount Return Value: java.lang.Integer
4: HealthState Return Value: weblogic.health.HealthState
5: HoggingThreadCount Return Value: java.lang.Integer
6: MinThreadsConstraintsCompleted Return Value: java.lang.Long
7: MinThreadsConstraintsPending Return Value: java.lang.Integer
8: Name Return Value: java.lang.String
9: Parent Return Value: javax.management.ObjectName
10: PendingUserRequestCount Return Value: java.lang.Integer
11: QueueLength Return Value: java.lang.Integer
12: SharedCapacityForWorkManagers Return Value:
java.lang.Integer
13: StandbyThreadCount Return Value: java.lang.Integer
14: Suspended Return Value: java.lang.Boolean
15: Throughput Return Value: java.lang.Double
16: Type Return Value: java.lang.String

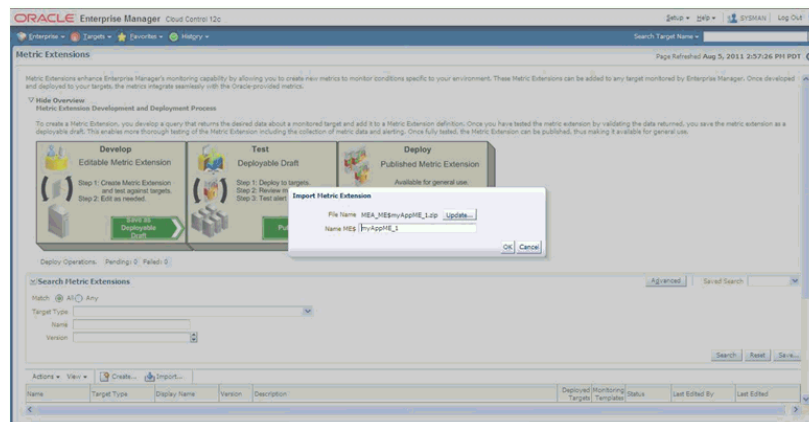
```

```
Select one or more items as comma separated indices: 5,13
Number of possible columns in the resultant metric are 2.
Enter the name for this metric column at index=0 : [HoggingThreadCount]
Is this column a KEY Column <y/n>? [n]
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [HoggingThreadCount]
Enter the NLSID for column: [HoggingThreadCount]
Enter the UNIT for column "HoggingThreadCount": [millisec, kb etc.. ]
Do you want to create a threshold for this column <y/n>? [n]
Enter the name for this metric column at index=1 : [StandbyThreadCount]
Is this column a KEY Column <y/n>? [n]
Is this column for SUMMARY_UI <y/n>? [n]
Enter the label for column: [StandbyThreadCount]
Enter the NLSID for column: [StandbyThreadCount]
Enter the UNIT for column "StandbyThreadCount": [millisec, kb etc.. ]
Do you want to create a threshold for this column <y/n>? [n]
Do you want periodic collection for this metric <y/n>? [n] y
Enter the collection interval in seconds: 300
Periodic collection interval is: 300 seconds.
Written the metadata xml file: ./metadata/ME#24#myAppME_1.xml.
Creating new file: ./collection/ME#24#myAppME_1.xml.
Updated the default collection file for j2ee_application at location
./collection/ME#24#myAppME_1.xml.
createMextArchive: Adding metadata
createMextArchive: Adding collection file
createMextArchive: Adding mea.xml file

Creating Metric Extension zip archive: ./MEA_ME$myAppME_1.zip
Please import this into Enterprise Manager Cloud Control using the console.
Exiting...
```

The previous session creates a ZIP file MEA_ME\$myAppME_1.zip in the directory specified by the -w option when `jmxcli` is invoked (or in current directory if -w is not specified).

Import this into the Enterprise Manager console as shown below. From the **Enterprise** menu, choose **Monitoring** and then **Metric Extensions** to access the Metric Extensions home page.



After the Management Extension Archive is imported as shown in the preceding example, it can be edited (and modified), tested, published and deployed.

17.11 Surfacing Metrics from a Standalone JVM or Oracle Coherence

Users can also use the mechanism outlined in previous section to create additional metrics that are not available out-of- box for Oracle Coherence or JVM targets and the data for which are available via JMX Mbean attributes.

17.11.1 Using the Enterprise Manager Console

The procedure is similar to the ones followed in previous section for extending metrics on j2ee_application target types except that you must select target type "JVM" or "Oracle Coherence xxx" in Step 3 for defining the Metric Extension on JVM or Oracle Coherence target types.

17.11.2 Using JMXCLI

The steps are similar to those for using JMXCLI to define a Metric Extension Archive for custom J2EE applications except that the start-up arguments when `jmxcli` is invoked as follows:

```
emctl jmxcli -t JVM -MEXT -h adc2180736 -p 6789 -m "*:*" -w /scratch/TEMP/
```

You must specify target type on which the Metric Extension is defined to be JVM or oracle_coherence as appropriate (instead of the default j2ee_application).

Using the Web Services Framework

In a service based architecture, Web services are Web-based applications that use XML standards and transport protocols to exchange data with clients. Web Service Definition Language (WSDL) is used to describe the available services and Simple Object Access Protocol (SOAP) is used to transfer the data in a request and response model over HTTP. XML is used to tag the data. These standard interfaces allow the Web service functionality to be exposed without the need for the client to know more specific information on how the Web service is implemented.

This chapter contains the following sections:

- [Using APIs to Write Java Clients](#)
- [About SOAPMessageBuilder](#)
- [About WebServiceInvoker](#)
- [Using the InvokeAPITester Sample Program](#)

18.1 Using APIs to Write Java Clients

The Web services framework offers two APIs in the following package to assist in the writing of Java clients that make Web service invocations:

```
oracle.sysman.emSDK.webservices.outbound
```

The SOAPMessageBuilder API can be used to build a payload SOAPMessage object by inserting items into the SOAP header and the SOAP body of the message. The resultant SOAPMessage can then be used by the WebServiceInvoker API to invoke a Web service method.

18.1.1 About SOAPMessageBuilder

This API offers different methods to construct a SOAP message. In addition, it allows the caller to set the actor, role, mustUnderstand, and relay attributes of the SOAP header element. Finally, the SOAP body of the message can be constructed with this API.

- `createSOAPMessage()` - create a empty SOAP message.
- `setHeaderElementActor()` - set the recipient of a SOAP header (SOAP1.1).
- `setHeaderElementRole()` - set the recipient of a SOAP header (SOAP 1.2).
- `setHeaderElementMustUnderstand()` - indicates whether a header entry is mandatory or optional for the recipient to understand.

- `setHeaderElementRelay()` - indicates whether the header is to be relayed to the next SOAP node if the current node doesn't understand the header.
- `addToSOAPHeader()` - add name/value pairs of elements and attributes to the SOAP header.
- `addSOAPHeader()` - add a DOM representation of a SOAP header.
- `addSOAPBody()` - add a DOM representation of a SOAP body.
- `addSOAPPayload()` - add entire SOAP message in XML format.
- `getSOAPMessage()` - returns constructed SOAP message.

18.1.2 About WebServiceInvoker

When a client program invokes a Web service method, there are various areas that need to be addressed, including:

- Security
- Setting up optional client-side SOAP message handlers
- Setting the attributes of the request
- Different modes of making the call

There are numerous predefined Web service security policies that are available through Oracle Web Service Manager with WebLogic Server in the areas of message-level and transport-level security.

As Enterprise Manager Web services support the `oracle/wss_username_token_service_policy`, the `WebServiceInvoker` API offers `setUpUNTAAuth(String username, String password)` to set up the corresponding client-side security. The API also has `setUpHttpBasicAuth(String username, String password)` for calls out to other web services that support that security scenario.

If desired, a client program can have client-side SOAP message handlers that modify parts of the SOAP message before it is actually sent to the server. To that end, custom SOAP message handlers that implement the `javax.xml.ws.handler.soap.SOAPHandler` interface can be written. To associate one with your client program, the `addClientSideSOAPMsgHandler(String handlerClass)` API can be called.

There are some basic attributes of a Web service method invocation that can be specified by the client. They include:

- connection timeout
(`com.sun.xml.ws.developer.JAXWSProperties.CONNECT_TIMEOUT`)
- request timeout
(`com.sun.xml.ws.developer.JAXWSProperties.REQUEST_TIMEOUT`)
- SOAP action (`javax.xml.ws.BindingProvider.SOAPACTION_URI_PROPERTY`)

These attributes can be set with these corresponding API calls:

- `setConnectTimeout(int connectTimeout)`
- `setRequestTimeout(int requestTimeout)`
- `setSOAPAction(String soapAction)`

`WebServiceInvoker` exposes four different ways to invoke a Web service:

- `invokeSyncWS(SOAPMessage request)`

calls the method and waits for a response.

- `invokeAsyncWS(SOAPMessage request, WebServiceAsyncProcessor asyncProcessor)`

calls the method and loops through the `process()` method of an implementation of `oracle.sysman.emSDK.webservices.outbound.WebServiceAsyncProcessor`.

- `invokeAsyncWithHandlerWS(SOAPMessage request, WebServiceAsyncProcessor asyncProcessor, AsyncHandler asyncHandler)`

Same as `invokeAsyncWS()`, except the results of the method invocation are processed by the `handleResponse()` method of an implementation of `javax.xml.ws.AsyncHandler`.

- `invokeOneWayWS(SOAPMessage request)`

Calls the method, immediately returns and doesn't expect a response back.

18.1.3 Using the InvokeAPITester Sample Program

The following is a sample test program that uses the `WebServiceInvoker` API to make Web service calls:

```
oracle.sysman.emSDK.samples.websvcs.outbound.InvokeAPITester
```

It is written to be flexible in that an input file of parameters is read and depending on the values of some of the parameters, the security mode and the invocation mode can be specified. The list of parameters that can be specified include:

- `protocol`
- `host`
- `port`
- `path`
- `name`

Name is the name in the `wsdl:definitions` tag in the WSDL.

- `username`

Used for security.

- `password`

Used for security.

- `calling mode`

Calling mode can be `sync`, `async`, `asynccb` or `oneway`.

- `security mode`

Security mode is `UNT`.

- `target namespace`

Target namespace is value of `targetNamespace` in `wsdl:definitions` tag in the WSDL for the Web service.

- `port name`

Port name is the name attribute of port in the WSDL.

- payload XML

Payload XML points to an XML file that contains the SOAP message for the request.

- default payload XML path

Default payload XML path is the directory in which the payload file is located.

- client-side SOAP message handler class

Handler class can specify a client-side SOAP message handler to manipulate the SOAP message before it is sent.

Note: The Web service endpoint is constructed as
protocol://host:port/path/name

Using Management Repository Views

Enterprise Manager repository views are used to access information in the Management Repository for further processing and presentation.

This chapter contains the following sections:

- [Overview](#)
- [Application Deployment Views](#)
- [Blackout Views](#)
- [Chargeback Views](#)
- [Compliance Views](#)
- [Compliance Real-time Monitoring Views](#)
- [Configuration Management Views](#)
- [Custom Configuration Specification Views](#)
- [Database Configuration Views](#)
- [Events Views](#)
- [Glassfish Views](#)
- [Hardware Views](#)
- [Inventory Views](#)
- [Job Views](#)
- [Linux Patching Views](#)
- [Management Template Views](#)
- [Metric Views](#)
- [Monitoring Views](#)
- [Operating System Views](#)
- [Oracle Home Directory Patching Views](#)
- [Oracle Home Directory Views](#)
- [Oracle WebLogic Server Views](#)
- [Oracle WebLogic Domain Views](#)
- [Oracle WebLogic Cluster Views](#)
- [Security Views](#)

- [Service Tag Views](#)
- [Storage Reporting Views](#)
- [Target Views](#)
- [VT Target Views](#)
- [Examples](#)

19.1 Overview

The Enterprise Manager Management Repository views provide access to target, metric, and monitoring information stored in the Management Repository. Accessing the repository will allow you to perform the following:

- Obtain pertinent application-specific information at the right level of granularity and density for a wider variety of users: IT staff, executives, developers.
- Send alerts for metric threshold violations.
- Perform historical analysis or additional computation on stored data.
- Seamless integration of Enterprise Manager alerts with user ticketing systems, such as iSupport and Remedy.

The Management Repository is the comprehensive source for all the management information for Enterprise Manager, with the key to extensibility being the repository's open schema. This open architecture allows users to customize how the information in the repository is used if Enterprise Manager's standard configuration does not meet their requirements. To facilitate easy access to information stored in the repository, Enterprise Manager supplies a comprehensive set of views rather than forcing the user to access repository base tables directly. Views buffer custom applications from any underlying changes to the repository schema and ensures up-stream applications will not break when the repository schema changes via patching or new releases.

19.1.1 Using Repository Views

Note: You must use the views that are documented in this guide and in the Extensibility Development Kit (EDK) only. Any other view that is not documented must not be used and backward compatibility for undocumented views and tables is not guaranteed.

Because the views are simple queries to a database, users can imbed these queries within any application code used to return information for further processing and/or display in the Enterprise Manager Cloud Control console.

As shown in [Example 19–1, "View Usage"](#), the Java code uses Enterprise Manager views to query the Management Repository rather than accessing the repository tables directly. For each of four time windows, there are four SQL statements with question marks (?) as placeholders for the parameters.

See Also: [Section 19.30, "Examples"](#) provides examples of how to use the Management Repository views.

Example 19–1 View Usage

```

public static final String hour_stmt =

"SELECT collection_timestamp, value " +
"FROM mgmt$metric_details " +
"WHERE target_type = ? and target_name = ? and metric_name = ? and metric_column=
? " +
"and collection_timestamp > sysdate - 1/24 " +
"ORDER BY collection_timestamp ";

public static final String day_stmt =

"SELECT rollup_timestamp, average " +
"FROM mgmt$metric_hourly " +
"WHERE target_type = ? and target_name = ? and metric_name = ? and metric_column=
? " +

"and rollup_timestamp > sysdate - 1 " +
"ORDER BY rollup_timestamp";

public static final String week_stmt =

"SELECT rollup_timestamp, average " +
"FROM mgmt$metric_daily " +
"WHERE target_type = ? and target_name = ? and metric_name = ? and metric_column=
? " +
"and rollup_timestamp > sysdate - 7 " +
"ORDER BY rollup_timestamp";

public static final String month_stmt =

"SELECT rollup_timestamp, average " +
"FROM mgmt$metric_daily " +
"WHERE target_type = ? and target_name = ? and metric_name = ? and metric_column=
? " +
"and rollup_timestamp > sysdate - 31 " +
"ORDER BY rollup_timestamp";

```

19.2 Application Deployment Views

This section provides a description of each application deployment view and its columns.

19.2.1 MGMT\$J2EE_APPLICATION

The MGMT\$J2EE_APPLICATION view displays general information about the Application configuration.

Table 19–1 MGMT\$J2EE_APPLICATION

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target
CM_TARGET_TYPE	The type of target: j2ee_application
CM_TARGET_NAME	The name of the target in Enterprise Manager
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected

Table 19–1 (Cont.) MGMT\$J2EE_APPLICATION

Column	Description
ECM_SNAPSHOT_ID	GUID of the snapshot
PATH	The fully resolved location of the application source files on the administration server
LOADORDER	A number value that indicates when the unit is deployed, relative to other DeployableUnits on a server, during startup
TYPE	Type of the module. The string value must match those defined by JSR 88: <i>Java EE Application Deployment</i> such as EAR or WAR.

19.2.2 MGMT\$J2EEAPP_EJBCOMPONENT

The MGMT\$J2EEAPP_EJBCOMPONENT view displays general information about the Enterprise JavaBeans (EJB) modules.

Table 19–2 MGMT\$J2EEAPP_EJBCOMPONENT

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target
CM_TARGET_TYPE	The type of target: j2ee_application
CM_TARGET_NAME	The name of the target in Enterprise Manager
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
NAME	Name of the EJB component
DEPLOYMENTORDER	Priority that the server uses when it deploys an item. The priority is relative to the other deployable items of same type.
KEEPGENERATED	Indicates whether KeepGenerated is enabled and whether EJB source files will be kept. Values: true, false.

19.2.3 MGMT\$J2EEAPP_JRFWS

The MGMT\$J2EEAPP_JRFWS view displays general information about the Java Required Files (JRF) Web Services configuration.

Table 19–3 MGMT\$J2EEAPP_JRFWS

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target
CM_TARGET_TYPE	The type of target: j2ee_application
CM_TARGET_NAME	The name of the target in Enterprise Manager
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
SERVICENAME	Name of JRF web service
WEBMODULE	Name of web module that contains the JRF web service

Table 19–3 (Cont.) MGMT\$J2EEAPP_JRFWS

Column	Description
JRFWEBSERVICEKEY	Key column computed as WebModuleName_WebServiceName where WebModuleName is the name of web module that contains the JRF web service and WebServiceName is the name of the JRF web service
DATABINDING	Data binding technology used by the web service port
EXPOSEWSDL	Specifies if the web service definition language (WSDL) is exposed for the service. Values: true, false.
METADATAEXCHANGE	Usage of WS-MetadataExchange for WSDL advertisement. Values: true, false
EXPOSETESTPAGE	Specifies whether the test page is exposed for the service. Values: true, false

19.2.4 MGMT\$J2EEAPP_JRFWSOPER

The MGMT\$J2EEAPP_JRFWSOPER view displays general information about the JRF Web Services Operation configuration.

Table 19–4 MGMT\$J2EEAPP_JRFWSOPER

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target
CM_TARGET_TYPE	The type of target: j2ee_application
CM_TARGET_NAME	The name of the target in Enterprise Manager
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
SERVICENAME	Name of the JRF web service
WEBMODULE	Name of the web module that contains the JRF web service
PORTNAME	Name of the JRF web service port
OPERATIONNAME	Name of the JRF web service port operation
SOAPACTION	SOAP action
ONEWAY	Indicates whether the operation is one-way. Values: true, false
INPUTENCODING	Operation input encoding
OUTPUTENCODING	Operation output encoding

19.2.5 MGMT\$J2EEAPP_JRFWSPOLICY

The MGMT\$J2EEAPP_JRFWSPOLICY view displays general information about the JRF Web Services Policy configuration.

Table 19–5 MGMT\$J2EEAPP_JRFWSPOLICY

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target
CM_TARGET_TYPE	The type of target: j2ee_application
CM_TARGET_NAME	The name of the target in Enterprise Manager

Table 19–5 (Cont.) MGMT\$J2EEAPP_JRFWSPOLICY

Column	Description
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
SERVICENAME	Name of the JRF web service
WEBMODULE	Name of the web module that contains the JRF web service
PORTNAME	Name of the JRF web service port
URI	Policy reference URI
CATEGORY	Category of the WS-Policy Reference. For example, security.
ENABLED	Specifies whether the policy references are enabled. Values: true, false.

19.2.6 MGMT\$J2EEAPP_JRFWSPORT

The MGMT\$J2EEAPP_JRFWSPORT view displays general information about the JRF web services port configuration.

Table 19–6 MGMT\$J2EEAPP_JRFWSPORT

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target
CM_TARGET_TYPE	The type of target: j2ee_application
CM_TARGET_NAME	The name of the target in Enterprise Manager
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
SERVICENAME	Name of the JRF web service
WEBMODULE	Name of web module that contains the JRF web service
PORTNAME	Name of the JRF web service port
AVAILABLE	Indicates if a port is available. Possible Values: <ul style="list-style-type: none"> ■ True ■ False
RESTSUPPORTED	Indicates whether the port supports REST. Possible Values: <ul style="list-style-type: none"> ■ True ■ False
LOGGINGLEVEL	The logging level for the web service port
MAXREQUESTSIZE	Largest size of message in bytes the port can accept
STYLE	SOAP binding style
SOAPVERSION	Version of the SOAP protocol the port supports

Table 19–6 (Cont.) MGMT\$J2EEAPP_JRFWSHORT

Column	Description
STATEFUL	Indicates if the port is stateful. Possible Values: <ul style="list-style-type: none"> ■ True ■ False
IMPLEMENTORTYPE	Implementor type of this port, such as JAXWS or JAXRPC
TRANSPORTS	Transports from which the web service port is available
ENDPOINTADDRESSURI	The sub-context of the HTTP URL of the web service port exposing an EJB(2.1)
POLICYSUBJECTNAME	The name of the policy subject
POLICYSUBJECTRESOURCEPATTERN	The resource pattern of the policy subject
POLICYATTACHMENTSUPPORT	Determines the class of the supported policies
POLICYSUBJECTTYPE	The type of the policy subject
LEGACYCONFIG	Indicates whether the port has legacy management configuration. Possible values: <ul style="list-style-type: none"> ■ True ■ False
IMPLEMENTORCLASS	The name of the user-provided class that implements the web service port
WSDLURI	The URI to the port WSDL definition
SCHEMAVALIDATEINPUT	Optional validation of input against WSDL schema. Possible values: <ul style="list-style-type: none"> ■ True ■ False
ASYNC	Specifies if async is available. Possible values: <ul style="list-style-type: none"> ■ True ■ False
ASYNCJNDIDESTRESPONSE	JMS queue name for saving responses
ASYNCJNDIDEST	JMS queue name for saving asynchronous requests
ASYNCCONNFACTRESPONSE	JMS connection factory name for saving responses
ASYNCCONNFACT	JMS connection factory name for saving asynchronous requests

19.2.7 MGMT\$J2EEAPP_WEBAPPCOMPONENT

The MGMT\$J2EEAPP_WEBAPPCOMPONENT view displays general information about the web modules.

Table 19–7 MGMT\$J2EEAPP_WEBAPPCOMPONENT

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target
CM_TARGET_TYPE	The type of target: j2ee_application
CM_TARGET_NAME	The name of the target in Enterprise Manager
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
NAME	Name of the web module
DEPLOYMENTORDER	A priority that the server uses to determine when it deploys an item. The priority is relative to other deployable items of the same type
CONTEXTPATH	Context path
SESSIONTIMEOUTSECS	Session timeout in seconds

19.2.8 MGMT\$J2EEAPP_WSCONFIG

The MGMT\$J2EEAPP_WSCONFIG view displays general information about the Web Service configuration.

Table 19–8 MGMT\$J2EEAPP_WSCONFIG

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target
CM_TARGET_TYPE	The type of target: j2ee_application
CM_TARGET_NAME	The name of the target in Enterprise Manager
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
SERVICENAME	Name of the web service
NAMEINWSDL	The "name" attribute of the "service" element in the WSDL that describes the Web service. It is specified at development time using the serviceName attribute of the @WebService JWS annotation
IMPLEMENTATIONTYPE	Implementation type of the service. The allowed values are: JAX-WS 2.0 JAX-RPC 1.1
URI	URI of this Web Service. The value corresponds to the final part of the endpoint address in the WSDL that describes the Web services

19.2.9 MGMT\$J2EEAPP_WSPORTCONFIG

The MGMT\$J2EEAPP_WSPORTCONFIG view displays general information about the Web Services Port configuration.

Table 19–9 MGMT\$J2EEAPP_WSPORTCONFIG

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target

Table 19–9 (Cont.) MGMT\$J2EEAPP_WSPORTCONFIG

Column	Description
CM_TARGET_TYPE	The type of target: j2ee_application
CM_TARGET_NAME	The name of the target in Enterprise Manager
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
SERVICENAME	The "name" attribute of the "service" element in the WSDL that describes the Web service. It is specified at development time using the serviceName attribute of the @WebService JWS annotation.
PORTNAME	Name of the web service port
TRANSPORTPROTOCOL	Transport protocol used to invoke this web service, such as HTTP, HTTPS, or JMS

19.3 Blackout Views

This section provides a description about each blackout view and its columns. Blackouts permit you to suspend monitoring on one or more targets in order to perform maintenance operations.

For examples of how to use these views, see [Section 19.30, "Examples"](#).

19.3.1 MGMT\$BLACKOUT_HISTORY

The MGMT\$BLACKOUT_HISTORY view displays a historical log of changes in the blackout state for a managed target. In addition, the view can be used to generate a list of targets that were in a blackout period for a specific period of time.

Table 19–10 MGMT\$BLACKOUT_HISTORY

Column	Description
BLACKOUT_NAME	The name of the blackout
CREATED_BY	The Enterprise Manager administrator who created the blackout
BLACKOUT_GUID	The unique global identifier (GUID) for the blackout
START_TIME	Start of the blackout period for the managed target
END_TIME	End of the blackout period for the managed target. If the target is currently in a blackout period, the END_TIMESTAMP date will be NULL.
TARGET_NAME	Name of the target where the metric was collected. The target name uniquely identifies a managed target within the Management Repository. The target name typically contains the name of the managed entity that was provided by the system or database administrator.
TARGET_TYPE	Types of targets may include databases, hosts, web servers, applications, or Application Servers. The definer of the collection definition at the Management Agent defines the target type. The target type defines the set of metrics that are collected for a managed target within the Management Repository.

Table 19–10 (Cont.) MGMT\$BLACKOUT_HISTORY

Column	Description
STATUS	Current status of the blackout Status Values: <ul style="list-style-type: none"> ■ 0: Scheduled ■ 1: Start Processing ■ 2: Start Partial ■ 4: Started ■ 5: Stop Pending ■ 6: Stop Failed ■ 7: Stop Partial ■ 8: Edit Failed ■ 9: Edit Partial ■ 10: Stopped ■ 11: Ended ■ 12: Partial Blackout ■ 13: Modify Pending

Usage Notes

Queries using this view will use an index if they reference the target_name, target_type, start_timestamp, or end_timestamp.

Typically, blackout history information retrieved using this view will be ordered by target_name, target_type, and start_timestamp.

19.3.2 MGMT\$BLACKOUTS

The MGMT\$BLACKOUTS view displays all blackout definition information along with current schedules.

Table 19–11 MGMT\$BLACKOUTS

Column	Description
BLACKOUT_NAME	The name of the blackout
BLACKOUT_GUID	The unique global identifier (GUID) of the blackout
REASON	Purpose of the blackout. Reasons are chosen from a predefined list by the report owner
DESCRIPTION	Detailed information about the blackout

Table 19–11 (Cont.) MGMT\$BLACKOUTS

Column	Description
STATUS	Current status of the blackout Status Values: <ul style="list-style-type: none"> 0: Scheduled 1: Start Processing 2: Start Partial 4: Started 5: Stop Pending 6: Stop Failed 7: Stop Partial 8: Edit Failed 9: Edit Partial 10: Stopped 11: Ended 12: Partial Blackout 13: Modify Pending
CREATED_BY	Administrator who created the blackout. CREATED_BY returns SYSTEM as the blackout owner if the blackout was created using the Enterprise Manager Command Line Interface.
LAST_START_TIME	Last time the blackout successfully started
LAST_END_TIME	Last time the blackout successfully ended
SCHEDULED_TIME	Possible values are: <ul style="list-style-type: none"> 0 - Immediate schedule 1 - Run once at specified time 2 - Run on interval 3 - Run daily 4 - Run on specified days of the week 5 - Run on specified days of the month 6 - Run on specified days of the year
SCHEDULE_START_TIME	Time the blackout is scheduled to start.
SCHEDULE_END_TIME	Time the blackout is scheduled to end
DURATION	Duration of the blackout in minutes

19.4 Chargeback Views

This section provides a description of each Chargeback view and its columns.

19.4.1 MGMT\$EMCT_CBA_CHARGE_HOURLY

This view provides hour aggregated metering and charge data. The configuration metrics and fixed charge metrics are at day aggregation level.

Table 19–12 MGMT\$EMCT_CBA_CHARGE_HOURLY

Column	Description
CONSUMER_NAME	Represents the internal cost center name to which the currently being charged target was assigned.
CONSUMER_DISPLAY_NAME	Represents the display name of cost center to which the currently being charged target was assigned.
COLLECTION_DATE	Represents the UTC date at which the charge item value was collected.
YEAR	Represents the year to which the collection date belongs.
MONTH_OF_YEAR	Represents the month of the year to which the collection date belongs.
DAY_OF_MONTH	Represents the day of the month to which the collection date belongs.
WEEK_OF_YEAR	Represents the week of the year to which the collection date belongs.
QUARTER_OF_YEAR	Represents the quarter of the year to which the current collection date belongs.
DAY_DATE	Represents the day date to which the current collection date belongs.
HOUR_OF_DAY	Represents the hour within the day to which the current collection date belongs.
TYPE_NAME	Represents the target type display name of the target that is being charged.
TARGET_NAME	Represents the internal name of the target that is being charged.
TARGET_DISPLAY_NAME	Represents display name of the target that is being charged.
HOST_NAME	Represents the host on which the currently being charged target is deployed.
ITEM	Represents the display name of the charge item for the charge being computed.
UNIT	Represents display name of the charge item unit for the charge being computed
CATEGORY	Represents the charge item category for the charge being computed. Possible values are <ul style="list-style-type: none"> ■ Instance ■ Service ■ CPU ■ Memory ■ Disk/Storage ■ Network ■ Software ■ Activity ■ Instance Uptime ■ Unclassified

Table 19–12 (Cont.) MGMT\$EMCT_CBA_CHARGE_HOURLY

Column	Description
USAGE_VALUE	Represents the usage value of the charge item. This column has a value if the charge item is a "number" data type. If the charge item aggregation type is "sum", then it will represent the sum value of the item within that hour. If the charge item aggregation type is "avg" then it represents the average value within that hour.
STRING_VALUE	Represents the string value of the charge item. This column has a value if the charge item has a "string" data type.
CHARGE_PLAN	Represents the name of the charge plan used to calculate the charge.
CHARGE_PLAN_CONFIG	Represents plan configuration within the charge plan used to calculate the charge.
RATE_TYPE	Represents the rate type of the charge item: <ul style="list-style-type: none"> ■ Config The charge item that is being charged was charged based on configuration. ■ Usage The charge item that is being charged was charged based on usage. ■ Flat The charge item that is being charged based on flat rate
RATE	Represents the charge rate expression defined in the charge plan configuration for the charge item that is being charged. If the charge item that is being charged is of universal item type, then this column represents the factor value with universal item defined rate.
CHARGE	Represents the computed charge value of the charge item that is being charged.
CHARGE_ADJUSTMENT	Represents the applicable charge adjustment value for the charge item that is being charged.
ADJUSTED_CHARGE	Represents the final charge value of the charge item that is being charged, after charge adjustments are taken into consideration.
UPTIME	Represents uptime (Hours) in the day of the target that is being charged. For "metric" items this value is null.

19.4.2 MGMT\$EMCT_CBA_CHARGE_DAILY

This view provides day aggregated charge data

Table 19–13 MGMT\$EMCT_CBA_CHARGE_DAILY

Column	Description
CONSUMER_NAME	Represents the internal cost center name to which the currently being charged target was assigned.
CONSUMER_DISPLAY_NAME	Represents the display name of cost center to which the currently being charged target was assigned.
COLLECTION_DATE	Represents the UTC date at which the charge item value was collected.

Table 19–13 (Cont.) MGMT\$EMCT_CBA_CHARGE_DAILY

Column	Description
YEAR	Represents the year to which the collection date belongs.
MONTH_OF_YEAR	Represents the month of year to which the collection date belongs.
DAY_OF_MONTH	Represents the day of month to which the collection date belongs.
WEEK_OF_YEAR	Represents the week of year to which the collection date belongs.
QUARTER_OF_YEAR	Represents the quarter of year to which the current collection date belongs.
DAY_DATE	Represents the day date to which the current collection date belongs.
TYPE_NAME	Represents the target type display name of the target that is being charged.
TARGET_NAME	Represents the internal name of the target that is being charged.
TARGET_DISPLAY_NAME	Represents display name of the target that is being charged.
HOST_NAME	Represents the host on which the currently being charged target is deployed.
ITEM	Represents the display name of the charge item for the charge being computed in the charge computing target.
UNIT	Represents the display name of the charge item unit for the charge being computed.
CATEGORY	Represents the charge item category for the charge being computed. Possible values are <ul style="list-style-type: none"> ■ Instance ■ Service ■ CPU ■ Memory ■ Disk/Storage ■ Network ■ Software ■ Activity ■ Instance Uptime ■ Unclassified
USAGE_VALUE	Represents the usage value of the charge item. This column has a value if the charge item has a "number" data type. If the charge item aggregation type is "sum", then it will represent the sum value of the item within that day. If the charge item aggregation type is "avg" then it represents average value within that day.
STRING_VALUE	Represents the string value of the charge item. This column has a value if the charge item has a "string" data type.
CHARGE_PLAN	Represents the name of the charge plan used to calculate the charge.
CHARGE_PLAN_CONFIG	Represents the plan configuration within the charge plan used to calculate the charge.

Table 19–13 (Cont.) MGMT\$EMCT_CBA_CHARGE_DAILY

Column	Description
RATE_TYPE	Represents the rate type of the charge item <ul style="list-style-type: none"> ■ Config The charge item that is being charged was charged based on the "Per Unit" basis. ■ Usage The charge item that is being charged was charged based on usage. ■ Flat The charge item that is being charged was based on a flat rate.
RATE	Represents charge rate expression defined in the charge plan configuration for the charge item that is being charged. If the charge item that is being charged is of universal item type, then this column represents the factor value with universal item defined rate.
CHARGE	Represents the computed charge value of the charge item that is being charged.
CHARGE_ADJUSTMENT	Represents the applicable charge adjustment value for the charge item that is being charged.
ADJUSTED_CHARGE	Represents final charge value of the charge item that is being charged, after charge adjustments are taken into consideration.
UPTIME	Represents uptime (Hours) in the day of the target that is being charged. For "metric" items this value is null

19.5 Compliance Views

This section provides a description of each compliance view and its columns. Compliance is the conformance to standards, or requirements, or both. Enterprise Manager Compliance Management provides the ability to evaluate the compliance of targets and systems as they relate to business best practices for configuration, security, and storage. This is accomplished by defining, customizing, and managing compliance frameworks, compliance standards, and compliance standard rules. In addition, it provides advice of how to change configuration to bring your targets and systems into compliance.

For examples of how to use these views, see [Section 19.30, "Examples"](#).

19.5.1 MGMT\$COMPLIANCE_STANDARD_RULE

The MGMT\$COMPLIANCE_STANDARD_RULE view contains the lists of all the compliance standard rules. A compliance standard rule is a test to determine if a configuration data change affects compliance. A compliance standard rule is mapped to one or more compliance standards.

Table 19–14 MGMT\$COMPLIANCE_STANDARD_RULE

Column	Description
RULE_NAME	Display name in English
DESCRIPTION	Description of the rule in English
TARGET_TYPE	Applicable target type of rule

Table 19–14 (Cont.) MGMT\$COMPLIANCE_STANDARD_RULE

Column	Description
REFERENCE_URL	Not used in this release
RATIONALE	Explains the importance of this rule, and the consequences of noncompliance
FIXTEXT	Explains the steps to bring the target into compliance with respect to this rule
WARNING	Cautionary or caveat note about this rule
RULE_TYPE	Type of rule. Possible values: <ul style="list-style-type: none"> ■ Repository ■ Agent ■ Monitoring
MESSAGE	Message recorded for new violation
CLEAR_MESSAGE	Message recorded for clear violation
SEVERITY	The severity of the rule Possible values: <ul style="list-style-type: none"> ■ Minor Warning ■ Warning ■ Critical
LIFECYCLE_STATE	Lifecycle status of the rule Possible values: <ul style="list-style-type: none"> ■ Development ■ Production ■ Draft
AUTHOR	Author of the rule
OWNER	Owner of the rule
IS_SYSTEM	Specifies whether the rule is system defined Possible values <ul style="list-style-type: none"> ■ False ■ True
RULE_DNAME_NLSID	NLSID of the rule display name for non-English users
DESCRIPTION_NLSID	NLSID of the rule description for non-English users
RATIONALE_NLSID	NLSID of the rule impact for non-English users
FIXTEXT_NLSID	NLSID of the rule recommendation for non-English users
WARNING_NLSID	NLSID of the rule warning for non-English users
RULE_TYPE_CODE	Code to represent the type of compliance standard rule. Possible values: <ul style="list-style-type: none"> ■ 1: Repository ■ 2: Agent ■ 3: Monitoring

Table 19–14 (Cont.) MGMT\$COMPLIANCE_STANDARD_RULE

Column	Description
SEVERITY_CODE	Code to represent the severity of the compliance standard rule. Possible values: <ul style="list-style-type: none"> ■ 18: Minor Warning ■ 20: Warning ■ 25: Critical
LIFECYCLE_STATE_CODE	Code to represent the status of the lifecycle of the compliance standard rule. Possible values: <ul style="list-style-type: none"> ■ 1: Development ■ 2: Production ■ 3: Draft
IS_SYSTEM_CODE	Code to represent whether the compliance standard rule is system defined. Possible values: <ul style="list-style-type: none"> ■ 0: False ■ 1: True

19.5.2 MGMT\$COMPLIANCE_STANDARD

The MGMT\$COMPLIANCE_STANDARD view contains the lists of all compliance standards. A compliance standard is a collection of checks or rules. It is the Enterprise Manager representation of a compliance control that must be tested against some set of IT infrastructure to determine if the control is being followed.

Table 19–15 MGMT\$COMPLIANCE_STANDARD

Column	Description
CS_NAME	Display name in English
TARGET_TYPE	Applicable target type of the compliance standard
AUTHOR	Author of the compliance standard
OWNER	Owner of the compliance standard
VERSION	Version of the compliance standard
KEYWORDS	Keywords associated with the compliance standard
LIFECYCLE_STATUS	Lifecycle status of the compliance standard Possible values: <ul style="list-style-type: none"> ■ Development ■ Production
AUTO_ENABLE	Specifies whether the compliance standard should be associated with applicable target automatically Possible values: <ul style="list-style-type: none"> ■ False ■ True
DESCRIPTION	Description of the compliance standard in English
REFERENCE_URL	Not used in this release

Table 19–15 (Cont.) MGMT\$COMPLIANCE_STANDARD

Column	Description
FRONT_MATTER	Introductory text of the compliance standard
REAR_MATTER	Concluding text of the compliance standard
NOTICE	Legal notice or copyright text about compliance standard
IS_SYSTEM	Specifies whether the compliance standard is system defined Possible values: <ul style="list-style-type: none"> False True
CS_DNAME_NLSID	NLSID of the standard display name for non-English users
LIFECYCLE_STATE_CODE	Code representing the status of the compliance standard lifecycle. Possible values: <ul style="list-style-type: none"> 1: Development 2: Production
AUTO_ENABLE_CODE	Code representing whether the compliance standard should be automatically associated with an applicable target. Possible values: <ul style="list-style-type: none"> 0: False 1: True
DESCRIPTION_NLSID	NLSID of the compliance standard description for non-English users
NOTICE_NLSID	NLSID of the legal notice or copyright text about the compliance standard for non-English users
IS_SYSTEM_CODE	Code representing whether the compliance standard is system defined. Possible values: <ul style="list-style-type: none"> 0: False 1: True
CS_TYPE	Type of compliance standard
CS_TYPE_CODE	Code representing the type of compliance standard. Possible values: <ul style="list-style-type: none"> 1: Repository 2: WebLogic server signature 3: Real-time monitoring

19.5.3 MGMT\$COMPLIANCE_STANDARD_GROUP

The MGMT\$COMPLIANCE_STANDARD_GROUP view contains the lists of the compliance standard groups.

Table 19–16 MGMT\$COMPLIANCE_STANDARD_GROUP

Column	Description
CSG_NAME	The display name in English
AUTHOR	Author of the compliance standard group
OWNER	Owner of the compliance standard group

Table 19–16 (Cont.) MGMT\$COMPLIANCE_STANDARD_GROUP

Column	Description
VERSION	The version of the compliance standard group
LIFECYCLE_STATUS	Lifecycle status of the compliance standard group Possible values: <ul style="list-style-type: none"> ■ Development ■ Production
DESCRIPTION	Description of the compliance standard group in English
REFERENCE_URL	Not used in this release
FRONT_MATTER	Introductory text of the compliance standard group
REAR_MATTER	Concluding text of the compliance standard group
NOTICE	Legal notice or copyright text about the compliance standard group
IS_SYSTEM	Specifies whether the compliance standard group is system defined Possible values: <ul style="list-style-type: none"> ■ False ■ True
CSG_DNAME_NLSID	NLSID of the compliance standard group display name for non-English users
LIFECYCLE_STATE_CODE	Code representing the status of the compliance standard group lifecycle. Possible values: <ul style="list-style-type: none"> ■ 1: Development ■ 2: Production
DESCRIPTION_NLSID	NLSID of the compliance standard group description for non-English users
NOTICE_NLSID	NLSID of the legal notice or copyright text about the compliance standard group for non-English users
IS_SYSTEM_CODE	Coderepresenting whether the compliance standard group is system defined. Possible values: <ul style="list-style-type: none"> ■ 0: False ■ 1: True

19.5.4 MGMT\$CS_EVAL_SUMMARY

The MGMT\$CS_EVAL_SUMMARY view contains the lists of all the root compliance standard scores.

Table 19–17 MGMT\$CS_EVAL_SUMMARY

Column	Description
CS_GUID	Unique identifier of compliance standard Note: You can obtain this value from the MGMT\$COMPLIANCE_STANDARD view.
TARGET_GUID	Unique identifier of target

Table 19–17 (Cont.) MGMT\$CS_EVAL_SUMMARY

Column	Description
CS_NAME	Internal name of the compliance standard
CS_INAME	English display name of the compliance standard
CS_AUTHOR	Author of the compliance standard
CS_VERSION	Version of the compliance standard
TARGET_NAME	Target name
TARGET_TYPE	Target type
COMPLIANT_RULES	Number of compliant rules in the compliance standard hierarchy for that target
CRITICAL_RULES	Number of critical rules in the compliance standard hierarchy for that target
WARN_RULES	Number of warning rules in the compliance standard hierarchy for that target
MWARN_RULES	Number of minor warning rules in the compliance standard hierarchy for that target
NON_COMPLIANT_RULES	Number of noncompliant rules in the compliance standard hierarchy for that target
ERROR_RULES	Number of error rules in the compliance standard hierarchy for that target
UNKNOWN_RULES	Number of unknown rules in the compliance standard hierarchy for that target
CRIT_VIOLATIONS	Total critical violations raised by compliance standard
WARN_VIOLATIONS	Total warning violations raised by compliance standard
MWARN_VIOLATIONS	Total minor warning violations raised by compliance standard
TOTAL_VIOLATIONS	Total violations raised by compliance standard
COMPLIANCE_SCORE_LEVEL	Specifies the compliance score level Possible values: <ul style="list-style-type: none"> ■ Compliant ■ Critical ■ Warning
LAST_EVALUATION_DATE	Last score evaluation date
COMPLIANCE_SCORE	Compliance score of standard
IS_SCORE_VALID	Specifies whether the compliance score is valid.
CS_TYPE	The type of compliance standard Possible values: <ul style="list-style-type: none"> ■ Repository ■ WebLogic Server Signature ■ Real-time Monitoring
CS_DNAME_NLSID	NLSID of the standard display name for non-English users

Table 19–17 (Cont.) MGMT\$CS_EVAL_SUMMARY

Column	Description
COMPLIANCE_SCORE_LEVEL_CODE	Represents compliance score level Possible values: <ul style="list-style-type: none"> 0: Compliant 1: Critical 2: Warning
IS_SCORE_VALID_CODE	Represents whether the compliance score is valid Possible values: <ul style="list-style-type: none"> 0: False 1: True
CS_TYPE_CODE	Represents the type of compliance standard Possible values: <ul style="list-style-type: none"> 1: Repository 2: WebLogic Server Signature 3: Real-time monitoring

19.5.5 MGMT\$COMPOSITE_CS_EVAL_SUMMARY

The MGMT\$COMPOSITE_CS_EVAL_SUMMARY view contains the list of all the compliance standard scores. Each row in the MGMT\$COMPOSITE_CS_EVAL_SUMMARY view represents the results for a top level compliance standard or top level target, and an included compliance standard or member target.

When you include a compliance standard within another top level compliance standard, the included standard must be of the same target type as the top level compliance standard. If the top level compliance standard is a composite target type, then the included standard can be one of the member target types of the composite target type.

Note: A root compliance standard is associated to a root target (of composite target type). Compliance standards are associated to member targets of the same applicable target type and target filter criteria.

Table 19–18 MGMT\$COMPOSITE_CS_EVAL_SUMMARY

Column	Description
ROOT_CS_GUID	Unique identifier of the root compliance standard Note: You can obtain this from the MGMT\$COMPLIANCE_STANDARD view.
RQS_GUID	Unique identifier of compliance standard within root compliance standard context
CS_GUID	Unique identifier of compliance standard Note: You can obtain this from the MGMT\$COMPLIANCE_STANDARD view.
ROOT_TARGET_GUID	Unique identifier of root target
TARGET_GUID	Unique identifier of target

Table 19–18 (Cont.) MGMT\$COMPOSITE_CS_EVAL_SUMMARY

Column	Description
ROOT_CS_NAME	Internal name of root compliance standard
ROOT_CS_INAME	English display name of root compliance standard
ROOT_CS_AUTHOR	Author of root compliance standard
ROOT_CS_VERSION	Version of root compliance standard
CS_NAME	Internal name of compliance standard
CS_INAME	Display name in English
CS_AUTHOR	Author of standard
CS_VERSION	Version of standard
ROOT_TARGET_NAME	Root target name
ROOT_TARGET_TYPE	Root target type
TARGET_NAME	Target name
TARGET_TYPE	Target type
COMPLIANT_RULES	Number of compliant rules in the compliance standard hierarchy for that target or member target
CRITICAL_RULES	Number of critical rules in the compliance standard hierarchy for that target or member target
WARN_RULES	Number of warning rules in the compliance standard hierarchy for that target or member target
MWARN_RULES	Number of minor warning rules in the compliance standard hierarchy for that target or member target
NON_COMPLIANT_RULES	Number of noncompliant rules in the compliance standard hierarchy for that target or member target
ERROR_RULES	Number of error rules in the compliance standard hierarchy for that target or member target
UNKNOWN_RULES	Number of unknown rules in the compliance standard hierarchy for that target or member target
CRIT_VIOLATIONS	Total critical violations raised by compliance standard
WARN_VIOLATIONS	Total warning violations raised by compliance standard
MWARN_VIOLATIONS	Total minor warning violations raised by compliance standard
TOTAL_VIOLATIONS	Total violations raised by compliance standard
SUPPRESSED_CRIT	Number of suppressed critical violations
SUPPRESSED_WARN	Number of suppressed warning violations
SUPPRESSED_MWARN	Number of suppressed minor warning violations
COMPLIANCE_SCORE_LEVEL	Compliance score level Possible values: <ul style="list-style-type: none"> ■ Compliant ■ Critical ■ Warning
LAST_EVALUATION_DATE	Last evaluation date
COMPLIANCE_SCORE	Compliance score of standard

Table 19–18 (Cont.) MGMT\$COMPOSITE_CS_EVAL_SUMMARY

Column	Description
ROOT_CS_NAME_NLSID	NLSID of the name of the root compliance standard for non-English users
CS_NAME_NLSID	NLSID of the name of the compliance standard for non-English users
COMPLIANCE_SCORE_LEVEL_CODE	Code representing the compliance score level. Possible values: <ul style="list-style-type: none"> 0: Compliant 1: Critical 2: Warning
IS_SCORE_VALID_CODE	Code representing whether the compliance score is valid. Possible values: <ul style="list-style-type: none"> 0: False 1: True

19.5.6 MGMT\$CS_RULE_EVAL_SUMMARY

The MGMT\$CS_RULE_EVAL_SUMMARY view contains the lists of all the compliance rule scores for the target.

Table 19–19 MGMT\$CS_RULE_EVAL_SUMMARY

Column	Description
ROOT_CS_GUID	Unique identifier of the root compliance standard Note: You can obtain this from the MGMT\$COMPLIANCE_STANDARD view.
RQS_GUID	Unique identifier of the rule within root compliance standard context
RULE_GUID	Unique identifier of the compliance rule Note: You can obtain this from the MGMT\$COMPLIANCE_STANDARD_RULE view.
ROOT_TARGET_GUID	Unique identifier of the root target
TARGET_GUID	Unique identifier of the target
ROOT_CS_NAME	Internal name of the root compliance standard
ROOT_CS_INAME	Display name of the root compliance standard in English
ROOT_CS_AUTHOR	Author of the root compliance standard
ROOT_CS_VERSION	Version of the root compliance standard
PARENT_CS_NAME	Internal name of the parent compliance standard
PARENT_CS_INAME	Display name of the parent compliance standard in English
PARENT_CS_AUTHOR	Author of the standard of parent compliance standard
PARENT_CS_VERSION	Version of the parent compliance standard
RULE_NAME	Display name of the rule in English
RULE_INAME	Internal name of the rule
ROOT_TARGET_NAME	Root target name

Table 19–19 (Cont.) MGMT\$CS_RULE_EVAL_SUMMARY

Column	Description
ROOT_TARGET_TYPE	Root target type
TARGET_NAME	Target name
TARGET_TYPE	Target type
TOTAL_VIOLATIONS	Total violations raised by the compliance rule within the compliance standard context
LAST_EVALUATION_DATE	Last evaluation date
COMPLIANCE_SCORE	Compliance score of the rule with respect to compliance standard context
IS_SCORE_VALID	Specifies whether the compliance score is valid
ROOT_CS_NAME_NLSID	NLSID of the name of the root compliance standard for non-English users
PARENT_CS_NAME_NLSID	NLSID of the name of the parent compliance standard for non-English users
IS_SCORE_VALID_CODE	Code representing whether the compliance score is valid. Possible values: <ul style="list-style-type: none"> ■ 0: False ■ 1: True

19.5.7 MGMT\$CS_GROUP_EVAL_SUMMARY

The MGMT\$CS_GROUP_EVAL_SUMMARY view contains the lists of all the compliance standard group scores.

Table 19–20 MGMT\$CS_GROUP_EVAL_SUMMARY

Column	Description
CSG_GUID	Unique identifier of compliance standard group Note: You can obtain this from the MGMT\$COMPLIANCE_STANDARD_GROUP view.
CSG_NAME	Internal name of compliance standard group
CSG_INAME	Display name in English
CSG_VERSION	Version of compliance standard group
CRITICAL_EVALUATIONS	Number of critical evaluations
WARNING_EVALUATIONS	Number of warning evaluations
COMPLIANT_EVALUATIONS	Number of compliant evaluations
CRITICAL_VIOLATIONS	Total critical violations
WARN_VIOLATIONS	Total warning violations
MWARN_VIOLATIONS	Total minor warning violations
COMPLIANCE_SCORE	Compliance score of compliance standard group

19.5.8 MGMT\$CS_TARGET_ASSOC

The MGMT\$CS_TARGET_ASSOC view contains the lists all the root compliance standard and target associations.

Table 19–21 *MGMT\$CS_TARGET_ASSOC*

Column	Description
CS_GUID	Unique identifier of compliance standard Note: You can obtain this from the MGMT\$COMPLIANCE_STANDARD view.
TARGET_GUID	Unique identifier of target
CS_NAME	Internal name of compliance standard
CS_INAME	Display name in English
CS_AUTHOR	Author of the standard
CS_VERSION	Version of the standard
TARGET_NAME	Target name
TARGET_TYPE	Target type
CRIT_THRESHOLD	Critical threshold value
WARN_THRESHOLD	Warning threshold value
STATUS	Status of the association Possible values: <ul style="list-style-type: none"> ■ Enabled ■ Disabled
CS_DNAME_NLSID	NLSID of the standard display name for non-English users
STATUS_CODE	Code representing the status of the association. Possible values: <ul style="list-style-type: none"> ■ 1: Enabled ■ 2: Disabled ■ 3: Pending Enable ■ 4: Pending Disable

19.5.9 MGMT\$CSR_CURRENT_VIOLATION

The MGMT\$CSR_CURRENT_VIOLATION view contains the active violations of all compliance rules.

Table 19–22 *MGMT\$CSR_CURRENT_VIOLATION*

Column	Description
ROOT_CS_GUID	Unique GUID of the root compliance standard
RQS_GUID	Unique GUID of rule inclusion within the root compliance standard
RULE_GUID	Unique GUID of the rule
ROOT_TARGET_GUID	Unique GUID of the root target
TARGET_GUID	Unique GUID of the target
POLICY_GUID	Unique GUID of the policy (repository rule)
KEY_VALUE	The key value of the violation
COLLECTION_TIMESTAMP	The timestamp when the violation occurred

Table 19–22 (Cont.) MGMT\$CSR_CURRENT_VIOLATION

Column	Description
VIOLATION_GUID	Unique GUID identifying the violation
VIOLATION_LEVEL	Specifies the priority level of the violation Possible values: <ul style="list-style-type: none"> ■ 18: Minor warning ■ 20: Warning ■ 25: Critical
RULE_TYPE	Specifies the type of compliance rule being violated Possible values: <ul style="list-style-type: none"> ■ 1: Repository rule ■ 2: Guardian rule ■ 3: Compliance real-time rule
ANNOTATED_FLAG	Not used in this release
MESSAGE	Violation message of the rule
MESSAGE_NLSID	NLSID of the violation message of the rule
MESSAGE_PARAMS	Violation message parameters
ACTION_MESSAGE_NLSID	Not used in this release
ACTION_MESSAGE_PARAMS	Not used in this release

19.5.10 MGMT\$CSR_VIOLATION_CONTEXT

The MGMT\$CSR_VIOLATION_CONTEXT view contains the violation context, that is extra columns defined in the rule to be collected for a violation

Table 19–23 MGMT\$CSR_VIOLATION_CONTEXT

Column	Description
VIOLATION_GUID	Unique GUID identifying the violation
COLLECTION_TIMESTAMP	Timestamp at which the violation occurred
COLUMN_NAME	The name of the column of the violation context
COLUMN_TYPE	Type of the column name Possible values: <ul style="list-style-type: none"> ■ 1: Numeric ■ 2: String
COLUMN_VALUE	Specifies the numeric value of the column Note: Applies only when COLUMN_TYPE is set to 1
COLUMN_STR_VALUE	Specifies the string value of the column Note: Applies only when COLUMN_TYPE is set to 2

19.5.11 MGMT\$EM_RULE_VIOL_CTXT_DEF

The MGMT\$EM_RULE_VIOL_CTXT_DEF view stores the violation context definition of compliance standard rules. Each row stores one violation column definition of a compliance standard rule.

Table 19–24 MGMT\$EM_RULE_VIOL_CTXT_DEF

Column	Description
RULE_GUID	Unique GUID of the compliance standard rule
COLUMN_INAME	Internal name of the column
COLUMN_DNAME	Display name of the column
COLUMN_DNAME_NLSID	The NLSID of the display name of the column
COLUMN_TYPE	Data type of the column. Possible values: <ul style="list-style-type: none"> 1: Number 2: String
COLUMN_POSITION	Position of the column within the violation context definition
IS_KEY	Specifies whether the column is a key column (1=key column)
IS_HIDDEN	Specifies whether to show or hide the violation column when viewing the violation in the rule violations UI. Possible values: <ul style="list-style-type: none"> 0: Show 1: Hide
LINK_TEMPLATE	Not used in the current release
LINK_ENCODE	Not used in the current release
IS_LINK_EM_PAGE	Not used in the current release

19.6 Compliance Real-time Monitoring Views

This section provides a description of each compliance real-time monitoring view and its columns.

For examples of how to use these views, see [Section 19.30, "Examples"](#).

19.6.1 MGMT\$CCC_ALL_OBS_BUNDLES

The MGMT\$CCC_ALL_OBS_BUNDLES view returns a summary of all observation bundles. Any query against this view should ensure that filtering is done on appropriate fields with *bundle_start_time* being the first to take advantage of partitions.

Table 19–25 MGMT\$CCC_ALL_OBS_BUNDLES

Column	Description
BUNDLE_ID	The bundle to which this observation belongs based on the rule bundle settings
TARGET	Target against which this observation was found
TARGET_TYPE	Type of the target
RULE_NAME	Name of the real-time Monitoring Compliance Standard Rule

Table 19–25 (Cont.) MGMT\$CCC_ALL_OBS_BUNDLES

Column	Description
ENTITY_TYPE	Entity type of the entity that had an action against it
USER_PERFORMING_ACTION	Name of the user that performed the action
BUNDLE_IN_VIOLATION	Boolean value if the bundle is in violation currently. This means at least one observation in the bundle is unauthorized. True means bundle is in violation
BUNDLE_START_TIME	Date of the first observation in this bundle
BUNDLE_CLOSE_TIME	Date when this bundle was closed
BUNDLE_CLOSE_REASON	Explanation of why this bundle was closed
DISTINCT_OBS_COUNT	Total number of observations in this bundle
AUTHORIZED_OBS_COUNT	Number of observations in this bundle that are currently authorized
UNAUTHORIZED_OBS_COUNT	Number of observations in this bundle that are currently unauthorized
UNAUTH_CLEARED_OBS_COUNT	Number of observations in this bundle that are currently cleared (at one point they were unauthorized)
UNAUDITED_OBS_COUNT	Number of observations in this bundle that are currently unaudited. They have not been evaluated manually or with Change Management integration to determine audit status

19.6.2 MGMT\$CCC_ALL_OBSERVATIONS

The MGMT\$CCC_ALL_OBSERVATIONS view returns all observations that have occurred. Any query against this view should ensure that filtering is done on appropriate fields with *action_time* being the first to take advantage of partitions.

Table 19–26 MGMT\$CCC_ALL_OBSERVATIONS

Column	Description
OBSERVATION_ID	Unique ID given to the observation when detected by the agent
BUNDLE_ID	Bundle this observation belongs to based on rule bundle settings
TARGET	Target this observation was found against
TARGET_TYPE	Type of the target
ENTITY_TYPE	Entity type of the entity that had an action against it
ACTION	Action that was observed
ACTION_TIME	Time the action occurred
USER_TYPE	Type of user that performed the action (that is, OS user versus DB user)
USER_PERFORMING_ACTION	Name of the user that performed the action
ORIGINAL_USER_NAME	Previous user name in the case of a SU/SUDO action (only applicable to some entity types)
AFFECTED_ENTITY_NAME	Name of the entity that was affected by this action (file name, and so on)

Table 19–26 (Cont.) MGMT\$CCC_ALL_OBSERVATIONS

Column	Description
AFFECTED_ENTITY_PREVIOUS_NAME	Name of the entity prior to the action. For example, for file rename actions, this would be the old file name.
SOURCE_HOST_IP	Source IP of a connection when an action comes from another host (only applicable to some entity types)
ACTION_PROCESS_ID	Process ID of the process that performed the action (only applicable to some entity types)
ACTION_PROCESS_NAME	Name of the process that performed the action (only applicable to some entity types)
ACTION_PARENT_PROCESS_ID	Process ID of the parent process of the process that performed the action (only applicable to some entity types)
ACTION_PARENT_PROCESS_NAME	Name of the parent process of the process that performed the action (only applicable to some entity types)
ENTITY_PREVIOUS_VALUE	Previous value of the entity (only applicable to some entity types)
ENTITY_NEW_VALUE	New value of the entity (only applicable to some entity types)
FILE_ENTITY_PREVIOUS_MD5_HASH	Previous MD5 hash value of the entity (only applicable to some entity types)
FILE_ENTITY_NEW_MD5_HASH	New MD5 hash value of the entity (only applicable to some entity types)
AUDIT_STATUS	Current audit status of the observation (unaudited, authorized, unauthorized, and so on)
AUDIT_STATUS_SET_DATE	Date the most recent audit status was set
AUDIT_STATUS_SET_BY_USER	User who set the most recent audit status

19.6.3 MGMT\$CCC_ALL_VIOLATIONS

The MGMT\$CCC_ALL_VIOLATIONS view returns all real-time monitoring violations caused by an observation bundle having at least one unauthorized observation in it.

Table 19–27 MGMT\$CCC_ALL_VIOLATIONS

Column	Description
RULE_TYPE	Type of rule Possible values: <ul style="list-style-type: none"> Repository WebLogic Server Signature Real-time Monitoring
SEVERITY	Severity level of the rule <ul style="list-style-type: none"> Info Warning Critical

Table 19–27 (Cont.) MGMT\$CCC_ALL_VIOLATIONS

Column	Description
ENTITY_TYPE	Entity type of the observation bundle and all observations inside that bundle
TARGET_TYPE	Target type of the observation bundle and all observations inside that bundle
RULE_NAME	Name of the rule that this violation is against
COMPLIANCE_STANDARD_NAME	Name of the compliance standard that this violation is against.
TARGET	Name of the target that this violation is against.
BUNDLE_ID	Internal ID of the observation bundle that is in violation. This observation bundle has one or more unauthorized observations in it
BUNDLE_START_TIME	Time that the observation bundle started
BUNDLE_CLOSE_TIME	Time that the observation bundle closed
USER_NAME	User name that performed the actions in this bundle
AUTHORIZED_OBS_COUNT	Number of authorized observations in the observation bundle involved in this violation
UNAUTHORIZED_OBS_COUNT	Number of unauthorized observations in the observation bundle involved in this violation.
UNAUTH_CLEARED_OBS_COUNT	Number of unauthorized-cleared observations in the observation bundle involved in this violation
ROOT_CS_ID	Root compliance standard ID. This is used for the internal representation of the violation context.
RQS_ID	Runtime compliance standard ID. This is used for the internal representation of the violation context
RULE_ID	Internal ID of the rule with the violation.
TARGET_ID	Internal ID of the target with the violation.
ROOT_TARGET_ID	Internal ID of the target hierarchy.

19.6.4 MGMT\$COMPLIANT_TARGETS

The MGMT\$COMPLIANT_TARGETS view returns all evaluation and violation details for all targets. This is the same data that is shown in the Compliance Summary dashboard regions for targets.

Table 19–28 MGMT\$COMPLIANT_TARGETS

Column	Description
TARGET_NAME	Name of the target
TARGET_TYPE	Target type of the target
CRIT_EVALS	Number of critical-level evaluations
WARN_EVALS	Number of warning-level evaluations
COMPLIANT_EVALS	Number of compliant evaluations
CRIT_VIOLATIONS	Number of critical-level violations
WARN_VIOLATIONS	Number of warning-level violations

Table 19–28 (Cont.) MGMT\$COMPLIANT_TARGETS

Column	Description
MWARN_VIOLATIONS	Number of minor warning-level violations
COMPLIANCE_SCORE	Current compliance score for the target
TARGET_ID	Internal representation of the target
TARGET_TYPE_INAME	Internal representation of the target type

19.6.5 MGMT\$COMPLIANCE_SUMMARY

The MGMT\$COMPLIANCE_SUMMARY view returns all evaluation and violation details for compliance standards and frameworks. This is the same data that is shown in the Compliance Summary dashboard regions for compliance standards and frameworks.

Table 19–29 MGMT\$COMPLIANCE_SUMMARY

Column	Description
ELEMENT_TYPE	Type of element (compliance standard, compliance framework)
ELEMENT_NAME	Display name of the compliance standard or compliance framework
CRIT_EVALS	Number of critical-level evaluations
WARN_EVALS	Number of warning-level evaluations
COMPLIANT_EVALS	Number of compliant evaluations
CRIT_VIOLATIONS	Number of critical-level violations
WARN_VIOLATIONS	Number of warning-level violations
MWARN_VIOLATIONS	Number of minor warning-level violations
COMPLIANCE_SCORE	Current compliance score for the compliance standard or framework
NON_COMPLIANT_SCORE	Current non-compliant score for the compliance standard or framework
AUTHOR	Author of the compliance standard or framework
VERSION	Version of the compliance standard or framework
ELEMENT_ID	Internal ID of the compliance standard or compliance framework
FRAMEWORK_ID	Internal ID of the compliance framework
ELEMENT_INAME	Internal representation of the compliance standard or framework

19.6.6 MGMT\$COMPLIANCE_TREND

The MGMT\$COMPLIANCE_TREND view returns the last 31 days compliance trend information for compliance frameworks and standards. This is the same data that is shown in the Compliance Summary dashboard trend regions for compliance standards and frameworks.

Table 19–30 MGMT\$COMPLIANCE_TREND

Column	Description
ELEMENT_TYPE	Type of element (compliance standard, compliance framework)
ELEMENT_ID	Internal ID representation of the compliance standard or framework
ELEMENT_NAME	Display name of the compliance standard or compliance framework
FRAMEWORK_ID	Internal ID representation of the compliance framework
AVG_COMPLIANCE_SCORE	Average compliance score over the last 31 days
DAILY_AVG_VIOLATIONS	Average number of violations per day over the last 31 days
SNAPSHOT_TS	The snapshot time stamp
TOTAL_EVALS	Total evaluations over the last 31 days
ELEMENT_INAME	Internal representation of the compliance standard or framework

19.7 Configuration Management Views

This section provides a description of each configuration management view and its columns.

19.7.1 MGMT\$CSA_COLLECTIONS

The MGMT\$CSA_COLLECTIONS view displays top-level information about all client configurations.

Table 19–31 MGMT\$CSA_COLLECTIONS

Column	Description
DISPLAY_TARGET_NAME	The display name of the client
CSACLIENT	The display name plus the custom keys, if they exist
COLLECTION_TIMESTAMP	The time at which the data was collected from the client
NET_IP	The actual IP address of the client
NET_EFFECTIVE_IP	The client IP address seen by the server
COLLECTION_MESSAGE	Error message generated while applet was running
OS_USER_NAME	The client's OS user name
HOSTNAME	The client's host name
DOMAIN	The client's domain
BOOT_DISK_VOLUME_SERIAL_NUM	The client's boot disk volume serial number
COMPLIANCE	The overall compliance score for the client (15=passed, 18=info, 20=warning, 25=critical)
APPID	The collection tag for this client configuration
NET_SUBNET	The client's subnet mask
NET_LATENCY_IN_MS	The client's HTTP response time with the server

Table 19-31 (Cont.) MGMT\$CSA_COLLECTIONS

Column	Description
NET_BANDWIDTH_IN_KBITPS	The client's download bandwidth from the server
BROWSER_TYPE	The name of the browser used to run CSA
BROWSER_VERSION	The version of the browser used to run CSA
BROWSER	A summary column that combines the browser name and version
BROWSER_JVM_VENDOR	The version of the JVM used to run the applet
BROWSER_JVM_VERSION	The version of the JVM used to run the applet
BROWSER_PROXY_SERVER	The proxy server used by the browser
BROWSER_PROXY_EXCEPTIONS	The client's browser proxy exceptions
BROWSER_CACHE_SIZE_IN_MB	The client browser's disk cache size
BROWSER_CACHE_UPDATE_FRQ	The browser's cache update policy
BROWSER_HTTP1_1_SUPPORT	Whether or not the browser supports HTTP 1.1
REFERRING_URL_HEADER	The URL from which the user came to CSA, minus the query string
REFERRING_URL_PARAMS	The query string of the URL from which the user came to CSA
REFURL	The complete URL from which the user came to CSA
CSA_URL_HEADER	The URL from which the user ran CSA, minus the query string
CSA_URL_PARAMS	The query string of the URL from which the user ran CSA
CSAURL	The complete URL from which the user ran CSA
DESTINATION_URL_HEADER	The destination URL, minus the query string
DESTINATION_URL_PARAMS	The query string of the destination URL
DESTURL	The complete destination URL
CONNECTION_TYPE	The estimated connection type, based on the download bandwidth (1=LAN, 2=cable, 3=dialup)
IS_WINDOWS_ADMIN	Whether or not the os user is a Windows administrator
WINDOWS_DOMAIN	The Windows domain of the host
BROWSER_PROXY_ENABLED	Whether or not the proxy server is enabled in the browser
AUTO_CONFIG_URL	The URL of the proxy auto-configuration script used by the browser
NUMBER_OF_COOKIES	The number of cookies collected by CSA
NUMBER_OF_CUSTOM_VALUES	The number of custom properties collected by CSA
HARDWARE	A summary of the system configuration, machine architecture, memory, disk space, and CPU

Table 19–31 (Cont.) MGMT\$CSA_COLLECTIONS

Column	Description
HARDWARE_VENDOR_NAME	The name of the hardware vendor, such as Dell
SYSTEM_CONFIG	The client's system configuration
MACHINE_ARCHITECTURE	The client's machine architecture
BUS_FREQ_IN_MHZ	The frequency of the motherboard's Front Side Bus (FSB)
MEMORY_SIZE_IN_MB	The total amount of physical memory
AVAIL_MEMORY_SIZE_IN_MB	The amount of available physical memory when CSA was run
LOCAL_DISK_SPACE_IN_GB	The total amount of disk space
AVAIL_LOCAL_DISK_SPACE_IN_GB	The available disk space
CPU_COUNT	The number of CPUs
SYSTEM_SERIAL_NUMBER	The host's serial number
MIN_CPU_SPEED_IN_MHZ	The minimum possible CPU speed
MAX_CPU_SPEED_IN_MHZ	The maximum possible CPU speed
CPU	The CPU vendor, implementation, and frequency
CPU_BOARD_COUNT	The number of CPU boards
IOCARD_COUNT	The number of IO cards
NIC_COUNT	The number of NICs
FAN_COUNT	The number of fans
POWER_SUPPLY_COUNT	The number of power supplies
SYSTEM_BIOS	The system BIOS
OPERATINGSYSTEM	A summary of the OS name, version, update level, address length, and distributor version
OS_NAME	The OS name
OS_VENDOR_NAME	The OS vendor name
OS_BASE_VERSION	The OS base version
OS_UPDATE_LEVEL	The OS update level
OS_DISTRIBUTOR_VERSION	The OS distributor version
MAX_SWAP_SPACE_IN_MB	The maximum amount of swap space
OS_ADDRESS_LENGTH_IN_BITS	The OS address length in bits
MAX_PROCESS_VIRTUAL_MEMORY	The maximum amount of virtual memory that can be allocated to a process
TIMEZONE	The time zone as reported in the registry
TIMEZONE_REGION	The time zone region as reported by the JVM
TIMEZONE_DELTA	The offset in minutes from GMT

Table 19–31 (Cont.) MGMT\$CSA_COLLECTIONS

Column	Description
NUMBER_OF_OS_PROPERTIES	The number of OS properties found
NUMBER_OF_OS_PATCHES	The number of OS patches found
NUMBER_OF_OS_FILESYSTEMS	The number of file systems found
NUMBER_OF_OS_REGISTERED_SW	The number of OS-registered software products found
SNAPSHOT_ID	The GUID of this configuration
TARGET_ID	The GUID of the collector target
INTERNAL_TARGET_NAME	The internal name of the client configuration
INTERNAL_TARGET_TYPE	oracle_csa_client
COLLECTION_DURATION	The amount of time it took to run CSA
LOADED_TIMESTAMP	The time at which the data was loaded into the repository
APPLET_VERSION	The version of the applet
TARGET_ID_METHOD	not used
CUSTOM_CLASS	The name of the custom class (if any)
CUSTOM_CLASS_VERSION	not used
KEY1	Custom key 1(optional)
KEY2	Custom key 2 (optional)
KEY3	Custom key 3 (optional)
PROXY_TARGET_NAME	The name of the collector target
PROXY_TARGET_DISPLAY_NAME	The display name of the collector target
PROXY_TARGET_ID	The GUID of the collector target
RULES_COUNT	The total number of rules evaluated (including rules with status of NA)
RULES_NA_COUNT	The number of rules that were not applicable
RULES_PASSED_COUNT	The number of rules that passed
RULES_INFO_COUNT	The number of rules that failed with status info
RULES_WARNING_COUNT	The number of rules that failed with status warning
RULES_CRITICAL_COUNT	The number of rules that failed with status critical

19.7.2 MGMT\$CSA_FAILED

The MGMT\$CSA_FAILED view displays all failed collections.

Table 19–32 MGMT\$CSA_FAILED

Column	Description
ID	The GUID of this failed collection
TIMESTAMP	The time at which this failed collection occurred

Table 19–32 (Cont.) MGMT\$CSA_FAILED

Column	Description
TIMEZONE_DELTA	The offset in minutes from GMT
SAVED_TIMESTAMP	The time at which the data was loaded in the repository
EFFECTIVE_IP	The effective IP address of the client
APPID	The collection tag
REFERRING_URL_HEADER	The URL from which the user was referred to CSA, minus the query string
REFERRING_URL_PARAMS	The query string of the URL from which the user was referred to CSA
CSA_URL_HEADER	The URL from which the user ran CSA, minus the query string
CSA_URL_PARAMS	The query string of the URL from which the user tried to run CSA
DESTINATION_URL_HEADER	The destination URL minus the query string
DESTINATION_URL_PARAMS	The query string of the destination URL
BROWSER_TYPE	The type of browser used to run CSA
BROWSER_VERSION	The version of the browser used to run CSA
BROWSER_JVM_VENDOR	The vendor of the JVM used to run CSA
BROWSER_JVM_VERSION	The version of the JVM used to run CSA
OS_ARCH	The OS architecture of the client as reported in the <code>ios.arch</code> Java system property
OS_NAME	The OS name of the client as reported in the <code>ios.name</code> Java system property
HTTP_REQUEST_USER_AGENT	The HTTP user-Agent header sent by the client
ERROR_CODE	The error condition that caused the failed collection (0=OS not supported, 1=browser not supported, 2=applet certificate refused by user, 3=other error)
ERROR_TEXT	Text that is collected along with the error code, such as a stack trace

19.7.3 MGMT\$CSA_HOST_OS_COMPONENTS

The MGMT\$CSA_HOST_OS_COMPONENTS view displays all OS components found on CSA client systems.

Table 19–33 MGMT\$CSA_HOST_OS_COMPONENTS

Column	Description
DISPLAY_TARGET_NAME	The display name of the client
TYPE	The type of the component
NAME	The name of the component
VERSION	The version of the component
DESCRIPTION	The description of the component
INSTALLATION_DATE	The date the component was installed

Table 19–33 (Cont.) MGMT\$CSA_HOST_OS_COMPONENTS

Column	Description
SNAPSHOT_ID	The snapshot ID of the client configuration
TARGET_ID	The target GUID of the collector target
COLLECTION_TIMESTAMP	The time at which the client configuration was collected

19.7.4 MGMT\$CSA_HOST_SW

The MGMT\$CSA_HOST_SW view displays all OS-registered software found on CSA hosts.

Table 19–34 MGMT\$CSA_HOST_SW

Column	Description
DISPLAY_TARGET_NAME	The display name of the client
NAME	The name of the software
VENDOR_NAME	The name of the software vendor
VERSION	The version of the software
INSTALLATION_DATE	The date on which the software as installed
INSTALLED_LOCATION	The location in which the software is installed
DESCRIPTION	The description of the software
VENDOR_SW_SPECIFIC_INFO	Any additional information provided by the vendor
SNAPSHOT_ID	The snapshot ID of the client configuration
TARGET_ID	The ID of the collector target
COLLECTION_TIMESTAMP	The time at which the data was collected

19.7.5 MGMT\$CSA_HOST_COOKIES

The MGMT\$CSA_HOST_COOKIES view displays the cookies collected with client configurations.

Table 19–35 MGMT\$CSA_HOST_COOKIES

Column	Description
DISPLAY_TARGET_NAME	The display name of the client
NAME	The name of the cookie
VALUE	The payload of the cookie
SNAPSHOT_ID	The snapshot ID of the client configuration
TARGET_ID	The ID of the collector target
COLLECTION_TIMESTAMP	The time at which the data was collected

19.7.6 MGMT\$CSA_HOST_CUSTOM

The MGMT\$CSA_HOST_CUSTOM view displays the custom properties collected with client configurations.

Table 19–36 MGMT\$CSA_HOST_CUSTOM

Column	Description
DISPLAY_TARGET_NAME	The display name of the client
TYPE	The category of the custom property
NAME	The name of the custom property
TYPE_UI	The display category of the custom property
NAME_UI	The display name of the custom property
VALUE	The value of the custom property
DISPLAY_UI	Should this property be displayed in the UI? Y or N
HISTORY_TRACKING	Not used
SNAPSHOT_ID	The snapshot ID of the client configuration
TARGET_ID	The ID of the collector target
COLLECTION_TIMESTAMP	The time at which the data was collected

19.7.7 MGMT\$CSA_HOST_RULES

The MGMT\$CSA_HOST_RULES view displays the rules that were evaluated with each client configuration.

Table 19–37 MGMT\$CSA_HOST_RULES

Column	Description
SNAPSHOT_ID	The snapshot ID of the client configuration
NAME	The name of the rule
DESCRIPTION	The description of the rule
STATUS	The status of the rule (-2=NA, 15=passed, 18=info, 20=warning, 25=critical)
MOREINFO	Any additional information for the rule

19.7.8 MGMT\$CSA_HOST_CPUS

The MGMT\$CSA_HOST_CPUS view displays information about the CPUs of CSA hosts. CSA assumes that in a multi-CPU host, all CPUs are identical.

Table 19–38 MGMT\$CSA_HOST_CPUS

Column	Description
DISPLAY_TARGET_NAME	The display name of the client
VENDOR_NAME	The name of the CPU vendor
FREQ_IN_MHZ	The clock frequency of the CPU
ECACHE_IN_MB	The size of the extended cache
IMPL	The CPU implementation
REVISION	The CPU revision
MASK	The CPU mask
NUMBER_OF_CPUS	The number of CPUs

Table 19–38 (Cont.) MGMT\$CSA_HOST_CPUS

Column	Description
SNAPSHOT_ID	The snapshot ID of the client configuration
TARGET_ID	The ID of the collector target
COLLECTION_TIMESTAMP	The time at which the data was collected

19.7.9 MGMT\$CSA_HOST_IOCARDS

The MGMT\$CSA_HOST_IOCARDS view displays all IO cards collected from client configurations.

Table 19–39 MGMT\$CSA_HOST_IOCARDS

Column	Description
DISPLAY_TARGET_NAME	The display name of the client
VENDOR_NAME	The name of the IO card vendor
NAME	The name of the IO card
FREQ_IN_MHZ	The frequency of the IO card bus
BUS	The bus type (PCI or AGP)
REVISION	The IO card revision
NUMBER_OF_IOCARDS	The number of cards
SNAPSHOT_ID	The snapshot ID of the client configuration
TARGET_ID	The ID of the collector target
COLLECTION_TIMESTAMP	The time at which the data was collected

19.7.10 MGMT\$CSA_HOST_NICS

The MGMT\$CSA_HOST_NICS view displays all network interface cards collected from client configurations.

Table 19–40 MGMT\$CSA_HOST_NICS

Column	Description
DISPLAY_TARGET_NAME	The display name of the client
NAME	The name of the NIC
DESCRIPTION	The description of the NIC
FLAGS	Flags set on the NIC – not applicable for Windows
MAX_TRANSFER_UNIT	The maximum transfer unit of the NIC
INET_ADDRESS	The IP address of the NIC
MASK	The subnet mask of the NIC
BROADCAST_ADDRESS	The broadcast address of the NIC
MAC_ADDRESS	The MAC address of the NIC
HOSTNAME_ALIASES	Any aliases for the host name that are stored in the NIC
DEFAULT_GATEWAY	The default gateway for the NIC
DHCP_ENABLED	Whether or not DHCP is enabled

Table 19–40 (Cont.) MGMT\$CSA_HOST_NICS

Column	Description
SNAPSHOT_ID	The snapshot ID of the client configuration
TARGET_ID	The ID of the collector target
COLLECTION_TIMESTAMP	The time at which the data was collected

19.7.11 MGMT\$CSA_HOST_OS_PROPERTIES

The MGMT\$CSA_HOST_OS_PROPERTIES view displays all OS properties, such as environment variables, found on CSA hosts.

Table 19–41 MGMT\$CSA_HOST_OS_PROPERTIES

Column	Description
DISPLAY_TARGET_NAME	The display name of the client
SOURCE	The source (e.g. the system environment) of the property
NAME	The name of the property
VALUE	The value of the property
SNAPSHOT_ID	The snapshot ID of the client configuration
TARGET_ID	The ID of the collector target
COLLECTION_TIMESTAMP	The time at which the data was collected

19.7.12 MGMT\$CSA_HOST_OS_FILESYSTEMS

The MGMT\$CSA_HOST_OS_FILESYSTEMS view displays all file systems found on CSA hosts.

Table 19–42 MGMT\$CSA_HOST_OS_FILESYSTEMS

Column	Description
DISPLAY_TARGET_NAME	The display name of the client
RESOURCE_NAME	The name of the file system
MOUNT_LOCATION	The location from which it is mounted
TYPE	The file system type
DISK_SPACE_IN_GB	The total disk space
AVAIL_DISK_SPACE_IN_GB	The available disk space
LOCAL_DRIVE	The Windows drive letter on which it is mounted
MOUNT_OPTIONS	The mount options
SNAPSHOT_ID	The snapshot ID of the client configuration
TARGET_ID	The ID of the collector target
COLLECTION_TIMESTAMP	The time at which the data was collected

19.8 Custom Configuration Specification Views

This section provides a description of each custom configuration specification (CCS) view and its columns.

19.8.1 MGMT\$CCS_DATA

The MGMT\$CCS_DATA view provides both current and saved data, that is data saved from configurations.

Table 19–43 MGMT\$CCS_DATA

Column	Description
CM_TARGET_GUID	The unique ID for the target
CM_TARGET_TYPE	The type of the target
CM_TARGET_NAME	The name of the target
CM_SNAPSHOT_TYPE	Type of snapshot
CCS_UI_NAME	Display CCS name
CCS_DRAFT_NUMBER	Draft number of the CCS for draft CCSs. (It is 0 for non-draft CCS)
LAST_COLLECTION_TIMESTAMP	The timestamp of the collection specified in the target's time zone
ECM_SNAPSHOT_ID	The Enterprise Configuration Management (ECM) snapshot ID that can be used to join with other ECM views
DATA_SOURCE_NAME	Depending on the value for EXPR_TYPE, this is one of the following: <ul style="list-style-type: none"> ■ File name (relative to the base path) ■ OS command name ■ Database SQL query name
CONTAINER	A slash (/) separated hierarchal container with additional identification information and order information. This column could be a single space but only if the attribute name and value are available and are at the top level of the hierarchy.
ATTR	Attribute name
ATTR_ORDER	The order of the attribute within its enclosing container
CONTAINER_ORDER	The order of the container in the data source contents
VALUE	Attribute value

19.8.2 MGMT\$CCS_DATA_SOURCE

The MGMT\$CCS_DATA_SOURCE view contains both current and saved data (that is, data from saved configurations). This view can be joined with MGMT\$CCS_DATA_SOURCE_VISIBLE on ECM_SNAPSHOT_ID and DATA_SOURCE_NAME.

Table 19–44 MGMT\$CCS_DATA_SOURCE

Column	Description
CM_TARGET_GUID	The unique ID for the target
CM_TARGET_TYPE	The type of the target
CM_TARGET_NAME	The name of the target
CM_SNAPSHOT_TYPE	Type of snapshot
CCS_UI_NAME	Display CCS name
CCS_DRAFT_NUMBER	Draft number of the CCS for draft CCSs. (It is 0 for non-draft CCS)

Table 19–44 (Cont.) MGMT\$CCS_DATA_SOURCE

Column	Description
LAST_COLLECTION_TIMESTAMP	The time stamp of the collection specified in the target's time zone
ECM_SNAPSHOT_ID	The Enterprise Configuration Management (ECM) snapshot ID that can be used to join with other ECM views
DATA_SOURCE_NAME	Depending on the value for <code>EXPR_TYPE</code> , this is one of the following: <ul style="list-style-type: none"> File name (relative to the base path) OS command name Database SQL query name
EXPR_TYPE	The type of expression Possible values: <ul style="list-style-type: none"> F: Files O: OS commands D: Database queries
SOURCE_ORDER	Numeric order in which the source was obtained
EXPR_NAME	The name of the expression <ul style="list-style-type: none"> For files, this can be the wildcarded path expression from the custom configuration specification that caused the file to be collected For OS commands and database queries, this is a user-specified name for the expression
EXPR_VALUE	The value of the expression <ul style="list-style-type: none"> For files, this is the same as the value for <code>EXPR_NAME</code> For OS commands, this is the actual command For database queries, this is the actual database query
FULL_PATH	The full path <ul style="list-style-type: none"> For files, this is the full file path For OS commands, this is the base directory path
CONTENTS_SIZE	Byte size of contents
HASH	Hash value for collected data
CONTENTS	Character large object (CLOB) contents column with raw contents for this data source
COLLECTION_ERROR_MSG	Any relevant error message for this data source during the collection
PARSING_ERROR_MSG	Any relevant error messages generated during parsing of data contents

19.8.3 MGMT\$CCS_DATA_VISIBLE

The `MGMT$CCS_DATA_VISIBLE` view contains both current and saved data (that is, data from saved configurations).

Table 19–45 *MGMT\$CCS_DATA_VISIBLE*

Column	Description
TARGET_GUID	The unique ID for the target
TARGET_NAME	Name of the target
TARGET_TYPE	Type of target
SNAPSHOT_TYPE	Snapshot type
CCS_UI_NAME	Display CCS name
CCS_DRAFT_NUMBER	Draft number of the CCS for draft CCSs. (It is 0 for non-draft CCS)
DISPLAY_TARGET_NAME	User-friendly display name of the target
DISPLAY_TARGET_TYPE	User-friendly display name of the target type
COLLECTION_TIMESTAMP	Time stamp of the collection specified in the time zone of the target
IS_CURRENT	Specifies whether the data is current or saved Possible values: <ul style="list-style-type: none"> ■ Y: Current data ■ N: Saved data
DESCRIPTION	Snapshot description provided by the user
CREATOR	For saved snapshots, the creator is the Enterprise Manager user who saved the snapshot
SAVED_TIMESTAMP	Time stamp of when the snapshot was saved specified in the time zone of the database
LAST_UPLOAD_TIMESTAMP	Last time (specified in the time zone of the database) when a collection was processed for this snapshot type.
ECM_SNAPSHOT_ID	The Enterprise Configuration Management (ECM) snapshot ID that can be used to join with other ECM views
DATA_SOURCE_NAME	Depending on the value for EXPR_TYPE, this is one of the following: <ul style="list-style-type: none"> ■ File name (relative to the base path) ■ OS command name ■ Database SQL query name
CONTAINER	A slash (/) separated hierarchal container with additional identification information and order information. This column could be a single space but only if the attribute name and value are available and are at the top level of the hierarchy.
ATTRIBUTE	Attribute name
VALUE	Attribute value
CONTAINER_ORDER	The order of the container in the data source contents
ATTRIBUTE_ORDER	The order of the attribute within its enclosing container

19.8.4 MGMT\$CCS_DATA

The MGMT\$CCS_DATA view is the same as the [MGMT\\$CCS_DATA_VISIBLE](#) view but it exposes the current most recently collected data only.

Table 19–46 *MGMT\$CCS_DATA*

Column	Description
CM_TARGET_GUID	The unique ID for the target
CM_TARGET_TYPE	Type of the target
CM_TARGET_NAME	Name of the target
CM_SNAPSHOT_TYPE	Type of snapshot
CCS_UI_NAME	Display CCS name
CCS_DRAFT_NUMBER	Draft number of the CCS for draft CCSs. (It is 0 for non-draft CCS)
LAST_COLLECTION_TIMESTAMP	The time stamp of the collection specified in the target's time zone
ECM_SNAPSHOT_ID	The Enterprise Configuration Management (ECM) snapshot ID that can be used to join with other ECM views
DATA_SOURCE_NAME	Depending on the value for EXPR_TYPE, this is one of the following: <ul style="list-style-type: none"> ■ File name (relative to the base path) ■ OS command name ■ Database SQL query name
CONTAINER	A slash (/) separated hierarchal container with additional identification information and order information. This column could be a single space but only if the attribute name and value are available and are at the top level of the hierarchy.
ATTR	Attribute name
ATTR_ORDER	The order of the attribute within its enclosing container
CONTAINER_ORDER	The order of the container in the data source contents
VALUE	Attribute value

19.9 Database Configuration Views

This section provides a description of each database configuration view and its columns, along with examples about how to use the views.

19.9.1 MGMT\$DB_TABLESPACES

The MGMT\$DB_TABLESPACES view displays configuration settings for tablespaces. Tablespace settings are collected from the sys.dba_tablespaces, dba_free_space, dba_data_files, dba_temp_files, and v\$temp_extent_pool tables.

Table 19–47 *MGMT\$DB_TABLESPACES*

Column	Description
HOST_NAME	Name of the target where the metrics will be collected
TARGET_NAME	Name of the database containing the data files
TARGET_TYPE	The type of target, for example, Oracle_database
TARGET_GUID	The unique ID for the database target
COLLECTION_TIMESTAMP	The date and time when the metrics were collected

Table 19–47 (Cont.) MGMT\$DB_TABLESPACES

Column	Description
TABSPACE_NAME	Name of the tablespace
CONTENTS	Tablespace contents: PERMANENT or TEMPORARY
STATUS	Tablespace status: ONLINE, OFFLINE, or READ ONLY
EXTENT_MANAGEMENT	Extent management tracking: DICTIONARY or LOCAL
ALLOCATION_TYPE	Type of extent allocation in effect for this tablespace
LOGGING	Default logging attribute
TABSPACE_SIZE	Current size of the tablespace in bytes
INITIAL_EXT_SIZE	Default initial extent size
NEXT_EXTENT	Next extent in the sequence
INCREMENT_BY	Default percent increase for extent size
MAX_EXTENTS	Default maximum number of extents
TABSPACE_USED_SIZE	Amount of data (in bytes) contained in the tablespace
SEGMENT_SPACE_MANAGEMENT	Indicates whether the free and used segment space in the tablespace is managed using free lists (MANUAL) or bitmaps (AUTO)
BLOCK_SIZE	Tablespace block size
MIN_EXTENTS	Default minimum number of extents
MIN_EXTLEN	Minimum extent size for this tablespace
BIGFILE	Indicates whether the tablespace is a bigfile tablespace (YES) or a smallfile tablespace (NO)

19.9.2 MGMT\$DB_DATAFILES

The MGMT\$DB_DATAFILES view displays the configuration settings for data files. The data file settings are collected from sources such as sys.dba_data_files, v\$datafile, sys.dba_free_space, sys.dba_tablespaces, sys.dba_temp_files, v\$tempfile.

Table 19–48 MGMT\$DB_DATAFILES

Column	Description
HOST_NAME	Name of the target where the metrics will be collected
TARGET_NAME	Name of the database containing the data files
TARGET_TYPE	The type of target, for example, Oracle_database
TARGET_GUID	The unique ID for the database target
COLLECTION_TIMESTAMP	The date and time when the metrics were collected
FILE_NAME	Name of the data file
TABSPACE_NAME	Name of the tablespace containing the data file
STATUS	Data file status: ACTIVE or NOT ACTIVE
FILE_SIZE	Size of the data file
AUTOEXTENSIBLE	Autoextensible indicator
INCREMENT_BY	Autoextension increment

Table 19–48 (Cont.) MGMT\$DB_DATAFILES

Column	Description
MAX_FILE_SIZE	Maximum file size in bytes
OS_STORAGE_ENTITY	OS level storage entity on which the file resides. For regular files it is the name of the file system on which the file resides. For character or raw files it is the name of the raw device
CREATE_BYTES	The initial size of the data file when it was created in bytes

19.9.3 MGMT\$DB_CONTROLFILES

The MGMT\$DB_CONTROLFILES view displays the configuration settings for database control files.

Table 19–49 MGMT\$DB_CONTROLFILES

Column	Description
HOST_NAME	Name of the target where the metrics will be collected
TARGET_NAME	Name of the database containing the data files
TARGET_TYPE	The type of target, for example, Oracle_ database
TARGET_GUID	The unique ID for the database target
COLLECTION_TIMESTAMP	The date and time when the metrics were collected
FILE_NAME	Name of the database control file.
STATUS	The type of control file: STANDBY - indicates database is in standby mode LOGICAL - indicates the database is a logical standby database (not a physical standby) CLONE - indicates a clone database BACKUP CREATED - indicates database is being recovered using a backup or created control file CURRENT - the control file changes to this type following a standby database activate or database open after recovery
CREATION_DATE	Control file creation date
SEQUENCE_NUM	Control file sequence number incremented by control file transactions
CHANGE_NUM	Last change number in the backup control file. Value is NULL if the control file is not a backup
MOD_DATE	Last timestamp in the backup control file. NULL if the control file is not a backup
OS_STORAGE_ENTITY	OS level storage entity on which the file resides. For regular files it is the name of the file system on which the file resides. For character or raw files it is the name of the raw device

19.9.4 MGMT\$DB_DBNINSTANCEINFO

The MGMT\$DB_DBNINSTANCEINFO view displays general information about database instance. The instance information is collected from v\$database, v\$version, v\$instance, global_name, database_properties and v\$nls_parameters.

Table 19–50 MGMT\$DB_DBNINSTANCEINFO

Column	Description
HOST_NAME	Name of the target host where the metrics will be collected
TARGET_NAME	Name of the database target from which the metrics are collected
TARGET_TYPE	The type of target, for example, Oracle_database
TARGET_GUID	The unique ID for the database target
COLLECTION_TIMESTAMP	The date and time when the metrics were collected
DATABASE_NAME	Name of the database
GLOBAL_NAME	Global name of the database
BANNER	Component name and version number
HOST	Name of the host system
INSTANCE_NAME	Name of the instance
STARTUP_TIME	Time when instance was started up
LOGINS	ALLOWED or RESTRICTED
LOG_MODE	The archive log mode, either ARCHIVELOG or NOARCHIVELOG
OPEN_MODE	Open mode information
DEFAULT_TEMP_TABLESPACE	Default temporary tablespace name
CHARACTERSET	NLS parameter value for NLS_CHARACTERSET
NATIONAL_CHARACTERSET	NLS parameter value for NLS_NCHAR_CHARACTERSET

Usage Notes

This information is collected through the dbconfig metric. However, as this metric is not run for standby databases, this table is not populated for standby targets.

19.9.5 MGMT\$DB_FEATUREUSAGE

The MGMT\$DB_FEATUREUSAGE view displays information about database feature usage.

Table 19–51 MGMT\$DB_FEATUREUSAGE

Column	Description
HOST	Name of the host target where the database feature usage information is collected
DATABASE_NAME	Name of the database where the database feature usage information is collected
INSTANCE_NAME	Name of the instance where the database feature usage information is collected
TARGET_TYPE	Either Oracle_database or rac_database
DBID	A unique number that identifies a database instance
NAME	The feature name

Table 19–51 (Cont.) MGMT\$DB_FEATUREUSAGE

Column	Description
CURRENTLY_USED	TRUE if the feature is currently in use, FALSE if the feature is not in use
DETECTED_USAGES	The number of times the feature has been used by the database
FIRST_USAGE_DATE	The date that the first usage of the feature occurred
LAST_USAGE_DATE	The date of the most recent usage of the feature
VERSION	The version number of the database
LAST_SAMPLE_DATE	The date that the database was last evaluated for feature usage
LAST_SAMPLE_PERIOD	The interval between the LAST_SAMPLE_DATE date and the database feature usage evaluation before that (by default, seven days)
SAMPLE_INTERVAL	The number of seconds between the LAST_SAMPLE_DATE date and the next database feature usage evaluation
TOTAL_SAMPLES	The total number of database feature usage evaluation samples that have been collected
AUX_COUNT	For Oracle internal use only
DESCRIPTION	The description of the feature

Usage Notes

This view can be used to gain an enterprise-wide view of database feature usage across all Oracle databases.

19.9.6 MGMT\$DB_INIT_PARAMS

The MGMT\$DB_INIT_PARAMS view displays initialization parameter settings for the database. Initialization parameter settings are collected from v\$parameter.

Table 19–52 MGMT\$DB_INIT_PARAMS

Column	Description
HOST_NAME	Name of the target where the metrics will be collected
TARGET_NAME	Name of the database target from which the metrics are collected
TARGET_TYPE	The type of target, such as Oracle_database
TARGET_GUID	The unique ID for the database target
COLLECTION_TIMESTAMP	The date and time when the metrics were collected
NAME	Name of the initialization parameter
ISDEFAULT	Indicates whether the parameter value is the default
VALUE	The parameter value
DATATYPE	The data type that the value string can be mapped to, for example, NUMBER, DATE, or TEXT

Usage Notes

This information is collected through the dbconfig metric. However, as this metric is not run for standby databases, this table is not populated for standby targets.

19.9.7 MGMT\$DB_LICENSE

The MGMT\$DB_LICENSE view displays database license configuration settings. Database license configuration settings are collected from v\$license.

Table 19–53 MGMT\$DB_LICENSE

Column	Description
HOST_NAME	The name of the host on which the database is running
TARGET_NAME	Name of the database containing the tablespace
TARGET_TYPE	The type of target, for example, Oracle_ database
TARGET_GUID	The unique ID for the database target
COLLECTION_TIMESTAMP	The date and time when the metrics were collected
SESSIONS_MAX	The maximum number of sessions allowed for the database
SESSIONS_WARNING	The number of sessions which will generate a warning for the database
SESSIONS_CURRENT	The current number of sessions for the database
SESSIONS_HIGHWATER	The highest water mark of sessions for the database
USERS_MAX	The maximum number of users for the database

Usage Notes

This view can be used to obtain database license configuration settings across all database targets.

19.9.8 MGMT\$DB_REDOLOGS

The MGMT\$DB_REDOLOGS view displays redo log configuration settings for the database. Redo log configuration settings are collected from the v\$log and v\$logfile tables.

Table 19–54 MGMT\$DB_REDOLOGS

Column	Description
HOST_NAME	Name of the target where the metrics will be collected
TARGET_NAME	Name of the database target from which the metrics are collected
TARGET_TYPE	The type of target, for example, Oracle_ database
TARGET_GUID	The unique ID for the database target
COLLECTION_TIMESTAMP	The date and time when the metrics were collected
GROUP_NUM	Redo log group identifier number

Table 19–54 (Cont.) MGMT\$DB_REDOLOGS

Column	Description
STATUS	<p>Log status:</p> <p>UNUSED - The online redo log has never been written to. This is the state of a redo log that was just added, or just after a RESETLOGS, when it is not the current redo log.</p> <p>CURRENT - This is the current redo log. This implies that the redo log is active. The redo log could be open or closed.</p> <p>ACTIVE - The log is active but is not the current log. It is needed for crash recovery. It may be in use for block recovery. It might or might not be archived.</p> <p>CLEARING - The log is being re-created as an empty log after an ALTER DATABASE CLEAR LOGFILE statement. After the log is cleared, the status changes to UNUSED.</p> <p>CLEARING_CURRENT - The current log is being cleared of a closed thread. The log can stay in this status if there is some failure in the switch such as an I/O error writing the new log header.</p> <p>INACTIVE - The log is no longer needed for instance recovery. It may be in use for media recovery. It might or might not be archived.</p>
MEMBERS	Number of members in the log group
FILE_NAME	Redo log file (member) name
ARCHIVED	Archive status either YES or NO
LOGSIZE	Size of the log file in bytes
SEQUENCE_NUM	Log sequence number
FIRST_CHANGE_SCN	Lowest SCN in the log
OS_STORAGE_ENTITY	OS level storage entity on which the file resides. For regular files it is the name of the file system on which the file resides. For character or raw files it is the name of the raw device.
THREAD_NUM	Log thread number

Usage Notes

Obtain redo log group or file configuration settings across all database targets.

19.9.9 MGMT\$DB_ROLLBACK_SEGS

The MGMT\$DB_ROLLBACK_SEGS view displays rollback segments configuration settings for the database. Rollback segments configuration settings are collected from the sys.dba_rollback_segs and v\$rollstat tables.

Table 19–55 MGMT\$DB_ROLLBACK_SEGS

Column	Description
HOST_NAME	Name of the target where the metrics will be collected
TARGET_NAME	Name of the database containing the data files
TARGET_TYPE	The type of target, for example, Oracle_ database
TARGET_GUID	The unique ID for the database target
COLLECTION_TIMESTAMP	The date and time when the metrics were collected

Table 19–55 (Cont.) MGMT\$DB_ROLLBACK_SEGS

Column	Description
ROLLNAME	Name of the rollback segment
STATUS	Rollback segment status
TABLESPACE_NAME	Name of the tablespace containing the rollback segment
EXTENTS	Number of extents in rollback segment
ROLLSIZE	Size in bytes of rollback segment. This values differs by the number of bytes in one database block from the value of the BYTES column of the ALL/DBA/USER_SEGMENTS views.
INITIAL_SIZE	Initial extent size in bytes
NEXT_SIZE	Secondary extent size in bytes
MAXIMUM_EXTENTS	Maximum number of extents
MINIMUM_EXTENTS	Minimum number of extents
PCT_INCREASE	Percent increase for extent size
OPTSIZE	Optimal size for rollback segments
AVEACTIVE	Current size of active extents averaged over time
WRAPS	Number of times rollback segment is wrapped
SHRINKS	Number of times the size of a rollback segment decreases
AVESHRINK	Average shrink size
HWMSIZE	High water mark of rollback segment size

Usage Notes

Obtain rollback segments configuration settings across all database targets.

19.9.10 MGMT\$DB_SGA

The MGMT\$DB_SGA view displays System Global Area (SGA) configuration settings. SGA settings are collected from the v\$sga and v\$sgastat tables.

Table 19–56 MGMT\$DB_SGA

Column	Description
HOST_NAME	Name of the target where the metrics will be collected
TARGET_NAME	Name of the database containing the datafiles
TARGET_TYPE	The type of target, for example, Oracle_ database
TARGET_GUID	The unique ID for the database target
COLLECTION_TIMESTAMP	The date and time when the metrics were collected
SGANAME	SGA component name
SGASIZE	SGA component size in kilobytes or megabytes

Usage Notes

Obtain System Global Area configuration settings across all database targets.

19.9.11 MGMT\$DB_TABLESPACES_ALL

The MGMT\$DB_TABLESPACES_ALL view displays configuration settings for tablespaces. Tablespace settings are collected from the sys.dba_tablespaces, dba_free_space, dba_data_files, dba_temp_files, and v\$temp_extent_pool tables.

Table 19–57 MGMT\$DB_TABLESPACES_ALL

Column	Description
HOST_NAME	Name of the target where the metrics will be collected
TARGET_NAME	The type of target, for example, Oracle_ database
TARGET_GUID	The unique ID for the database target
COLLECTION_TIMESTAMP	The date and time when the metrics were collected
TABLESPACE_NAME	Name of the tablespace
CONTENTS	Tablespace contents: PERMANENT or TEMPORARY
STATUS	Tablespace status: ONLINE, OFFLINE, or READ ONLY
EXTENT_MANAGEMENT	Extent management tracking: DICTIONARY or LOCAL
ALLOCATION_TYPE	Type of extent allocation in effect for this tablespace
LOGGING	Default logging attribute
TABLESPACE_SIZE	Current size of the tablespace in bytes
INITIAL_EXT_SIZE	Default initial extent size
INCREMENT_BY	Default percent increase for extent size
MAX_EXTENTS	Default maximum number of extents

Usage Notes

Obtain tablespace configuration settings across all database targets.

19.9.12 MGMT\$DB_OPTIONS

The MGMT\$DB_OPTIONS view displays whether or not the option is currently LOADED and ACTIVE, or either the option does not exist or is NOT LOADED or INACTIVE. Options settings are collected by checking user name and status in the sys.dba_users and dba_registry tables.

Table 19–58 MGMT\$DB_OPTIONS

Column	Description
HOST_NAME	Name of the target where the metrics will be collected
TARGET_NAME	Name of the database containing the data files
TARGET_TYPE	The type of target, for example, Oracle_ database
TARGET_GUID	The unique ID for the database target
COLLECTION_TIMESTAMP	The date and time when the metrics were collected
NAME	Name of the database option
SELECTED	If the option is currently LOADED and ACTIVE (TRUE), or either the option does not exist or is NOT LOADED or INACTIVE (FALSE)

19.10 Events Views

This section provides a description of each event view and its columns.

For examples of how to use views, see [Section 19.30, "Examples"](#).

19.10.1 MGMT\$INCIDENTS

The MGMT\$INCIDENTS view provides a view of the attributes of the incident including its summary message.

Table 19–59 MGMT\$INCIDENTS

Column	Description
INCIDENT_ID	The unique RAW ID of an incident
INCIDENT_NUM	The end-user visible ID of the incident
SUMMARY_MSG	Summary message of the incident
SEVERITY	The severity of the incident
IS_ESCALATED	Specifies whether the issue is escalated. Possible values: <ul style="list-style-type: none"> 1: Yes 0: No
ESCALATION_LEVEL	If the incident is escalated, then this value specifies the escalation level. This value can be between level 1 and level 5.
PRIORITY	The priority level of the incident. Possible values: <ul style="list-style-type: none"> None Urgent Very High High Medium Low
RESOLUTION_STATE	The resolution state of the issue.
OWNER	The owner of the issue. If there is no owner, then this value is "-".
IS_ACKNOWLEDGED	Specifies whether the incident is acknowledged. Possible values: <ul style="list-style-type: none"> 1: Yes 0: No
IS_SUPPRESSED	Specifies whether the incident is suppressed. Possible values: <ul style="list-style-type: none"> 1: Yes 0: No
LAST_ANNOTATION_SEQ	The sequence ID of the last annotation entered for this issue
CREATION_DATE	The date the incident was created
LAST_UPDATED_DATE	The date when this incident was updated last
EVENT_COUNT	The number of events associated with this incident

Table 19–59 (Cont.) MGMT\$INCIDENTS

Column	Description
OPEN_STATUS	Specifies the status of the incident. Possible values: <ul style="list-style-type: none"> ■ 1: Open incidents ■ 0: Closed incidents
CLOSED_DATE	The date when the incident is closed (if it is closed)
SRC_COUNT	The number of unique target or source object combinations to which events in this incident belong
TARGET_GUID	The unique ID of a target associated with the incident This value is set only when all the events in the incident belong to the same target or source object combination. It is set to null when the events belong to multiple sources.
SOURCE_OBJ_TYPE	The source object or entity type to which all events in the incident belong (if they all belong to the same target or source object combination). Set to null when the events belong to multiple sources.
ADR_RELATED	Indicates if the incident is a Oracle diagnostic incident. Possible values: <ul style="list-style-type: none"> ■ 0: No ■ 1: Yes
TICKET_ID	Ticket associated with this incident (can be null)
TICKET_STATUS	Status of the ticket associated with this incident (can be null)
SR_ID	ID of the service request associated with this problem (if any)
PROBLEM_ID	The unique RAW ID of the related problem (if any)
PROBLEM_NUM	The end-user visible ID of the related problem (if any)

19.10.2 MGMT\$INCIDENT_CATEGORY

The MGMT\$INCIDENT_CATEGORY view is the incident view for the mapping between incidents and categories. An incident can have multiple categories associated with it.

Table 19–60 MGMT\$INCIDENT_CATEGORY

Column	Description
INCIDENT_ID	The unique RAW ID of an incident
CATEGORY_NAME	Name of the category
OPEN_STATUS	Specifies the status of the incident. Possible values: <ul style="list-style-type: none"> ■ 1: Open incidents ■ 0: Closed incidents
CLOSED_DATE	The date when the incident is closed

19.10.3 MGMT\$INCIDENT_TARGET

The MGMT\$INCIDENT_TARGET view is the incident view for the mapping between incidents and targets. An incident can be made of multiple events and these events could be from different targets.

Table 19–61 MGMT\$INCIDENT_TARGET

Column	Description
INCIDENT_ID	The unique RAW ID of an incident
TARGET_GUID	The unique ID of a target (can be null)
OPEN_STATUS	Specifies the status of the incident. Possible values: <ul style="list-style-type: none"> ■ 1: Open incidents ■ 0: Closed incidents
CLOSED_DATE	The date when the incident is closed

19.10.4 MGMT\$INCIDENT_ANNOTATION

The MGMT\$INCIDENT_ANNOTATION view is the view for the mapping between the incidents and annotations. Each incident can have multiple annotations.

Table 19–62 MGMT\$INCIDENT_ANNOTATION

Column	Description
INCIDENT_ID	The unique RAW ID of an incident
ANNOTATION_SEQ	The order ID in which the annotation is added
ANNOTATION_MSG	The annotation message
ANNOTATION_DATE	The time stamp when the annotation is made
ANNOTATION_TYPE	The type of the annotation, that is, whether it is user or system generated. Possible values: <ul style="list-style-type: none"> ■ USER ■ SYSTEM
ANNOTATION_USER	The user that added the annotation. If the annotation is system-generated, then this value is set to "-".
OPEN_STATUS	Specifies the status of the incident. Possible values: <ul style="list-style-type: none"> ■ 1: Open incidents ■ 0: Closed incidents
CLOSED_DATE	The date when the incident is closed

19.10.5 MGMT\$EVENTS_LATEST

The MGMT\$EVENTS_LATEST view shows the details of the latest state of all events in a given sequence of events. A sequence is a series of raw events that are related to the same source and reporting on the same issue. For example, for a given host, if the CPU utilization goes from warning to critical and then to warning again, then these three events are correlated into a single sequence with three raw events with warning as the latest state.

Table 19–63 *MGMT\$EVENTS_LATEST*

Column	Description
EVENT_SEQ_ID	The unique RAW ID of an event sequence
EVENT_ID	The unique RAW ID of the latest event in the sequence
EVENT_CLASS	The event class to which this event belongs
SEVERITY	The severity of the event
LAST_ANNOTATION_SEQ	The sequence ID of the last annotation entered for this sequence
MSG	The event message
EVENT_NAME	The internal event name describing the nature of the events
INCIDENT_ID	The incident ID to which this event belongs (if any)
INCIDENT_NUM	The end-user readable number or ID for the incident
TARGET_GUID	The target GUID to which the events of the sequence belong. If the sequence is not related to any target, then this value is set to NULL.
SOURCE_OBJ_TYPE	The source object or entity type to which the events of the sequence belong. Default value is NULL
SOURCE_OBJ_ID	The source object or entity GUID to which the events of the sequence belong. Default value is NULL
OPEN_STATUS	The status of the event sequence. The event sequence is considered open if the severity of the last event is a non-clear severity. Possible values: <ul style="list-style-type: none"> ■ 1: Open ■ 0: Closed
CLOSED_DATE	The date when the event is marked as closed, that is, when the event sequence is cleared
CREATION_DATE	The date the event sequence was created
LAST_UPDATED_DATE	The date when this event sequence was last updated

19.10.6 MGMT\$EVENTS

The MGMT\$EVENTS view shows the details of all the raw events in a given sequence of events. A sequence is a series of raw events that are related to the same source and reporting on the same issue. For example, for a given host, if the CPU utilization goes from warning to critical and then to warning again, then these three events are correlated into a single sequence with three raw events with warning as the latest state.

Table 19–64 *MGMT\$EVENTS*

Column	Description
EVENT_SEQ_ID	The unique RAW ID of the event sequence
EVENT_ID	The unique RAW ID of the event
SIGNATURE_ID	The ID of the unique signature of raw events that is used to correlate all raw events to a sequence
EVENT_CLASS	The event class to which this event belongs
SEVERITY	The severity of the raw event

Table 19–64 (Cont.) MGMT\$EVENTS

Column	Description
LAST_ANNOTATION_SEQ	The sequence ID of the last annotation entered for this sequence
MSG	The event message
EVENT_NAME	The internal event name describing the nature of the events
INCIDENT_ID	The incident ID to which this event belongs (if applicable)
INCIDENT_NUM	The end-user readable number or ID for the incident
TARGET_GUID	The target GUID to which the event belongs. If the sequence is not related to any target, then this value is set to NULL.
SOURCE_OBJ_TYPE	The source object or entity type to which the event belongs. Default value is NULL
SOURCE_OBJ_ID	The source object or entity GUID to which the event belongs. Default value is NULL
OPEN_STATUS	The status of the event sequence. The event sequence is considered open if the severity of the last event is a non-clear severity. Possible values: <ul style="list-style-type: none"> ■ 1: Open ■ 0: Closed
CLOSED_DATE	The date when the event is marked as closed, that is, when the event sequence is cleared
REPORTED_DATE	The date when the event was reported

19.10.7 MGMT\$EVENT_ANNOTATION

The MGMT\$EVENT_ANNOTATION view is the view for the mapping between events and annotations. Each event can have multiple annotations.

Note: Annotations are associated with the sequence and *not* with the individual raw events

Table 19–65 MGMT\$EVENT_ANNOTATION

Column	Description
EVENT_SEQ_ID	The unique RAW ID of an event sequence
EVENT_INSTANCE_ID	The unique RAW ID of an event instance
ANNOTATION_SEQ_NUM	The order ID in which the annotations is added
ANNOTATION_DATE	The time stamp when the annotation is made
ANNOTATION_TYPE	The type of the annotation, that is, whether it is user or system generated. Possible values: <ul style="list-style-type: none"> ■ USER ■ SYSTEM
ANNOTATION_USER	The user which added the annotation
ANNOTATION_MSG	The annotation message

Table 19–65 (Cont.) MGMT\$EVENT_ANNOTATION

Column	Description
OPEN_STATUS	The status of the event sequence. The event sequence is considered open if the severity of the last event is a non-clear severity. Possible values: <ul style="list-style-type: none"> ■ 1: Open ■ 0: Closed
CLOSED_DATE	The date when the event is marked as closed, that is, when the event sequence is cleared

19.10.8 MGMT\$PROBLEMS

The MGMT\$PROBLEMS view provides a view of the attributes of the problem including its summary message.

Table 19–66 MGMT\$PROBLEMS

Column	Description
PROBLEM_ID	The unique RAW ID of the problem
PROBLEM_NUM	The end-user visible ID of the problem
SUMMARY_MSG	Summary message of the problem
SEVERITY	The severity of the problem
IS_ESCALATED	Specifies whether the issue is escalated. Possible values: <ul style="list-style-type: none"> ■ 1: Yes ■ 0: No
ESCALATION_LEVEL	If the problem is escalated, then this value specifies the escalation level. This value can be between level 1 and level 5.
PRIORITY	The priority level of the incident. Possible values: <ul style="list-style-type: none"> ■ None ■ Urgent ■ Very High ■ High ■ Medium ■ Low
RESOLUTION_STATE	The resolution state of the issue
OWNER	The owner of the issue. If there is no owner, then this value is "-".
IS_ACKNOWLEDGED	Specifies whether the problem is acknowledged. Possible values: <ul style="list-style-type: none"> ■ 1: Yes ■ 0: No

Table 19–66 (Cont.) MGMT\$PROBLEMS

Column	Description
IS_SUPPRESSED	Specifies whether the problem is suppressed. Possible values: <ul style="list-style-type: none"> ■ 1: Yes ■ 0: No
LAST_ANNOTATION_SEQ	The sequence ID of the last annotation entered for this issue
CREATION_DATE	The date the problem was created
LAST_UPDATED_DATE	The date when this problem was updated last
INC_COUNT	The number of incidents associated with this problem
OPEN_STATUS	Specifies the status of the problem. Possible values: <ul style="list-style-type: none"> ■ 1: Open ■ 0: Closed
CLOSED_DATE	The date when the problem is closed (if it is closed)
TARGET_GUID	The unique ID of a target (can be null). This value is set only when all the incidents in the problem belong to the same target or source object combination. It is set to null when the incidents belong to multiple sources. Note: For this release, problems can be associated with a single target only
PROBLEM_KEY	Unique signature of this problem
SR_ID	ID of the service request associated with this problem, if any
BUG_ID	ID of the bug associated with this problem, if any

19.10.9 MGMT\$PROBLEM_ANNOTATION

The MGMT\$PROBLEM_ANNOTATION view is the view for the mapping between problems and annotations. Each problem can have multiple annotations.

Table 19–67 MGMT\$PROBLEM_ANNOTATION

Column	Description
PROBLEM_ID	The unique RAW ID of a problem
ANNOTATION_SEQ	The order ID in which the annotations is added
ANNOTATION_MSG	The annotation message
ANNOTATION_DATE	The time stamp when the annotation is made
ANNOTATION_TYPE	The type of the annotation, either user or system generated. Valid values are 'USER' or 'SYSTEM'
ANNOTATION_USER	The user which added the annotation
OPEN_STATUS	Specifies the status of the problem. Possible values: <ul style="list-style-type: none"> ■ 1: Open ■ 0: Closed
CLOSED_DATE	The date when the problem is closed

19.11 Glassfish Views

This section provides a description of each event view and its columns.

For examples of how to use views, see [Section 19.30, "Examples"](#).

19.11.1 MGMT\$EMAS_GLASSFISH_DOMAIN

The MGMT\$EMAS_GLASSFISH_DOMAIN view displays general information about the glassfish domain target.

Table 19–68 MGMT\$EMAS_GLASSFISH_DOMAIN

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target.
CM_TARGET_TYPE	The type of target: glassfish_domain
CM_TARGET_NAME	The name of target in Enterprise Manager
CM_SNAPSHOT_TYPE	The type of snapshot which collected the config metric
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
DOMAINNAME	Name of the domain
FULLVERSION	Version of the domain
CONFIGDIR	Config directory of the domain
INSTALLDIR	Install directory of the domain
DEBUGPORT	Debug port of the node
SECUREADMINENABLED	Secure admin enabled or not

19.11.2 MGMT\$EMAS_GLASSFISH_NODES

The MGMT\$EMAS_GLASSFISH_NODES view displays general information about the nodes configured on a glassfish domain target.

Table 19–69 MGMT\$EMAS_GLASSFISH_NODES

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target.
CM_TARGET_TYPE	The type of target: glassfish_domain
CM_TARGET_NAME	The name of target in Enterprise Manager
CM_SNAPSHOT_TYPE	The type of snapshot which collected the config metric
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
NODENAME	Name of the node
NODEDIR	Directory of the node
NODEHOST	Host of the node
TYPE	Type of the node
INSTALLDIR	Install directory of the node

19.11.3 MGMT\$EMAS_GLASSFISH_SERVER

The MGMT\$EMAS_GLASSFISH_SERVER view displays general information about the glassfish server target.

Table 19–70 MGMT\$EMAS_GLASSFISH_SERVER

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target.
CM_TARGET_TYPE	The type of target: glassfish_j2eeserver
CM_TARGET_NAME	The name of target in Enterprise Manager
CM_SNAPSHOT_TYPE	The type of snapshot which collected the config metric
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
SERVERNAME	Name of the server
HOST	Host of the server
SERVERVERSION	The version of the server
SERVICEURL	The ServiceURL of the server
LISTEPORT	The listenport of the server
JAVAVENDOR	The name of the java vendor
JAVAVERSION	The version of the java

19.11.4 MGMT\$EMAS_GLASSFISH_SVR_PROP

The MGMT\$EMAS_GLASSFISH_SVR_PROP view displays the system properties of the glassfish server target.

Table 19–71 MGMT\$EMAS_GLASSFISH_SVR_PROP

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target.
CM_TARGET_TYPE	The type of target: glassfish_j2eeserver
CM_TARGET_NAME	The name of target in Enterprise Manager
CM_SNAPSHOT_TYPE	The type of snapshot which collected the config metric
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
NAME	The system property name
TYPE	The system property value

19.11.5 MGMT\$EMAS_GLASSFISH_NW_LSTNR

The MGMT\$EMAS_GLASSFISH_NW_LSTNR view displays the network listeners of the glassfish server target.

Table 19–72 MGMT\$EMAS_GLASSFISH_NW_LSTNR

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target.
CM_TARGET_TYPE	The type of target: glassfish_j2eeserver
CM_TARGET_NAME	The name of target in Enterprise Manager
CM_SNAPSHOT_TYPE	The type of snapshot which collected the config metric
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
NAME	The name of the network listener
PROTOCOL	The protocol of the network listener
TRANSPORT	The name of the transport
ADDRESS	The address of the listener
ENABLED	The enabled flag of the network listener
SECURITYENABLED	The security enabled flag of the network listener
PORT	The port of the network listener

19.11.6 MGMT\$EMAS_GLASSFISH_DATASOURCE

The MGMT\$EMAS_GLASSFISH_DATASOURCE view displays information about the JDBC Datasource.

Table 19–73 MGMT\$EMAS_GLASSFISH_DATASOURCE

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target.
CM_TARGET_TYPE	The type of target: glassfish_j2eeserver
CM_TARGET_NAME	The name of target in Enterprise Manager
CM_SNAPSHOT_TYPE	The type of snapshot which collected the config metric
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
JDBC_DATASOURCE_NAME	The name of the JDBC datasource
JNDI_NAME	The JNDI name of the datasource
POOL_NAME	The name of the connection pool
MIN_POOL_SIZE	The initial and minimum pool size
MAX_POOL_SIZE	The maximum pool size
POOL_RESIZE_QUANTITY	The pool resize quantity
STATEMENT_CACHE_SIZE	The statement cache size
IDLE_TIMEOUT	The idle timeout

19.11.7 MGMT\$EMAS_GLASSFISH_DS_PROP

The MGMT\$EMAS_GLASSFISH_DS_PROP view displays the properties of a JDBC datasource for the glassfish server target.

Table 19–74 MGMT\$EMAS_GLASSFISH_DATASOURCE

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target.
CM_TARGET_TYPE	The type of target: glassfish_j2eeserver
CM_TARGET_NAME	The name of target in Enterprise Manager
CM_SNAPSHOT_TYPE	The type of snapshot which collected the config metric
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
JDBC_DATASOURCE_NAME	The name of the JDBC datasource
NAME	The property name
VALUE	The property value

19.12 Hardware Views

This section provides a description of each hardware view and its columns.

For examples of how to use views, see [Section 19.30, "Examples"](#).

19.12.1 MGMT\$HW_CPU_DETAILS

The MGMT\$HW_CPU_DETAILS view returns a summary of hardware CPU details.

Table 19–75 MGMT\$HW_CPU_DETAILS

Column	Description
TARGET_TYPE	Type of target for this metric
TARGET_NAME	The name of the target
VENDOR_NAME	The name of the hardware vendor
FREQUENCY_IN_MHZ	The frequency measured in MHz
ECACHE_IN_MB	The size of the ecache measured in MB
IMPL	The details of the implementation
REVISION	The revision details
MASK	The mask details
INSTANCE_COUNT	This is a count of the CPU devices
NUM_CORES	The number of cores per physical CPU
IS_HYPERTHREADING_ENABLED	Defines whether hyperthreading is enabled for this physical CPU (set to 0 or 1)
SIBLINGS	Total number of logical processors for this physical CPU
SNAPSHOT_GUID	The globally unique identifier of the snapshot
TARGET_GUID	The globally unique identifier of the target

Table 19–75 (Cont.) MGMT\$HW_CPU_DETAILS

Column	Description
LAST_COLLECTION_TIMESTAMP	The date-time of the last collection

19.12.2 MGMT\$HW_NIC

The MGMT\$HW_NIC view returns a summary of hardware network interface card (NIC) information.

Table 19–76 MGMT\$HW_NIC

Column	Description
TARGET_TYPE	Type of target for this metric
TARGET_NAME	The name of the target
HOST_NAME	The name of the host
NAME	The NIC name
INET_ADDRESS	The NIC address
MAX_TRANSFER_UNIT	The NIC maximum transfer unit
BROADCAST_ADDRESS	The NIC broadcast address
FLAGS	The NIC flags
MASK	The NIC masks
MAC_ADDRESS	The NIC MAC address
MAC_ADDRESS_STD	The STD NIC MAC address
DHCP_ENABLED	Defines whether DHCP is enabled (set to Y or N)
HOST_ALIASES	The NIC host aliases
INET6_ADDRESSES	The Ipv6 addresses of the host
SNAPSHOT_GUID	The globally unique identifier of the snapshot
TARGET_GUID	The globally unique identifier of the target
LAST_COLLECTION_TIMESTAMP	The date-time of the last collection
IS_PHYSICAL	Defines whether the NIC is physical or not (set to 0 or 1)

19.12.3 MGMT\$HW_NIC_BONDS

The MGMT\$HW_NIC_BONDS view returns a summary of hardware network interface card (NIC) bonds information.

Table 19–77 MGMT\$HW_NIC_BONDS

Column	Description
TARGET_TYPE	Type of target for this metric
TARGET_NAME	The name of the target
BOND_NAME	The name of the bond
PRIMARY_SLAVE	The primary slave of the bond
SECONDARY_SLAVES	The secondary slaves of the bond

Table 19–77 (Cont.) MGMT\$HW_NIC_BONDS

Column	Description
BOND_MODE	The mode of the bond, for example, Balanced
OPTIONS	The options used when the bond is created
SNAPSHOT_GUID	The globally unique identifier of the snapshot
TARGET_GUID	The globally unique identifier of the target
LAST_COLLECTION_TIMESTAMP	The date-time of the last collection

19.12.4 MGMT\$HW_IO_DEVICES

The MGMT\$HW_IO_DEVICES view returns a summary of IO device details.

Table 19–78 MGMT\$HW_IO_DEVICES

Column	Description
TARGET_TYPE	Type of target for this metric
TARGET_NAME	The name of the target
VENDOR_NAME	The vendor name
NAME	The name of the IO device
FREQ_IN_MHZ	The frequency in MHz
BUS	The bus type
REVISION	The revision of the IO device
INSTANCE_COUNT	This is a count of the IO devices that have the same vendor name, name, and so on
SNAPSHOT_GUID	The globally unique identifier of the snapshot
TARGET_GUID	The globally unique identifier of the target
LAST_COLLECTION_TIMESTAMP	The date-time of the last collection

19.13 Inventory Views

This section provides a description of each inventory view and its columns.

For examples of how to use these views, see [Section 19.30, "Examples"](#).

19.13.1 MGMT\$TARGET

The MGMT\$TARGET view displays information about the managed targets that are known to the Management Repository. These targets may or may not be actively monitored.

Table 19–79 MGMT\$TARGET

Column	Description
TARGET_NAME	Name of the target where the metric was collected. The target name uniquely identifies a managed target within the Management Repository. The target name typically contains the name of the managed entity that was provided by the system or database administrator.

Table 19–79 (Cont.) MGMT\$TARGET

Column	Description
TARGET_TYPE	The type of the target. Types of targets may include databases, hosts, web servers, applications, or Application Servers. The definer of the collection definition at the Management Agent defines the target type. The target type defines the set of metrics that are collected for a managed target within the Management Repository.
TARGET_GUID	The unique global identifier for the target.
TYPE_VERSION	The target type meta version of the metadata set. Metadata versions may be updated when applying patches or upon new releases of Enterprise Manager Grid Control.
TYPE_QUALIFIER1-5	Up to five qualifiers can be used to distinguish different metric definitions based on different system configurations. Example qualifier entries may include operating system version, database version, or RAC configuration.
EMD_URL	The URL address of the Management Agent that is managing the target
TIMEZONE_REGION	The time zone region in which the target operates
DISPLAY_NAME	User-friendly name for the target
HOST_NAME	Name of the host where the target is running. For composite targets or targets that span a host, this column will be NULL.
LAST_METRIC_LOAD_TIME	Timestamp when information for this target was last loaded into the Management Repository. If metrics have not been loaded into the Management Repository for the target, this column will be NULL.
TYPE_DISPLAY_NAME	User-friendly name of the target type

Usage Notes

- Display a list of the targets known to the Management Repository.
- Display administration and monitoring information in the context of a managed target.
- Order the targets by last load time for customers to get a sense on how recent the information is for a target in the Management Repository. To access this information in an ordered way, customers should use the appropriate ORDER BY clause with the view.
- Access to this view will use an index if the query references the target name and target type.
- There is an implicit assumption that customers will not use this view to identify the targets that are owned by a Management Agent or the targets that reside on a specific host.

19.13.2 MGMT\$TARGET_TYPE

The MGMT\$TARGET_TYPE view displays metric descriptions for a given target name and target type. This information is available for the metrics for the managed targets that have been loaded into the Management Repository. Metrics are specific to the target type.

Table 19–80 *MGMT\$TARGET_TYPE*

Column	Description
TARGET_NAME	Name of the target where the metric was collected. The target name uniquely identifies a managed target within the Management Repository. The target name typically contains the name of the managed entity that was provided by the system or database administrator.
TARGET_TYPE	The target type defines the set of metrics that are applicable for the target
TARGET_GUID	The unique global identifier for the target
TYPE_VERSION	The target type meta version of the metadata set. Metadata versions may be updated when applying patches or upon new releases of Enterprise Manager Grid Control.
TYPE_QUALIFIER1-5	Up to five qualifiers can be used to distinguish different metric definitions based on different system configurations. Example qualifier entries may include operating system version, database version, or RAC configuration.
METRIC_NAME	The name of the metric that is being defined
METRIC_COLUMN	<p>For table metrics, the metric column contains the name of the column in the table that is being defined. If the metric that is being defined is not a table metric, the value in this column is a single space.</p> <p>For example, if a table describing the MGMT\$TARGET_TYPE view is to be defined as a table metric, Column Name, Data Type, and Description would be metric columns.</p>
KEY_COLUMN	<p>For table metrics, the key column contains the name of the column in the table that represents the primary key. Values in this column must uniquely identify rows in the table. If the metric that is being defined is not a table metric, the value in this column is a single space;</p> <p>For example, the Column Name would be the key column if this table describing the MGMT\$TARGET_TYPE view was being defined as a table metric.</p>
METRIC_TYPE	<p>A DECODE of the internal numeric type of the metric that is being defined. This column will contain one of the following values:</p> <ul style="list-style-type: none"> ■ Number ■ String ■ Table ■ Raw ■ External ■ Repository Metric
METRIC_LABEL	A intuitive display name for the metric that is being defined
COLUMN_LABEL	For table metrics, the column label contains a user understandable display name for the metric column
DESCRIPTION	A description of the metric that is being defined
DESCRIPTION_NLSID	The NLSid of the description of the metric
UNIT	The unit of the metric that is being defined
UNIT_NLSID	The NLSid of the unit of the metric being defined

Table 19–80 (Cont.) MGMT\$TARGET_TYPE

Column	Description
SHORT_NAME	This is a shortened version of the metric display name for the "dense" UI concept
SHORT_NAME_NLSID	The NLSid of the short name of the metric being defined

Usage Notes

- List the set of metrics that have been defined for a target type.
- Display intuitive metric names and associated attributes such as unit in a general way during portal, web application, or custom 4GL report generation.
- Access to this view will use an index if the query references the metric name, metric column. The query should also qualify the target name and target type in order to restrict the amount of information returned.

19.13.3 MGMT\$TARGET_TYPE_DEF

The MGMT\$TARGET_TYPE_DEF view displays definition information for a target type.

Table 19–81 MGMT\$TARGET_TYPE_DEF

Column	Description
TARGET_TYPE	The target type defines the set of metrics that are applicable for the target
TYPE_DISPLAY_NAME	User-friendly name of the target type
TARGET_TYPE_GUID	The unique global identifier (GUID) of the target type
MAX_TYPE_META_VER	The maximum version of the target type stored in the Management Repository

19.13.4 MGMT\$TARGET_ASSOCIATIONS

The MGMT\$TARGET_ASSOCIATIONS view displays the various associations between targets. This view can be used to find all types of associations for a given target.

Table 19–82 MGMT\$TARGET_ASSOCIATIONS

Column	Description
ASSOC_DEF_NAME	Name of the association definition
SOURCE_TARGET_NAME	Target name of the target to which the association is being defined
SOURCE_TARGET_TYPE	The target type of the target for which the association is being defined. "ANY" can be used to specify that any target type be used.
ASSOC_TARGET_NAME	Target Name of the target which is being associated with the source target
ASSOC_TARGET_TYPE	The target type of the associated target. "ANY" can be used to specify that any target type be used

Table 19–82 (Cont.) MGMT\$TARGET_ASSOCIATIONS

Column	Description
SCOPE_TARGET_NAME	The target under whose scope the association is valid This applies to non-global associations only. For example: A database may be part of a composite target only for a particular service.
SCOPE_TARGET_TYPE	The target type for which the association is valid. This applies to non-global associations only.
ASSOCIATION_TYPE	The type of association

Usage Notes

- Can be used to list the associations defined for a specific target.
- Queries using this view will use an index if either (source_target_name, source_target_type) or (assoc_target_name, assoc_target_type) is specified.

19.13.5 MGMT\$TARGET_MEMBERS

The MGMT\$TARGET_MEMBERS view displays the list of direct members for a target.

Table 19–83 MGMT\$TARGET_MEMBERS

Column	Description
AGGREGATE_TARGET_NAME	Target name of the aggregate target
AGGREGATE_TARGET_TYPE	Target type of the aggregate target
AGGREGATE_TARGET_GUID	Target GUID of the aggregate target
MEMBER_TARGET_NAME	Target name of the member target
MEMBER_TARGET_TYPE	Target type of the member target
MEMBER_TARGET_GUID	Target GUID of the member target

Usage Notes

- Find the members for a aggregate target.
- Find the aggregate targets for which a given target is a direct member.
- Queries, which specify values for (AGGREGATE_TARGET_NAME, AGGREGATE_TARGET_TYPE) or (MEMBER_TARGET_NAME, MEMBER_TARGET_TYPE) will use index.
- Joins using AGGREGATE_TARGET_GUID and MEMBER_TARGET_GUID will be efficient.

19.13.6 MGMT\$TARGET_FLAT_MEMBERS

The MGMT\$TARGET_FLAT_MEMBERS view displays the list of all direct and indirect members of the target.

Table 19–84 MGMT\$TARGET_FLAT_MEMBERS

Column	Description
AGGREGATE_TARGET_NAME	The target name of the aggregate target
AGGREGATE_TARGET_TYPE	The target type of the aggregate target
AGGREGATE_TARGET_GUID	Target GUID of the aggregate target
MEMBER_TARGET_NAME	Target Name of the member target
MEMBER_TARGET_TYPE	Target type of the member target
MEMBER_TARGET_GUID	Target GUID of the member target

Usage Notes

- Find the members for an aggregate target.
- Find the aggregate targets for which a given target is a member either directly or indirectly.
- Queries, which specify values for (AGGREGATE_TARGET_NAME, AGGREGATE_TARGET_TYPE) or (MEMBER_TARGET_NAME, MEMBER_TARGET_TYPE), will use index.
- Joins using AGGREGATE_TARGET_GUID and MEMBER_TARGET_GUID will be the most efficient on this view.

19.13.7 MGMT\$TARGET_TYPE_PROPERTIES

The MGMT\$TARGET_TYPE_PROPERTIES view displays the default list of properties that are applicable to the target based on the target type to which the target belongs.

Table 19–85 MGMT\$TARGET_TYPE_PROPERTIES

Column	Description
TARGET_NAME	Name of the target
TARGET_TYPE	Name of the target type
PROPERTY_NAME	Name of the property, such as is_aggregate, is_service, IsBaselineable
PROPERTY_VALUE	Value of the property

Usage Notes

List the properties applicable to the target and the default values.

19.13.8 MGMT\$TARGET_PROPERTIES

The MGMT\$TARGET_PROPERTIES view displays detailed target properties.

Table 19–86 MGMT\$TARGET_PROPERTIES

Column	Description
TARGET_NAME	Name of the target where the metric was collected. The target name uniquely identifies a managed target within the Management Repository. The target name typically contains the name of the managed entity that was provided by the system or database administrator.
TARGET_TYPE	The target type defines the set of metrics that are applicable for the target
TARGET_GUID	The unique global identifier for the target
PROPERTY_NAME	The name of the target property being defined
PROPERTY_VALUE	The value of the target property being defined
PROPERTY_TYPE	The type of the target property being defined. Possible values are: INSTANCE, if the property is applicable to the target instance. DYNAMIC, if the property is calculated dynamically.

19.14 Job Views

A job is a unit of work that you define to automate commonly-run tasks. This section provides a description of each job view and its columns.

19.14.1 MGMT\$CA_TARGETS

The MGMT\$CA_TARGETS view provides basic information about a Corrective Action (CA).

Table 19–87 MGMT\$CA_TARGETS

Column	Description
CA_NAME	Name of the CA
CA_ID	Unique ID of the CA
CA_OWNER	Owner of the CA
CA_DESCRIPTION	Description of the CA
JOB_TYPE	Job type of the CA
TARGET_NAME	Name of the target associated with the CA
TARGET_TYPE	Type of target associated with the CA
TARGET_GUID	Unique ID of the target associated with the CA
IS_BROKEN	Specifies whether the CA is broken Possible values: <ul style="list-style-type: none"> ■ 1: CA is broken ■ 0: CA is not broken

19.14.2 MGMT\$CA_EXECUTIONS

The MGMT\$CA_EXECUTIONS view provides a summary of the Corrective Actions (CA) executions along with the status and targets for each execution.

Table 19–88 *MGMT\$CA_EXECUTIONS*

Column	Description
CA_NAME	Name of the CA
CA_OWNER	Owner of the CA
CA_ID	Unique ID of the CA
JOB_TYPE	Job type of the CA
EXECUTION_ID	Execution ID of the CA
SCHEDULED_TIME	Scheduled time for the CA execution
START_TIME	Start time of the CA execution (in Coordinated Universal Time (UTC))
END_TIME	End time of the CA execution (in UTC)
TRIGGERING_SEVERITY	Severity that triggered the CA
TARGET_NAME	Name of the target on which the CA executed Note: This value can be different from the target with which the CA is associated
TARGET_TYPE	Type of target on which the CA executed
TARGET_GUID	Unique ID of the target on which the CA executed
TIMEZONE_REGION	The time zone region associated with the execution

Table 19–88 (Cont.) MGMT\$CA_EXECUTIONS

Column	Description
STATUS	<p>Represents the status of the execution</p> <p>Possible values:</p> <ul style="list-style-type: none"> ■ Scheduled: The execution is scheduled ■ Running: The execution has steps that have ran already or are running currently ■ Error: The execution encountered internal errors and terminated ■ Failed: Some steps of the execution ran into failures ■ Succeeded: The execution ran as expected ■ Suspended By User: The user suspended the execution ■ Suspended: Agent Unreachable: The execution cannot continue because the Management Agent cannot be contacted ■ Stopped: The execution is stopped explicitly ■ Suspended on Lock: The execution cannot continue because it is waiting for a logical lock to be obtained ■ Suspended on Event: The execution cannot continue because it is waiting for an internal event (such as a Management Agent restart) or a timeout to occur ■ Suspended on Blackout: The execution cannot proceed because the target it is supposed to run against is under blackout ■ Suspend Pending: The user initiated a suspension of the execution but the execution is still running because some steps cannot be suspended ■ Stop Pending: The user initiated a stop of the execution but the execution is waiting for some steps that could not be stopped ■ Inactive: This status is not used in the current release ■ Queued: The execution is submitted against a queue and there are executions that must complete before this execution can complete ■ Waiting: This execution tracks the next schedule compared to the current scheduled, running, or suspended execution ■ Skipped: The execution did not start and its corresponding schedule was skipped. The skip could be due to many reasons, such as overshooting the start grace period, previous schedule not completed when the scheduled time of this execution passed, or the OMS is down. These executions have no corresponding steps ■ Reassigned: The owner of the job has changed and the new owner has not updated the job to claim ownership. ■ Missing Credentials: The execution is blocked waiting for the user to supply target credentials ■ Action Required: The execution is blocked waiting for user action ■ Suspended on Broken Target: The execution cannot proceed because the corresponding target is broken (it has metric collection issues)
STATUS_INTERNAL	Internal status of the execution.

Table 19–88 (Cont.) MGMT\$CA_EXECUTIONS

Column	Description
STATUS_CODE	Status code of the execution. The meaning of this column varies depending on the job type and the steps ran for the execution. This code usually maps to the exit code of the step that rolls up the status.
STATUS_BUCKET	The status bucket to which the execution corresponds. Note: This value can change from one release to another or between patchsets. Oracle recommends that you do not use this value as filtering criteria
STATE_CYCLE	Provides a lifecycle representation of the status of the execution. Possible values: <ul style="list-style-type: none"> ■ SCHEDULED: The execution has no steps that are executed. All steps corresponding to the execution are either scheduled to be picked up shortly or later in the future ■ RUNNING: The execution has at least one step that ran or is running while no steps are blocked waiting for external processing ■ SUSPENDED: The execution is blocked and is waiting for external processing or an external event. For example, the user might have suspended the execution or a step of the execution might be waiting for some other system to respond (such as clearing of blackout, or a Management Agent restart). Such an execution cannot proceed until a timeout occurs or the event that the execution is waiting on occurs. ■ FINISHED: The execution has reached a terminal status where no further execution is possible. These executions usually cannot be manipulated in any way (except for failed executions that can be retried, but the retry operation creates a new execution, and does not affect the current execution)

19.14.3 MGMT\$JOBS

The MGMT\$JOBS view displays information about a job including the job's schedule.

Table 19–89 MGMT\$JOBS

Column	Description
JOB_NAME	The unique name for the job
JOB_ID	The unique system identifier for the job
JOB_OWNER	The owner or creator of the job
JOB_DESCRIPTION	Optional text describing the job function
JOB_TYPE	The job type. For example, multi-task, SQL script or OS Command.
TARGET_TYPE	The type of target the job was submitted against. Applies to single-target jobs only
IS_LIBRARY	Indicates whether or not the job is part of the job library
IS_RESTARTABLE	Indicates whether the job can be restarted. "0" indicates the job is not restartable. "1" indicates the job is not restartable. By default, a job is not restartable if the original job owner is deleted and the job is transferred to another administrator.

Table 19–89 (Cont.) MGMT\$JOBS

Column	Description
START_TIME	The scheduled start time. For daily, days of week and days of month schedules, the start_time denotes when the job should start.
END_TIME	For all periodic schedules, the last date (and time) to run the job. For daily, day of week and day of month schedules, only the date portion is used.
TIMEZONE_TYPE	Possible values are: <ul style="list-style-type: none"> ■ 1 - Repository (deprecated) ■ 2 - Agent ■ 3 - Specified Offset/Region (offset from GMT) ■ 4 - Specified Offset/Region (time zone region name)
TIMEZONE_REGION	The specified time zone region
SCHEDULE_TYPE	Possible values are: <ul style="list-style-type: none"> ■ 1 - One Time ■ 2 - Interval ■ 3 - Daily ■ 4 - Weekly ■ 5 - Monthly ■ 6 - Yearly
INTERVAL	If schedule_type is interval (2), this is the interval at which the job repeats, in minutes
EXECUTION_HOURS	Indicates the time of day at which the job will execute. Hours are specified using the 24-hour format (0 to 23)
EXECUTION_MINUTES	Indicates the time of day at which the job will execute. Minutes are specified as a number between 0 and 59
MONTHS	For days-of-year job schedules, this indicates the “month” in the schedule
DAYS	For day-of-week/month or day(s) of the week job schedules, this indicates the “day” of the week/month. Days-of-week specified as numbers 1 (Sunday) to 7 (Saturday). Days-of-month specified as numbers 1 to 31.

19.14.4 MGMT\$JOB_TARGETS

The MGMT\$JOB_TARGETS view displays the target(s) the job was submitted against.

Table 19–90 MGMT\$JOB_TARGETS

Column	Description
JOB_NAME	The unique name for the job
JOB_OWNER	The owner or creator of the job
JOB_ID	Unique id of the submitted job
JOB_TYPE	The job type. For example, multi-task, SQL script or OS Command.
TARGET_NAME	Name of the target the job was submitted against

Table 19–90 (Cont.) MGMT\$JOB_TARGETS

Column	Description
TARGET_TYPE	The type of target the job was submitted against. Applies to single-target jobs only.
TARGET_GUID	The unique global identifier for the target

19.14.5 MGMT\$JOB_EXECUTION_HISTORY

The MGMT\$JOB_EXECUTION_HISTORY view displays a summary of job executions along with their status and targets for each execution.

Table 19–91 MGMT\$JOB_EXECUTION_HISTORY

Column	Description
JOB_NAME	The unique name for the job
JOB_OWNER	The owner of the job
JOB_ID	The unique system identifier for the job
JOB_TYPE	The job type. For example, multi-task, SQL script or OS Command.
EXECUTION_ID	The unique execution identifier
SCHEDULED_TIME	The scheduled time of job execution (using the time zone specified by TIMEZONE_REGION)
START_TIME	The actual time (UTC) the job executed
END_TIME	The actual time (UTC) the job ended execution
TARGET_NAME	Name of the target the job was submitted against
TARGET_TYPE	The type of target the job was submitted against. Applies to single-target jobs only.
TARGET_GUID	The unique global identifier for the target
TIMEZONE_REGION	The time zone region specifying to the job execution

Table 19–91 (Cont.) MGMT\$JOB_EXECUTION_HISTORY

Column	Description
STATUS	<p>The current status of the job execution</p> <p>Valid values:</p> <ul style="list-style-type: none"> ■ Scheduled: The job execution is scheduled ■ Running: The job execution had steps that are run already or are running at the moment ■ Error: The job execution encountered internal errors and stopped unexpectedly ■ Failed: Some of the steps of the job execution failed ■ Succeeded: The job execution completed as expected ■ Suspended By User: The user suspended the job execution ■ Suspended: Agent Unreachable: The job execution cannot continue because it cannot contact the Management Agent ■ Stopped: The job execution is stopped ■ Suspended on Lock: The job execution cannot proceed because it waiting for a logical lock to be obtained ■ Suspended on Event: The job execution cannot proceed because it is waiting for an internal event (such as restarting the Management Agent) or a timeout to occur ■ Suspended on Blackout: The job execution cannot proceed because the target that the job is running against is under blackout ■ Suspend Pending: The user initiated a suspension but the job execution is still running because some steps could not be suspended ■ Stop Pending: The user initiated a stop but the job execution is waiting for some steps that could not be stopped ■ Inactive: This status is not used in the current release ■ Queued: The job execution is submitted against a queue and there are other jobs that must complete before this job execution can proceed ■ Waiting: The job execution tracks the next schedule compared to the current scheduled or running or suspended job execution ■ Skipped: The job execution did not start and its corresponding schedule is skipped ■ Reassigned: The job owner is changed and the new owner has not updated to job to claim ownership ■ Missing Credentials: The job execution is blocked and is waiting for the user to supply target credentials ■ Action Required: The job execution is blocked waiting for user action ■ Suspended on Broken Target: The job execution cannot proceed because the corresponding target is broken
STATUS_INTERNAL	The internal integer status of the job execution
STATUS_CODE	The status code of the job execution. This integer usually maps to the exit code of the step that rolls up the status.

Table 19–91 (Cont.) MGMT\$JOB_EXECUTION_HISTORY

Column	Description
STATUS_BUCKET	The status bucket to which the execution corresponds Note: This value can change between patchsets and Oracle recommends that you do not use this value as filtering criteria
STATE_CYCLE	Provides a lifecycle representation of the status of the execution Valid values: <ul style="list-style-type: none"> ■ SCHEDULED: The execution has no steps that are executed yet ■ RUNNING: The execution has at least one step that ran or is running and no steps are blocked ■ SUSPENDED: The execution has blocked steps that are waiting for external processing or an external event ■ FINISHED: The execution has reached where no further execution is possible
SOURCE_EXECUTION_ID	The execution which was retried and caused this execution (EXECUTION_ID). For executions that are not retried, the SOURCE_EXECUTION_ID is the same as the EXECUTION_ID.
RETRIED	This value is set to 1 if this execution is retried. Otherwise, it is set to 0

19.14.6 MGMT\$JOB_STEP_HISTORY

The MGMT\$JOB_STEP_HISTORY view displays step-level details of job executions.

Table 19–92 MGMT\$JOB_STEP_HISTORY

Column	Description
JOB_NAME	The unique name for the job
JOB_OWNER	The owner or creator of the job
JOB_ID	The unique system identifier for the job
EXECUTION_ID	The unique execution identifier
STEP_NAME	The name of the job step
START_TIME	The start time of the job step
END_TIME	The end time of the job step
STATUS	The current status of the job execution
TARGET_NAME	Name of the target the job was submitted against
TARGET_TYPE	The type of target the job was submitted against. Applies to single-target jobs only.
TARGET_GUID	The unique global identifier for the target
OUTPUT	Generated job output
STATUS_INTERNAL	The internal integer status of the job step
STATUS_CODE	The status code of the job execution. This integer usually maps to the exit code of the step that rolls up the status
STEP_ID	The step ID of the job step
STEP_TYPE	The step type of the job step

19.14.7 MGMT\$JOB_ANNOTATIONS

The MGMT\$JOB_ANNOTATIONS view displays a summary of annotations for changes in job status.

Table 19–93 MGMT\$JOB_ANNOTATIONS

Column	Description
JOB_NAME	The unique name for the job
JOB_OWNER	The owner or creator of the job
JOB_STATUS	The job status. Possible values are as follows: <ul style="list-style-type: none"> ■ 1: Scheduled ■ 2: Executing ■ 3: Aborted ■ 4: Failed ■ 5: Completed ■ 6: Suspended ■ 7: Agent Down ■ 8: Stopped ■ 9: Suspended/Lock ■ 11: Suspended/Blackout ■ 13: Suspend Pending ■ 15: Queued ■ 16: Failed ■ 17: Waiting ■ 18: Skipped
OCCURRENCE_TIMESTAMP	The time at which the state change occurred
ANNOTATION_MESSAGE	Annotation text
ANNOTATION_TIMESTAMP	The time the annotation was created
ANNOTATED_BY	Enterprise Manager administrator who authored the annotation

19.14.8 MGMT\$JOB_NOTIFICATION_LOG

The MGMT\$JOB_NOTIFICATION_LOG view displays details of notification deliveries for changes in job status.

Table 19–94 MGMT\$JOB_NOTIFICATION_LOG

Column	Description
JOB_NAME	The unique name for the job
JOB_OWNER	The owner or creator of the job

Table 19–94 (Cont.) MGMT\$JOB_NOTIFICATION_LOG

Column	Description
JOB_STATUS	<p>The job status.</p> <p>Possible values:</p> <ul style="list-style-type: none"> ■ 1: Scheduled ■ 2: Executing ■ 3: Aborted ■ 4: Failed ■ 5: Completed ■ 6: Suspended ■ 7: Agent Down ■ 8: Stopped ■ 9: Suspended/Lock ■ 11: Suspended/Blackout ■ 13: Suspend Pending ■ 15: Queued ■ 16: Failed ■ 17: Waiting ■ 18: Skipped
OCCURRENCE_TIMESTAMP	The time at which the state change occurred
DELIVERY_MESSAGE	The message indicating the success or failure of the notification delivery
DELIVERY_TIMESTAMP	The time at which the log message was created

19.15 Linux Patching Views

This section provides a description of each Linux patching view and its columns.

19.15.1 MGMT\$HOSTPATCH_HOSTS

The MGMT\$HOSTPATCH_HOSTS view displays information required to generate compliance reports.

Table 19–95 MGMT\$HOSTPATCH_HOSTS

Column	Description
HOST_NAME	Host name
GROUP_NAME	The group the host belongs to
OUT_OF_DATE_PACKAGES	Number of Packages which have a newer version available
ROGUE_PACKAGES	The packages that are not supposed to be installed on the host

19.15.2 MGMT\$HOSTPATCH_GROUPS

The MGMT\$HOSTPATCH_GROUPS view displays additional information about a group, the maturity level which is set by the administrator and the packages which need the host to be rebooted on application.

Table 19–96 MGMT\$HOSTPATCH_GROUPS

Column	Description
GROUP_NAME	The (unique) name of the group
MATURITY_LEVEL	The maturity level of the group. This is set by the administrator.
NEED_REBOOT_PKGS	Comma separated list of packages which need the system to be rebooted on application

19.15.3 MGMT\$HOSTPATCH_GRP_COMPL_HIST

The MGMT\$HOSTPATCH_GRP_COMPL_HIST view displays information required to generate compliance history reports.

Table 19–97 MGMT\$HOSTPATCH_GRP_COMPL_HIST

Column	Description
GROUP_NAME	Name of the group
TOTAL_HOSTS	Number of hosts in the group
COMPLIANT_HOSTS	Number of compliant hosts in the group
LAST_CHECKED_ON	Date on which this record was collected

19.15.4 MGMT\$HOSTPATCH_HOST_COMPL

The MGMT\$HOSTPATCH_HOST_COMPL view displays information required to generate advisory reports.

Table 19–98 MGMT\$HOSTPATCH_HOST_COMPL

Column	Description
HOST_NAME	Host name
PKG_NAME	Package name
VERSION	Version of the package
IS_OUT_OF_DATE	If out of date
IS_ROGUE	If it is rogue

19.16 Management Template Views

This section provides a description of each Management Template view and its columns.

19.16.1 MGMT\$TEMPLATES

The MGMT\$TEMPLATES views displays details of all the management templates stored in the Management Repository.

Table 19–99 MGMT\$TEMPLATES

Column	Description
TARGET_TYPE	The target type defines the set of metrics and policies that are applicable for the target
TEMPLATE_NAME	The name of the template

Table 19–99 (Cont.) MGMT\$TEMPLATES

Column	Description
TEMPLATE_GUID	The unique global identifier for the template
DESCRIPTION	The description of the template
OWNER	Enterprise Manager administrator who owns the template
IS_PUBLIC	The flag to specify whether the template is accessible to all EM administrators
CREATED_DATE	The date or time when the template is created in the repository
LAST_UPDATED_DATE	The date or time when the template was last modified in the repository
LAST_UPDATED_BY	The Enterprise Manager administrator who last updated the template

19.16.2 MGMT\$TEMPLATE_POLICY_SETTINGS

The MGMT\$TEMPLATE_POLICY_SETTINGS view displays policy settings for management templates.

Table 19–100 MGMT\$TEMPLATE_POLICY_SETTINGS

Column	Description
TARGET_TYPE	The target type defines the set of metrics and policies that are applicable for the target
TEMPLATE_NAME	The name of the template
TEMPLATE_GUID	The unique global identifier for the template
POLICY_NAME	The name of the policy that is associated with the template
POLICY_GUID	The unique global identifier for the policy
CATEGORY	The name of the category the policy Refer to MGMT\$METRIC_CATEGORIES for the list of all categories
KEY_VALUE	The key value for which the alert has been recorded. For composite keys, this is the first part of the key
KEY_VALUE2	For composite keys, this is the second part of the key for which the alert has been recorded
KEY_VALUE3	For composite keys, this is the third part of the key for which the alert has been recorded
KEY_VALUE4	For composite keys, this is the fourth part of the key for which the alert has been recorded
KEY_VALUE5	For composite keys, this is the fifth part of the key for which the alert has been recorded
KEY_OPERATOR	Specifies whether the key_value columns have any SQL wildcards. For single key column metrics, the value is 1 if the key_value has wildcard characters, 0 otherwise. For metrics with multiple keys, a list of operators for all key columns will be stored here. For example, a metric with 3 keys (k1, k2, k3) where K1 and K2 use wildcards and K3 uses exact match, then 011 is stored in this column
PARAMETER_NAME	The name of the parameter

Table 19–100 (Cont.) MGMT\$TEMPLATE_POLICY_SETTINGS

Column	Description
PREVENT_OVERRIDE	The is a flag that specifies that the metric rows with this key_value has a template override flag. Once the template override flag is ON, any template application will not update the threshold/parameter values.
POLICY_THRESHOLD	The threshold value configured for the policy parameter
ACTION_TYPE	The corrective action type configured. Possible values are: No-Action: when no action is configured Corrective-Action: when a repository side corrective action is configured Agent-Fixit Job: when an Agent side fix-it job is configured.
ACTION_JOB_TYPE	Specifies the job type of the corrective action when ACTION_TYPE is "Corrective-Action"
ACTION_JOB_NAME	Specifies the job name of the corrective action when ACTION_TYPE is "Corrective-Action"
ACTION_JOB_OWNER	Specifies the job owner of the corrective action when ACTION_TYPE is "Corrective-Action"

19.16.3 MGMT\$TEMPLATE_METRICCOLLECTION

The MGMT\$TEMPLATE_METRICCOLLECTIONS view displays information on the metric collections defined for a template.

Table 19–101 MGMT\$TEMPLATE_METRICCOLLECTION

Column	Description
TEMPLATE_NAME	The name of the template
TARGET_TYPE	The target type defines the set of metrics and policies that are applicable for the target
TEMPLATE_GUID	The unique global identifier for the template
METRIC_NAME	The name of the metric for which the template collection is configured
METRIC_COLUMN	The name of the metric column for which the template collection is configured
METRIC_GUID	The unique global identifier for the metric column
COLLECTION_NAME	The name of the collection
IS_REPOSITORY	Indicates whether this is a repository-side collection. A repository-side collection has a PL/SQL evaluation procedure that is responsible for calculating the metric values.

Table 19–101 (Cont.) MGMT\$TEMPLATE_METRICCOLLECTION

Column	Description
FREQUENCY_CODE	The metric collection frequency type. Possible values are: <ul style="list-style-type: none"> 1: One Time 2: Interval 3: Daily 4: Weekly 5: Monthly 6: Yearly 7: On Demand
COLLECTION_FREQUENCY	The frequency of the metric collection. Value displayed is dependent on the frequency code: <ul style="list-style-type: none"> For One Time, the start date-time is stored in DD-MON-YY HH24:MI format. For Interval type, the frequency in minutes is stored. For Daily/Weekly/Monthly/Yearly types, the hour and minute of collection is stored in HH24:MI format. For On-Demand type, On-Demand is stored.
UPLOAD_POLICY	The frequency with which the metric data is uploaded/stored.

19.16.4 MGMT\$TEMPLATE_METRIC_SETTINGS

The MGMT\$TEMPLATE_METRIC_SETTINGS view displays management template settings.

Table 19–102 MGMT\$TEMPLATE_METRIC_SETTINGS

Column	Description
TEMPLATE_NAME	The name of the template
TARGET_TYPE	The target type defines the set of metrics and policies that are applicable for the target
TEMPLATE_GUID	The unique global identifier for the template
METRIC_NAME	The name of the metric for which the template collection is configured
METRIC_COLUMN	The name of the metric column for which the template collection is configured
METRIC_GUID	The unique global identifier for the metric column
COLLECTION_NAME	The name of the collection
CATEGORY	The name of the category the policy. Refer to MGMT\$METRIC_CATEGORIES for the list of all categories.
KEY_VALUE	The key value for which the alert has been recorded. For composite keys, this is the first part of the key.
KEY_VALUE2	For composite keys, this is the second part of the key for which the alert has been recorded
KEY_VALUE3	For composite keys, this is the third part of the key for which the alert has been recorded

Table 19–102 (Cont.) MGMT\$TEMPLATE_METRIC_SETTINGS

Column	Description
KEY_VALUE4	For composite keys, this is the fourth part of the key for which the alert has been recorded
KEY_VALUE5	For composite keys, this is the fifth part of the key for which the alert has been recorded
KEY_OPERATOR	<p>Specifies whether the key_value columns have any SQL wildcards.</p> <p>For single key column metrics, the value is 1 if the key_value has wildcard characters, 0 otherwise.</p> <p>For metrics with multiple keys, a list of operators for all key columns will be stored here. For example, a metric with 3 keys (k1, k2, k3) where K1 and K2 use wildcards and K3 uses exact match, then 011 is stored in this column</p>
PREVENT_OVERRIDE	The is a flag that specifies that the metric rows with this key_value has a template override flag. Once the template override flag is ON, any template application will not update the threshold/parameter values.
WARNING_OPERATOR	<p>Defines the warning threshold condition to be applied</p> <ul style="list-style-type: none"> ■ 0 - GT ■ 1 - EQ ■ 2 - LT ■ 3 - LE ■ 4 - GE ■ 5 - CONTAINS ■ 6 - NE ■ 7 - MATCH : regular expression
WARNING_THRESHOLD	The value of the warning threshold
CRITICAL_OPERATOR	<p>Defines the critical threshold condition to be applied</p> <ul style="list-style-type: none"> ■ 0 - GT ■ 1 - EQ ■ 2 - LT ■ 3 - LE ■ 4 - GE ■ 5 - CONTAINS ■ 6 - NE ■ 7 - MATCH : regular expression
CRITICAL_THRESHOLD	The value of the critical threshold
OCCURRENCE_COUNT	The number of times the test has to trigger to raise a violation
WARNING_ACTION_TYPE	<p>The warning corrective action type configured. Possible values are:</p> <ul style="list-style-type: none"> ■ No-Action: when no action is configured ■ Corrective-Action: when a repository side corrective action is configured ■ Agent-Fixit Job: when a Management Agent side fix-it job is configured

Table 19–102 (Cont.) MGMT\$TEMPLATE_METRIC_SETTINGS

Column	Description
WARNING_ACTION_JOB_TYPE	Specifies the job type of the warning corrective action when WARNING_ACTION_TYPE is “Corrective-Action”
WARNING_ACTION_JOB_OWNER	Specifies the job owner of the warning corrective action when WARNING_ACTION_TYPE is “Corrective-Action”
WARNING_ACTION_JOB_NAME	Specifies the job name of the warning corrective action when WARNING_ACTION_TYPE is “Corrective-Action”
CRITICAL_ACTION_TYPE	The critical corrective action type configured. Possible values are: <ul style="list-style-type: none"> ■ No-Action: when no action is configured ■ Corrective-Action: when a repository side corrective action is configured ■ Agent-Fixit Job: when a Management Agent side fix-it job is configured
CRITICAL_ACTION_JOB_TYPE	Specifies the job type of the critical corrective action when CRITICAL_ACTION_TYPE is “Corrective-Action”
CRITICAL_ACTION_JOB_OWNER	Specifies the job owner of the critical corrective action when CRITICAL_ACTION_TYPE is “Corrective-Action”
CRITICAL_ACTION_JOB_NAME	Specifies the job name of the critical corrective action when CRITICAL_ACTION_TYPE is “Corrective-Action”

19.17 Metric Views

This section provides a description of each metric view and its columns.

19.17.1 MGMT\$METRIC_CATEGORIES

The MGMT\$METRIC_CATEGORIES view displays the list of classes and categories to which the metric belongs. It can be used to classify the metric based on the class (such as service or functional) and category within the class (such as security or configuration under functional class or usage or performance under service class).

Table 19–103 MGMT\$METRIC_CATEGORIES

Column	Description
TARGET_TYPE	Defines the target type of the metric being categorized
TYPE_VERSION	Defines the version of the target type
METRIC_NAME	Defines the name of the metric
METRIC_COLUMN	For table metrics, the metric column contains the name of the column in the table that is being defined. If the metric that is being defined is not a table metric, then the value in this column is a single space.
METRIC_GUID	A unique ID for the metric
METRIC_CLASS_NAME	Name of the metric class to which this metric belongs
METRIC_CATEGORY_NAME	Name of the category of the metric in the class
METRIC_CATEGORY_NLSID	The NLS ID of the category which is used by Enterprise Manager to translate the name to different languages

19.17.2 MGMT\$METRIC_COLLECTION

The MGMT\$METRIC_COLLECTION view provides the metric thresholds details, frequency, upload policy, and thresholds per target.

Table 19–104 MGMT\$METRIC_COLLECTION

Column	Description
TARGET_NAME	Target where the metrics will be collected
TARGET_TYPE	Defines the set of metrics that are applicable for the target
TARGET_GUID	The unique id of the target
METRIC_NAME	Name of the metric
METRIC_COLUMN	For table metrics, the metric column contains the name of the column in the table that is being defined. If the metric that is being defined is not a table metric, then the value in this column is a single space.
WARNING_OPERATOR	<p>The operator for the warning threshold. This is used in the reporting environment to create a line on the graph representing the warning threshold appropriately.</p> <p>Possible values:</p> <ul style="list-style-type: none"> ■ 0 - GT ■ 1 - EQ ■ 2 - LT ■ 3 - LE ■ 4 - CONTAINS ■ 5 - NE ■ 6 - MATCH : regular expression
CRITICAL_OPERATOR	<p>The operator for the critical threshold. This is used in the reporting environment to create a line on the graph representing the critical threshold appropriately.</p> <p>Possible values:</p> <ul style="list-style-type: none"> ■ 0 - GT ■ 1 - EQ ■ 2 - LT ■ 3 - LE ■ 4 - CONTAINS ■ 5 - NE ■ 6 - MATCH : regular expression
WARNING_THRESHOLD	Value for the warning severity
CRITICAL_THRESHOLD	Value for the critical severity
OCCURENCE_COUNT	The number of occurrences of a warning, critical, or clear severity before a severity record is generated
WARNING_COUNT	The number of consecutive times a metric value has exceeded the warning threshold
CRITICAL_COUNT	The number of consecutive times a metric value has exceeded the critical threshold

19.17.3 MGMT\$METRIC_ERROR_CURRENT

The MGMT\$METRIC_ERROR_CURRENT view associates current metric errors pertaining to a metric.

Table 19–105 MGMT\$METRIC_ERROR_CURRENT

Column	Description
TARGET_NAME	Name of the target for which the metric collection error occurred
TARGET_TYPE	Target type defines the set of metrics that are applicable for the target
TARGET_GUID	The unique ID of the target for which the metric collection error occurred
METRIC_NAME	The underlying metric for which the metric collection error occurred
METRIC_GUID	A unique ID for the metric. It can be used to associate metric information with metric data information during reporting
METRIC_LABEL	User display name of the metric for which the error occurred
COLL_NAME	Name of the collection collecting the metric for which the error occurred
COLLECTION_TIMESTAMP	Time when the collection error occurred
ERROR_TYPE	Indicates the type of error that happened during the collection of the metric Possible values: <ul style="list-style-type: none"> ■ ERROR ■ WARNING
ERROR_MESSAGE	Text of the error message

19.17.4 MGMT\$METRIC_ERROR_HISTORY

The MGMT\$METRIC_ERROR_HISTORY view displays the history of metric collection errors.

Table 19–106 METRIC_ERROR_HISTORY

Column	Description
TARGET_NAME	The name of the target for which the metric collection error occurred
TARGET_TYPE	Target type defines the set of metrics that are applicable for the target
TARGET_GUID	Target GUID of the target for which the metric collection error occurred
METRIC_NAME	Name of the metric for which the error occurred
METRIC_GUID	A unique ID for the metric. It can be used to associate metric information with metric data information during reporting
METRIC_LABEL	User display name of the metric for which the error occurred
COLL_NAME	Name of the collection collecting the metric for which the error occurred
COLLECTION_TIMESTAMP	The time when the collection error occurred

Table 19–106 (Cont.) METRIC_ERROR_HISTORY

Column	Description
ERROR_TYPE	Indicates the type of error that happened during the collection of the metric Possible values are: <ul style="list-style-type: none"> ■ ERROR ■ WARNING
ERROR_MESSAGE	Text of the error message

19.18 Monitoring Views

This section provides a description of each monitoring view and its columns.

For examples of how to use these views, see [Section 19.30, "Examples"](#).

19.18.1 MGMT\$ALERT_CURRENT

MGMT\$ALERT_CURRENT displays current information for any alerts that are logged in the Management Repository that are in a non-clear state. Only the most recent open alert in a non-clear status for a given metric is displayed through this view.

Table 19–107 MGMT\$ALERT_CURRENT

Column	Description
TARGET_NAME	Name of the target where the metric was collected. The target name uniquely identifies a managed target within the Management Repository. The target name typically contains the name of the managed entity that was provided by the system or database administrator.
TARGET_TYPE	The target type defines the set of metrics that are applicable for the target
TARGET_GUID	The unique global identifier for the target
VIOLATION_GUID	Unique identifier for the alert
METRIC_NAME	Name of the metric being defined
METRIC_COLUMN	For table metrics, the metric column contains the name of the column in the table that is being defined. If the metric that is being defined is not a table metric, the value in this column is a single space. For example, if this table describing the MGMT\$TARGET_TYPE view was being defined as a table metric, Column Name, Data Type, and Description would be metric columns.
METRIC_LABEL	An intuitive display name for the metric that is being defined
COLUMN_LABEL	For table metrics, the column label contains an intuitive display name for the metric column
KEY_VALUE	The key value for which the alert has been recorded. For composite keys, this is the first part of the key
KEY_VALUE2	For composite keys, this is the second part of the key for which the alert has been recorded
KEY_VALUE3	For composite keys, this is the third part of the key for which the alert has been recorded

Table 19–107 (Cont.) MGMT\$ALERT_CURRENT

Column	Description
KEY_VALUE4	For composite keys, this is the fourth part of the key for which the alert has been recorded
KEY_VALUE5	For composite keys, this is the fifth part of the key for which the alert has been recorded
COLLECTION_TIMESTAMP	The date-time when the alert condition was detected by the Management Agent
ALERT_STATE	<p>A user readable description of the internal alert code that is sent from the Management Agent to identify the state of the alert condition. A alert record is transferred to the repository from the Management Agent each time the metric threshold is crossed in either direction, or if the Management Agent is restarted. The value of this column will contain one of the following strings:</p> <ul style="list-style-type: none"> ■ Warning ■ Critical <p>If the metric's alert condition goes into a clear state, it will no longer be visible from this view.</p>
VIOLATION_TYPE	<p>A user readable description of the type of violation. Possible values are:</p> <ul style="list-style-type: none"> ■ Threshold Violation, when the alert is triggered based on a metric threshold ■ Availability, when the alert is triggered for an availability metric ■ Policy Violation, when the alert is triggered based on a policy violation
MESSAGE	An optional message that is generated when the alert is created that provides additional information about the alert condition
MESSAGE_NLSID	The NLSID of the alert message
MESSAGE_PARAMS	Contains the URL encoded parameters separated by "&" to be used to format the alert message
ACTION_MESSAGE	Suggested action message in English for this alert
ACTION_MESSAGE_NLSID	The NLS ID of the action message
ACTION_MESSAGE_PARAMS	Contains the URL encoded parameters for translating action message
TYPE_DISPLAY_NAME	The display name of the target type

Usage Notes

- List the current alerts that are in a non-clear state for a metric, set of metrics, or for a managed target. If the user is only interested in non-clear alerts, counts or selects, using this view provide better performance than using the MGMT\$ALERT_DETAILS view.
- Access to this view will use an index if the query references the member target name, target type, metric name, and metric column or a subset of these columns if they are included as listed above from left to right.

19.18.2 MGMT\$TARGET_METRIC_COLLECTIONS

The MGMT\$TARGET_METRIC_COLLECTIONS view displays information about the metric collections.

Table 19–108 MGMT\$TARGET_METRIC_COLLECTIONS

Column	Description
TARGET_NAME	Name of the target where the metric was collected. The target name uniquely identifies a managed target within the Management Repository. The target name typically contains the name of the managed entity that was provided by the system or database administrator.
TARGET_TYPE	The target type defines the set of metrics that are applicable for the target
TARGET_GUID	Unique global identifier (GUID) for the target
METRIC_NAME	Name of the metric being defined
METRIC_COLUMN	For table metrics, the metric column contains the name of the column in the table that is being defined. If the metric that is being defined is not a table metric, the value in this column is a single space. For example, if the table describing the MGMT\$TARGET_TYPE view is defined as a table metric, then Column Name, Data Type, and Description would be metric columns.
METRIC_GUID	Unique global identifier (GUID) for the metric. This ID can be used to associate metric information with metric data information during reporting.
COLLECTION_NAME	The name of the collection
IS_ENABLED	Indicates whether the collection is currently enabled <ul style="list-style-type: none"> 0=not enabled 1=enabled
IS_REPOSITORY	Indicates whether this is a repository-side collection. A repository-side collection has a PL/SQL evaluation procedure that is responsible for calculating the metric values.
FREQUENCY_CODE	The metric collection frequency type. Possible values are: <ul style="list-style-type: none"> 1: One Time 2: Interval 3: Daily 4: Weekly 5: Monthly 6: Yearly 7: On Demand
COLLECTION_FREQUENCY	Frequency of the metric collection. Value displayed is dependent on the frequency code: <ul style="list-style-type: none"> For One Time, the start date-time is stored in DD-MON-YY HH24:MI format. For Interval type, the frequency in minutes is stored. For Daily/Weekly/Monthly/Yearly types, the hour and minute of collection is stored in HH24:MI format. For On-Demand type, On-Demand is stored.

Table 19–108 (Cont.) MGMT\$TARGET_METRIC_COLLECTIONS

Column	Description
UPLOAD_POLICY	The frequency with which the metric data is uploaded or stored

Usage Notes

List the metric collections for a given target.

19.18.3 MGMT\$TARGET_METRIC_SETTINGS

The MGMT\$TARGET_METRIC_SETTINGS view displays information about the current metric setting stored for all targets in the Management Repository. This view provides information for both Management Agent-side and Management Repository-side metrics.

Table 19–109 MGMT\$TARGET_METRIC_SETTINGS

Column	Description
TARGET_NAME	Name of the target where the metric was collected. The target name uniquely identifies a managed target within the Management Repository. The target name typically contains the name of the managed entity that was provided by the system or database administrator.
TARGET_TYPE	The target type defines the set of metrics that are applicable for the target
TARGET_GUID	The unique global identifier (GUID) for the target
METRIC_NAME	Name of the metric being defined
METRIC_COLUMN	For table metrics, the metric column contains the name of the column in the table that is being defined. If the metric that is being defined is not a table metric, the value in this column is a single space. For example, if the table describing the MGMT\$TARGET_TYPE view is defined as a table metric, then Column Name, Data Type, and Description would be metric columns.
METRIC_GUID	The unique global identifier for the metric. This ID can be used to associate metric information with metric data information during reporting.
COLLECTION_NAME	The name of the collection
CATEGORY	The name of the category the metric Refer to MGMT\$METRIC_CATEGORIES for the list of all metric categories.
KEY_VALUE	The key value of the metric setting. For composite keys, this is the first part of the key. If the thresholds are not for a table metric, or the thresholds apply for all rows in the metric column, then the value in this column will contain a single space.
KEY_VALUE2	For composite keys, this is the second part of the key
KEY_VALUE3	For composite keys, this is the third part of the key
KEY_VALUE4	For composite keys, this is the fourth part of the key
KEY_VALUE5	For composite keys, this is the fifth part of the key

Table 19–109 (Cont.) MGMT\$TARGET_METRIC_SETTINGS

Column	Description
KEY_OPERATOR	<p>Specifies whether the key_value columns have any SQL wildcards.</p> <p>For single key column metrics, the value is 1 if the key_value has wildcard characters, 0 otherwise.</p> <p>For metrics with multiple keys, a list of operators for all key columns will be stored here. For example, a metric with 3 keys (k1, k2, k3) where K1 and K2 use wildcards and K3 uses exact match, then 011 is stored in this column.</p>
HAS_ACTIVE_BASELINE	The is a flag that specifies that the metric rows with this key_value has an active baseline and any user updates to thresholds or parameter values should be ignored.
PREVENT_OVERRIDE	The is a flag that specifies that the metric rows with this key_value has a template override flag. Once the template override flag is ON, any template application will not update the threshold or parameter values.
WARNING_OPERATOR	<p>Defines the warning threshold condition to be applied</p> <ul style="list-style-type: none"> ■ 0 - GT ■ 1 - EQ ■ 2 - LT ■ 3 - LE ■ 4 - GE ■ 5 - CONTAINS ■ 6 - NE ■ 7 - MATCH : regular expression
WARNING_THRESHOLD	The warning threshold value
CRITICAL_OPERATOR	<p>Defines the critical threshold condition to be applied</p> <ul style="list-style-type: none"> ■ 0 - GT ■ 1 - EQ ■ 2 - LT ■ 3 - LE ■ 4 - GE ■ 5 - CONTAINS ■ 6 - NE ■ 7 - MATCH : regular expression
CRITICAL_THRESHOLD	The critical threshold value
OCCURRENCE_COUNT	The number of times the test has to trigger to raise a violation
WARNING_ACTION_TYPE	Specifies the job type of the warning corrective action when WARNING_ACTION_TYPE is "Corrective-Action"
WARNING_ACTION_JOB_OWNER	Specifies the job owner of the warning corrective action when WARNING_ACTION_TYPE is "Corrective-Action"
WARNING_ACTION_JOB_NAME	Specifies the job name of the warning corrective action when WARNING_ACTION_TYPE is "Corrective-Action"

Table 19–109 (Cont.) MGMT\$TARGET_METRIC_SETTINGS

Column	Description
CRITICAL_ACTION_TYPE	The critical corrective action type configured. Possible values are: <ul style="list-style-type: none"> ■ No-Action: when no action is configured ■ Corrective-Action: when a repository side corrective action is configured ■ Agent-Fixit Job: when an Agent side fix-it job is configured.
CRITICAL_ACTION_JOB_TYPE	Specifies the job type of the critical corrective action when WARNING_ACTION_TYPE is "Corrective-Action"
CRITICAL_ACTION_JOB_OWNER	Specifies the job owner of the critical corrective action when WARNING_ACTION_TYPE is "Corrective-Action"
CRITICAL_ACTION_JOB_NAME	Specifies the job name of the critical corrective action when WARNING_ACTION_TYPE is "Corrective-Action"

Usage Notes

- List all the metric setting for a given target.
- List the metric settings for a given target and metric.
- List the corrective actions assigned for a given target-metric.

19.18.4 MGMT\$AVAILABILITY_CURRENT

The MGMT\$AVAILABILITY_CURRENT view displays information about the most recent target availability information stored in the Management Repository.

Table 19–110 MGMT\$AVAILABILITY_CURRENT

Column	Description
TARGET_NAME	Name of the target where the metric was collected. The target name uniquely identifies a managed target within the Management Repository. The target name typically contains the name of the managed entity that was provided by the system or database administrator.
TARGET_TYPE	The target type defines the set of metrics that are applicable for the target
START_TIMESTAMP	The time when the target availability status change was first detected
AVAILABILITY_STATUS	Current target availability status. This column contains one of the following values: <ul style="list-style-type: none"> ■ Target Down ■ Target Up ■ Metric Error ■ Agent Down ■ Unreachable ■ Blackout ■ Pending/Unknown

Usage Notes

Get the current availability status of a given target.

19.18.5 MGMT\$AVAILABILITY_HISTORY

The MGMT\$AVAILABILITY_HISTORY view displays detailed historical information about changes in the availability status for a target over time.

Table 19–111 MGMT\$AVAILABILITY_HISTORY

Column	Description
TARGET_NAME	Name of the target where the metric was collected. The target name uniquely identifies a managed target within the Management Repository. The target name typically contains the name of the managed entity that was provided by the system or database administrator.
TARGET_TYPE	The target type defines the set of metrics that are applicable for the target
TARGET_GUID	The unique global identifier for the target
START_TIMESTAMP	The time when the target availability status change was first detected
END_TIMESTAMP	The time when the target availability status change was last detected
AVAILABILITY_STATUS	Target availability status. This column will contain one of the following values: <ul style="list-style-type: none"> Target Down Target Up Metric Error Agent Down Unreachable Blackout Pending/Unknown

Usage Notes

Access to this view will use an index if the query references the member TARGET_NAME, TARGET_TYPE and the START_TIMESTAMP.

19.18.6 MGMT\$ALERT_HISTORY

The MGMT\$ALERT_HISTORY view displays historical information for any alerts that are logged in the Management Repository.

Table 19–112 MGMT\$ALERT_HISTORY

Column	Description
TARGET_NAME	Name of the target where the metric was collected. The target name uniquely identifies a managed target within the Management Repository. The target name typically contains the name of the managed entity that was provided by the system or database administrator.
TARGET_TYPE	The target type defines the set of metrics that are applicable for the target
TARGET_GUID	The unique global identifier for the target
VIOLATION_GUID	Unique identifier for the alert
METRIC_NAME	Name of the metric being defined

Table 19–112 (Cont.) MGMT\$ALERT_HISTORY

Column	Description
METRIC_COLUMN	<p>For table metrics, the metric column contains the name of the column in the table that is being defined. If the metric that is being defined is not a table metric, the value in this column is a single space.</p> <p>For example, if this table describing the MGMT\$TARGET_TYPE view was being defined as a table metric, Column Name, Data Type, and Description would be metric columns.</p>
METRIC_LABEL	An intuitive display name for the metric that is being defined
COLUMN_LABEL	For table metrics, the column label contains an intuitive display name for the metric column
KEY_VALUE	The key value for which the alert has been recorded. For composite keys, this is the first part of the key.
KEY_VALUE2	For composite keys, this is the second part of the key for which the alert has been recorded
KEY_VALUE3	For composite keys, this is the third part of the key for which the alert has been recorded
KEY_VALUE4	For composite keys, this is the fourth part of the key for which the alert has been recorded
KEY_VALUE5	For composite keys, this is the fifth part of the key for which the alert has been recorded
COLLECTION_TIMESTAMP	The date-time when the alert condition was detected by the Management Agent
ALERT_STATE	<p>A user readable description of the internal alert code that is sent from the Management Agent to identify the state of the alert condition. A alert record is transferred to the repository from the Management Agent each time the metric threshold is crossed in either direction, or if the Management Agent is restarted. The value of this column will contain one of the following strings:</p> <ul style="list-style-type: none"> ■ Warning ■ Critical <p>If the metric's alert condition goes into a clear state, it will no longer be visible from this view.</p>
ALERT_DURATION	The time, in hours, from when the alert condition was first detected until it was cleared
MESSAGE	An optional message that is generated when the alert is created that provides additional information about the alert condition
MESSAGE_NLSID	The NLSID of the alert message
MESSAGE_PARAMS	Contains the URL encoded parameters separated by "&" to be used to format the alert message
ACTION_MESSAGE	Suggested action message in English for this alert
ACTION_MESSAGE_NLSID	The NLS ID of the action message
ACTION_MESSAGE_PARAMS	Contains the URL encoded parameters for translating action message

Table 19–112 (Cont.) MGMT\$ALERT_HISTORY

Column	Description
VIOLATION_TYPE	An intuitive description of the type of violation. Possible values are: <ul style="list-style-type: none"> Threshold Violation: When the alert is triggered based on a metric threshold Availability: When the alert is triggered for an availability metric Policy Violation: When the alert is triggered based on a policy violation
TYPE_DISPLAY_NAME	The display name of the target type

19.18.7 MGMT\$METRIC_DETAILS

The MGMT\$METRIC_DETAILS view displays a rolling 7 day window of individual metric samples. These are the metric values for the most recent sample that has been loaded into the Management Repository plus any earlier samples that have not been aggregated into hourly statistics.

Table 19–113 MGMT\$METRIC_DETAILS

Column	Description
TARGET_NAME	Name of the target where the metric was collected. The target name uniquely identifies a managed target within the Management Repository. The target name typically contains the name of the managed entity that was provided by the system or database administrator.
TARGET_TYPE	The target type defines the set of metrics that are applicable for the target
TARGET_GUID	The unique global identifier for the target
METRIC_NAME	Name of the metric being defined
METRIC_TYPE	A DECODE of the internal numeric type of the metric that is being defined. This column will contain one of the following values: <ul style="list-style-type: none"> Number String Table Raw External
METRIC_COLUMN	For table metrics, the metric column contains the name of the column in the table that is being defined. If the metric that is being defined is not a table metric, the value in this column is a single space. For example, if this table describing the MGMT\$TARGET_TYPE view was being defined as a table metric, Column Name, Data Type, and Description would be metric columns.
METRIC_LABEL	An intuitive display name for the metric that is being defined
COLUMN_LABEL	For table metrics, the column label contains an intuitive display name for the metric column
COLLECTION_TIMESTAMP	The date-time when the alert condition was detected by the Management Agent

Table 19–113 (Cont.) MGMT\$METRIC_DETAILS

Column	Description
VALUE	Since current metric values can be a numeric or a string type, this column returns the value of the metric as a string. If the user of the view is restricting the query to numeric metric values, they can use the TO_NUMBER SQL function to return the values in numeric form.
KEY_VALUE	The key value for which the alert has been recorded. For composite keys, this is the first part of the key.
KEY_VALUE2	For composite keys, this is the second part of the key for which the alert has been recorded
KEY_VALUE3	For composite keys, this is the third part of the key for which the alert has been recorded
KEY_VALUE4	For composite keys, this is the fourth part of the key for which the alert has been recorded
KEY_VALUE5	For composite keys, this is the fifth part of the key for which the alert has been recorded

Usage Notes

- Show the individual values for a metric over time.
- Identify time periods when abnormal samples for metric were collected.
- Calculate the correlation coefficient between two or more metrics.
- Provide metric values that are associated with an alert.
- Queries using this view will use an index if the queries use the target name, the target type, metric name, metric column, and key value, or if they are based upon the collection_timestamp.

19.18.8 MGMT\$METRIC_CURRENT

The MGMT\$METRIC_CURRENT view displays information on the most recent metric values that have been loaded into the Management Repository.

Table 19–114 MGMT\$METRIC_CURRENT

Column	Description
TARGET_NAME	Name of the target where the metric was collected. The target name uniquely identifies a managed target within the Management Repository. The target name typically contains the name of the managed entity that was provided by the system or database administrator.
TARGET_TYPE	The target type defines the set of metrics that are applicable for the target
TARGET_GUID	The unique global identifier for the target
METRIC_NAME	Name of the metric being defined

Table 19–114 (Cont.) MGMT\$METRIC_CURRENT

Column	Description
METRIC_TYPE	<p>A DECODE of the internal numeric type of the metric that is being defined. This column will contain one of the following values:</p> <ul style="list-style-type: none"> ■ Number ■ String ■ Table ■ Raw ■ External
METRIC_COLUMN	<p>For table metrics, the metric column contains the name of the column in the table that is being defined. If the metric that is being defined is not a table metric, the value in this column is a single space.</p> <p>For example, if this table describing the MGMT\$TARGET_TYPE view was being defined as a table metric, Column Name, Data Type, and Description would be metric columns.</p>
METRIC_LABEL	An intuitive display name for the metric that is being defined
COLUMN_LABEL	For table metrics, the column label contains an intuitive display name for the metric column
COLLECTION_TIMESTAMP	The date-time when the alert condition was detected by the Management Agent
VALUE	Since current metric values can be a numeric or a string type, this column returns the value of the metric as a string. If the user of the view is restricting the query to numeric metric values, they can use the TO_NUMBER SQL function to return the values in numeric form.
KEY_VALUE	The key value for which the alert has been recorded. For composite keys, this is the first part of the key.
KEY_VALUE2	For composite keys, this is the second part of the key for which the alert has been recorded
KEY_VALUE3	For composite keys, this is the third part of the key for which the alert has been recorded
KEY_VALUE4	For composite keys, this is the fourth part of the key for which the alert has been recorded
KEY_VALUE5	For composite keys, this is the fifth part of the key for which the alert has been recorded

Usage Notes

- Retrieve the most recent value for a metric that is stored in the Management Repository.
- Retrieve the latest metrics for a target or metric for a specific time period.
- Queries using this view will use an index if the queries use target name, the target type, metric name, metric column, and key value, or if they are based upon the collection_timestamp.

19.18.9 MGMT\$METRIC_HOURLY

The MGMT\$METRIC_HOURLY view displays metric statistics information that have been aggregated from the individual metric samples into hourly time periods. For

example, if a metric is collected every 15 minutes, the 1 hour rollup would aggregate the 4 samples into a single hourly value by averaging the 4 individual samples together. The current hour of statistics may not be immediately available from this view. The timeliness of the information provided from this view is dependent on when the query against the view was executed and when the hourly rollup table was last refreshed.

Table 19–115 *MGMT\$METRIC_HOURLY*

Column	Description
TARGET_NAME	Name of the target where the metric was collected. The target name uniquely identifies a managed target within the Management Repository. The target name typically contains the name of the managed entity that was provided by the system or database administrator.
TARGET_TYPE	The target type defines the set of metrics that are applicable for the target
TARGET_GUID	The unique global identifier for the target
METRIC_NAME	Name of the metric being defined
METRIC_COLUMN	For table metrics, the metric column contains the name of the column in the table that is being defined. If the metric that is being defined is not a table metric, the value in this column is a single space. For example, if this table describing the MGMT\$TARGET_TYPE view was being defined as a table metric, Column Name, Data Type, and Description would be metric columns.
METRIC_LABEL	An intuitive display name for the metric that is being defined
COLUMN_LABEL	For table metrics, the column label contains an intuitive display name for the metric column
KEY_VALUE	The key value for which the alert has been recorded. For composite keys, this is the first part of the key
KEY_VALUE2	For composite keys, this is the second part of the key for which the alert has been recorded
KEY_VALUE3	For composite keys, this is the third part of the key for which the alert has been recorded
KEY_VALUE4	For composite keys, this is the fourth part of the key for which the alert has been recorded
KEY_VALUE5	For composite keys, this is the fifth part of the key for which the alert has been recorded
ROLLUP_TIMESTAMP	The rollup timestamp identifies the start of the rollup period. For the one-hour rollups, samples that fall within the hourly boundaries from minute 00 through minute 59 inclusive will be combined. For example, samples from 12:00 AM through 12:59 AM would be combined into a single aggregated record with a rollup timestamp of "date" 12:00 AM.
SAMPLE_COUNT	The number of non NULL samples for the metric that were aggregated
AVERAGE	The average of the metric values for the samples that have been included in the rollup period
MINIMUM	The minimum value for the metric for the samples that have been included in the rollup period

Table 19–115 (Cont.) MGMT\$METRIC_HOURLY

Column	Description
MAXIMUM	The maximum value for the metric for samples that have been included in the rollup period
STANDARD_DEVIATION	The standard deviation for the metric values that have been included in the rollup period

Usage Notes

- This view provides the best level of granularity to show changes in a metric's value over the course of a day.
- Identify hourly time periods when a metric or sets of metrics are maximized.
- Understand how the variability of a metric over a one hour time period.
- Identify the values of the collected metrics for a target when a particular hour has been identified as problematic.
- Queries using this view will use an index if the queries use the target_name, the metric_name, or if they are based upon the rollup_timestamp.

19.18.10 MGMT\$METRIC_DAILY

The MGMT\$METRIC_DAILY view displays metric statistics that have been aggregated from the samples collected over the previous twenty-four hour time period. The timeliness of the information provided from this view is dependent on when the query against the view was executed and when the hourly rollup table was last refreshed.

Table 19–116 MGMT\$METRIC_DAILY

Column	Description
TARGET_NAME	Name of the target where the metric was collected. The target name uniquely identifies a managed target within the Management Repository. The target name typically contains the name of the managed entity that was provided by the system or database administrator.
TARGET_TYPE	The target type defines the set of metrics that are applicable for the target
TARGET_GUID	The unique global identifier for the target
METRIC_NAME	Name of the metric being defined
METRIC_COLUMN	For table metrics, the metric column contains the name of the column in the table that is being defined. If the metric that is being defined is not a table metric, the value in this column is a single space. For example, if this table describing the MGMT\$TARGET_TYPE view was being defined as a table metric, Column Name, Data Type, and Description would be metric columns.
METRIC_LABEL	An intuitive display name for the metric that is being defined
COLUMN_LABEL	For table metrics, the column label contains an intuitive display name for the metric column
KEY_VALUE	The key value for which the alert has been recorded. For composite keys, this is the first part of the key

Table 19–116 (Cont.) MGMT\$METRIC_DAILY

Column	Description
KEY_VALUE2	For composite keys, this is the second part of the key for which the alert has been recorded
KEY_VALUE3	For composite keys, this is the third part of the key for which the alert has been recorded
KEY_VALUE4	For composite keys, this is the fourth part of the key for which the alert has been recorded
KEY_VALUE5	For composite keys, this is the fifth part of the key for which the alert has been recorded
ROLLUP_TIMESTAMP	The rollup timestamp identifies the start of the rollup period. For the one-hour rollups, samples that fall within the hourly boundaries from minute 00 through minute 59 inclusive will be combined. For example, samples from 12:00 AM through 12:59 AM would be combined into a single aggregated record with a rollup timestamp of "date" 12:00 AM.
SAMPLE_COUNT	The number of non-NULL samples for the metric that were aggregated
AVERAGE	The average of the metric values for the samples that have been included in the rollup period
MINIMUM	The minimum value for the metric for the samples that have been included in the rollup period
MAXIMUM	The maximum value for the metric for samples that have been included in the rollup period
STANDARD_DEVIATION	The standard deviation for the metric values that have been included in the rollup period

Usage Notes

- This view provides the best granularity to show changes in a metric's value over the course of a week or month.
- Understand trends in metric values.
- Queries using this view will use an index if the queries use the target_name, the metric_name, or if they are based upon the rollup_timestamp.

19.19 Operating System Views

This section provides a description of each Operating System view and its columns.

For examples of how to use these views, see [Section 19.30, "Examples"](#).

19.19.1 MGMT\$OS_SUMMARY

The MGMT\$OS_SUMMARY view contains the summary of targets installed in the Oracle home directories.

Table 19–117 MGMT\$OS_SUMMARY

Column	Description
VENDOR_NAME	The name of the vendor
BASE_VERSION	The OS base version
UPDATE_LEVEL	The OS update level

Table 19–117 (Cont.) MGMT\$OS_SUMMARY

Column	Description
DISTRIBUTOR_VERSION	The OS distributor version
MAX_SWAP_SPACE_IN_MB	The maximum amount of swap space
SNAPSHOT_GUID	The globally unique identifier of the operating system snapshot
ADDRESS_LENGTH_IN_BITS	The OS address length in bits
TARGET_GUID	The globally unique identifier of the target
PLATFORM_ID	The platform ID of the host
TARGET_NAME	The name of the target
TARGET_TYPE	The type of the target
START_TIMESTAMP	The time when the target availability status change was first detected.
RUN_LEVEL	The run level of the operating system
DEFAULT_RUN_LEVEL	The default run level of the operating system
PLATFORM_VERSION_ID	The platform version ID number of the application system
DBM_MEMBER	Indicates whether the host is part of an Exadata configuration
EXALOGIC_MEMBER	Indicates whether the host is part of an Exalogic configuration

19.19.2 MGMT\$OS_COMPONENTS

The MGMT\$OS_COMPONENTS view returns performance information for host OS components.

Table 19–118 MGMT\$OS_COMPONENTS

Column	Description
TARGET_NAME	The name of this target
COMPONENT_NAME	The name of the software component.
TARGET_TYPE	The type of target for this view
TARGET_GUID	The globally unique identifier for the target
SNAPSHOT_GUID	The globally unique identifier for the snapshot
START_TIMESTAMP	The date-time when the data was first collected
INSTALLATION_DATE	The installation date of the component
VERSION	The version of the component
DESCRIPTION	The description of the component

19.19.3 MGMT\$OS_HW_SUMMARY

The MGMT\$OS_HW_SUMMARY view displays summary information for both operating systems and hardware.

Table 19–119 MGMT\$OS_HW_SUMMARY

Column	Description
TARGET_NAME	Type of the target for this metric
DOMAIN	The domain of the host
OS_NAME	The operating system name
SYSTEM_CONFIGURATION	A summary of the system configuration information
MACHINE_ARCHITECTURE	A summary of the system architecture
CLOCK_FREQUENCY_IN_MHZ	The clock frequency measured in MHz
MEMORY_SIZE_IN_MB	The memory size measured in MB
LOCAL_DISK_SPACE_IN_GB	The local disk space measured in GBs
CPU_COUNT	The number of CPUs
HARDWARE_VENDOR_NAME	The name of the hardware vendor
OS_VENDOR_NAME	The name of the system vendor
OS_DISTRIBUTOR_VERSION	The distribution version
SNAPSHOT_GUID	The globally unique identifier of the configuration snapshot
TARGET_GUID	The globally unique identifier of the target
PHYSICAL_CPU_COUNT	The number of physical CPUs
LOGICAL_CPU_COUNT	The number of logical CPUs
PLATFORM_ID	The identification number of the platform
TARGET_TYPE	The type of target
LAST_COLLECTION_TIMESTAMP	The date-time of the last collection
OS_RUN_LEVEL	The run level of the operating system
OS_DEFAULT_RUN_LEVEL	The default run level of the operating system
HOST_ID	The host ID number
OS_PLATFORM_VERSION_ID	The operating system platform version number
OS_DBM_MEMBER	Indicates whether the host is part of an Exadata configuration
OS_EXALOGIC_MEMBER	Indicates whether the host is part of an Exalogic configuration
VIRTUAL	The identification for the given host is virtual or physical
SYSTEM_SERIAL_NUMBER	The system serial number of the host

19.19.4 MGMT\$OS_PATCH_SUMMARY

The MGMT\$OS_PATCH_SUMMARY view provides a summary of the patches applied to the operating system.

Table 19–120 MGMT\$OS_PATCH_SUMMARY

Column	Description
TARGET_NAME	Type of the target for this metric
VENDOR_NAME	The name of the vendor
BASE_VERSION	The base version of the operating system
UPDATE_LEVEL	The update level of the operating system
DISTRIBUTOR_VERSION	The distributor version of the OS
MAX_SWAP_SPACE_IN_MB	The maximum swap space measured in MB
SNAPSHOT_GUID	The globally unique identifier of the snapshot
NUM_PATCHES	The number of OS patches found
TARGET_GUID	The globally unique identifier of the target
TARGET_TYPE	The type of the target
START_TIMESTAMP	The date-time of the last collection
NAME	The name of the patch

19.19.5 MGMT\$OS_FS_MOUNT

The MGMT\$OS_FS_MOUNT view displays performance information for mounted file systems.

Table 19–121 MGMT\$OS_FS_MOUNT

Column	Description
TARGET_TYPE	Type of the target for this metric
RESOURCE_NAME	The name of the mounted resource
TYPE	The file system mount
MOUNT_LOCATION	The mount location
MOUNT_OPTIONS	The mount options
SNAPSHOT_GUID	The globally unique identifier of the snapshot
TARGET_GUID	The globally unique identifier of the target
START_TIMESTAMP	The date-time of the last collection

19.19.6 MGMT\$OS_KERNEL_PARAMS

The MGMT\$OS_KERNEL_PARAMS view returns a summary for operating system kernel parameters.

Table 19–122 MGMT\$OS_KERNEL_PARAMS

Column	Description
TARGET_TYPE	Type of the target for this metric
TARGET_NAME	The name of this target
TARGET_GUID	The globally unique identifier of the target
VALUE	The value of the parameter

Table 19–122 (Cont.) MGMT\$OS_KERNEL_PARAMS

Column	Description
NAME	The name of the parameter
SOURCE	The source of the parameter
START_TIMESTAMP	The date-time of the last collection
HOST	The name of the host

19.19.7 MGMT\$OS_PATCHES

The MGMT\$OS_PATCHES view returns a summary of the operating system patches.

Table 19–123 MGMT\$OS_PATCHES

Column	Description
TARGET_TYPE	The type of target
TARGET_NAME	The name of the target
TARGET_GUID	The globally unique identifier for the configuration target
START_TIMESTAMP	The time-date of the last collection
SNAPSHOT_GUID	The globally unique identifier of the snapshot
VENDOR_NAME	The name of the vendor
NAME	The name of the patch

19.19.8 MGMT\$OS_PROPERTIES

The MGMT\$OS_PROPERTIES view returns a summary of the operating system properties.

Table 19–124 MGMT\$OS_PROPERTIES

Column	Description
TARGET_TYPE	Type of the target for this metric
TARGET_NAME	The name of the target
NAME	The name of the property
SOURCE	The source of the property
VALUE	The value of the property
SNAPSHOT_GUID	The globally unique identifier of the snapshot
START_TIMESTAMP	The date-time of the last collection

19.19.9 MGMT\$OS_MODULES

The MGMT\$OS_MODULES view returns a summary of the operating system module details.

Table 19–125 MGMT\$OS_MODULES

Column	Description
TARGET_TYPE	Type of the target for this metric
NAME	The name of the module

Table 19–125 (Cont.) MGMT\$OS_MODULES

Column	Description
SIZE_IN_BYTES	The size in bytes of the module
REFERRING_MODULES	The list of referring modules
SNAPSHOT_GUID	The globally unique identifier of the snapshot
TARGET_NAME	The name of the target
START_TIMESTAMP	The date-time of the last collection

19.19.10 MGMT\$OS_LIMITS

The MGMT\$OS_LIMITS view returns a summary of operating system limit values

Table 19–126 MGMT\$OS_LIMITS

Column	Description
TARGET_TYPE	Type of target for this metric
SNAPSHOT_GUID	The globally unique identifier of the snapshot
TARGET_NAME	The target name
TARGET_GUID	The globally unique identifier of the target
START_TIMESTAMP	The date-time of the last collection

19.19.11 MGMT\$OS_INIT_SERVICES

The MGMT\$OS_INIT_SERVICES view returns a summary of operating system init service details.

Table 19–127 MGMT\$OS_INIT_SERVICES

Column	Description
TARGET_TYPE	Type of target for this metric
TARGET_NAME	The name of the target
APPLICATION_ID	The application ID of the service
RUN_STATE	The run state of the service
MAPPER_VERSION	The mapper version of the service
SNAPSHOT_GUID	The globally unique identifier of the snapshot
TARGET_GUID	The globally unique identifier of the target
START_TIMESTAMP	The date-time of the last collection

19.20 Oracle Home Directory Patching Views

This section provides a description of each Oracle home directory patching view.

19.20.1 MGMT\$EM_HOMES_PLATFORM

The MGMT\$EM_HOMES_PLATFORM view displays the platform information about the home directories. If the home directory does not have an ARU platform ID, then the platform of the Operating System is considered as the platform of the home directory.

Table 19–128 MGMT\$EM_HOMES_PLATFORM

Column	Description
HOME_ID	Unique ID for the home directory
PLATFORM_ID	If the home directory has an ARU platform it is used, otherwise the platform ID of the host is picked
PLATFORM	The platform corresponding to the platform_id

19.20.2 MGMT\$HOMES_AFFECTED

The MGMT\$HOMES_AFFECTED view displays the list of home directories, vulnerable to bugs, which are fixed by the critical patches released. The number of alerts which are applicable to the home directory are calculated.

Table 19–129 MGMT\$HOMES_AFFECTED

Column	Description
HOST	Host name
HOME_DIRECTORY	Home directory location
TARGET_GUID	Unique ID for target
ALERTS	Number of alerts for this home directory

19.20.3 MGMT\$APPL_PATCH_AND_PATCHSET

The MGMT\$APPL_PATCH_AND_PATCHSET view displays the list of interim patches and patchsets that are applicable to the home directories.

Table 19–130 MGMT\$APPL_PATCH_AND_PATCHSET

Column	Description
PATCH_ID	The patch ID
TYPE	Patch or patchset
PRODUCT	The product pertaining to the patch
PATCH_RELEASE	Release version
PLATFORM	Platform on which patch is applicable
ADVISORY	The alert name
HOST_NAME	Host name
HOME_LOCATION	Home directory location
PATCH_GUID	Unique ID for the patch or patchset
TARGET_GUID	Unique ID for the target

19.20.4 MGMT\$APPLIED_PATCHES

The MGMT\$APPLIED_PATCHES view displays the list of patches that have been applied on the home directories along with the installation time. Each patch can fix more than one bug. The bugs are listed in a comma-separated string.

Table 19–131 MGMT\$APPLIED_PATCHES

Column	Description
PATCH	Patch name
BUGS	The bugs fixed by this patch
INSTALLATION_TIME	Time of installation (time zone of the target)
HOST	Host name
HOME_LOCATION	Home location
HOME_NAME	Name of the home
CONTAINER_GUID	Name of the home
TARGET_GUID	Unique ID for target

19.20.5 MGMT\$APPLIED_PATCHSETS

The MGMT\$APPLIED_PATCHSETS view displays the list of patchsets that have been applied on the home directories along with the installation time.

Table 19–132 MGMT\$APPLIED_PATCHSETS

Column	Description
VERSION	The version to which the home will get upgraded to when this patchset is applied
NAME	Patchset external name
TIMESTAMP	Time of Installation (time zone of the target)
HOST	Host name
HOME_LOCATION	Home location
HOME_NAME	Name of the home directory
CONTAINER_GUID	Name of the home directory
TARGET_GUID	Unique ID for target

19.21 Oracle Home Directory Views

This section provides a description of each Oracle home directory view and its columns.

19.21.1 MGMT\$OH_HOME_INFO

The MGMT\$OH_HOME_INFO view contains properties of the Oracle home targets.

Table 19–133 MGMT\$OH_HOME_INFO

Column	Description
ECM_SNAPSHOT_ID	ECM Snapshot ID of the current snapshot for Oracle home target
TARGET_NAME	Name of Oracle home target
TARGET_GUID	Oracle home target globally unique identifier
HOST_NAME	Name of the host on which the Oracle home directory is installed
EMD_URL	EMD_URL of the agent monitoring this Oracle home target

Table 19–133 (Cont.) MGMT\$OH_HOME_INFO

Column	Description
HOME_LOCATION	Complete path to the Oracle home directory
OUI_HOME_NAME	OUI home name
OUI_HOME_GUID	OUI Oracle home globally unique identifier. This is unique across all Oracle product installations.
HOME_TYPE	Type of the HOME ('O' [OUI] or 'W' [WebLogic])
HOME_POINTER	OUI Central Inventory / Composite Home / BEA Home that contains this home
IS_CLONABLE	Is this home clonable? [0/1]
IS_CRS	Is it a Cluster Ready Services (CRS) home [0/1]
ARU_ID	ARU Platform ID of the Oracle home directory
OUI_PLATFORM_ID	OUI Platform ID of the host
HOME_SIZE	Size of the Oracle home directory (in KBytes)
HOME_RW_STATUS	Read write status of home[NRNRW/RO/WO/RW]
ORACLE_BASE	Oracle Base (for OUI homes only)
OH_OWNER_ID	Oracle home owner ID
OH_OWNER	Oracle home owner
OH_GROUP_ID	Oracle home group ID
OH_GROUP	Oracle home group
OH_OWNER_GROUPS_ID	Semi colon separated list of groups IDs to which the Oracle home owner belong
OH_OWNER_GROUPS	Semi colon separated list of groups to which the Oracle home owner belongs

19.21.2 MGMT\$OH_DEP_HOMES

The MGMT\$OH_DEP_HOMES view contains information about other homes on which an Oracle home depends.

Table 19–134 MGMT\$OH_DEP_HOMES

Column	Description
ECM_SNAPSHOT_ID	ECM Snapshot ID of the current snapshot for Oracle home target
TARGET_NAME	Name of Oracle home target
TARGET_GUID	Oracle home target globally unique identifier
HOST_NAME	Name of the host on which Oracle home is installed
HOME_LOCATION	Complete path to the Oracle home
HOME_NAME	OUI home name
DEP_HOME_LOCATION	Install location of dependee home

19.21.3 MGMT\$OH_CRS_NODES

The MGMT\$OH_CRS_NODES view contains information about member nodes of a CRS Oracle home.

Table 19–135 MGMT\$OH_CRS_NODES

Column	Description
ECM_SNAPSHOT_ID	ECM Snapshot ID of the current snapshot for Oracle home target
TARGET_NAME	Name of Oracle home target
TARGET_GUID	Oracle home target globally unique identifier
HOST_NAME	Name of the host on which Oracle home is installed
HOME_LOCATION	Complete path to the Oracle home
HOME_NAME	OUI home name
NODE	Node host name

19.21.4 MGMT\$OH_CLONE_PROPERTIES

The MGMT\$OH_CLONE_PROPERTIES view contains information about clone properties of an Oracle home.

Table 19–136 MGMT\$OH_CLONE_PROPERTIES

Column	Description
ECM_SNAPSHOT_ID	ECM Snapshot ID of the current snapshot for the Oracle home target
TARGET_NAME	Name of the Oracle home target
TARGET_GUID	Oracle home target globally unique identifier
HOST_NAME	Name of the host on which the Oracle home is installed
HOME_LOCATION	Complete path to the Oracle home
HOME_NAME	OUI home name
PROPERTY_NAME	Clone property name
PROPERTY_VALUE	Property value

19.21.5 MGMT\$OH_COMPONENT

The MGMT\$OH_COMPONENT view contains information about components installed in an Oracle home.

Table 19–137 MGMT\$OH_COMPONENT

Column	Description
ECM_SNAPSHOT_ID	ECM Snapshot ID of the current snapshot for the Oracle home target
TARGET_NAME	Name of the Oracle home target
TARGET_GUID	Oracle home target globally unique identifier
HOST_NAME	Name of the host on which the Oracle home is installed
HOME_LOCATION	Complete path to the Oracle Home
HOME_NAME	OUI home name
COMPONENT_NAME	Component name
VERSION	Current version of component
BASE_VERSION	Component base version

Table 19–137 (Cont.) MGMT\$OH_COMPONENT

Column	Description
INSTALL_TIME	Installation time of component
IS_TOP_LEVEL	Is it a top level component [0/1]
EXTERNAL_NAME	External name of the component
DESCRIPTION	A brief description of the component
LANGUAGES	Languages supported by this component installation
INSTALLED_LOCATION	Component install location
INSTALLER_VERSION	Installer version
MIN_DEINSTALLER_VERSION	Minimum OUI version required to deinstall this component

19.21.6 MGMT\$OH_COMP_INST_TYPE

The MGMT\$OH_COMP_INST_TYPE view contains Install Type information about components installed in an Oracle home directory.

Table 19–138 MGMT\$OH_COMP_INST_TYPE

Column	Description
ECM_SNAPSHOT_ID	ECM Snapshot ID of the current snapshot for the Oracle home target
TARGET_NAME	Name of the Oracle home target
TARGET_GUID	Oracle home target globally unique identifier
HOST_NAME	Name of the host on which Oracle home is installed
HOME_LOCATION	Complete path to the Oracle home
HOME_NAME	OUI home name
COMPONENT_NAME	Component name
COMPONENT_VERSION	Component base version
NAME_ID	Install type name ID
INSTALL_TYPE_NAME	Install type name
DESC_ID	Install type desc ID

19.21.7 MGMT\$OH_COMP_DEP_RULE

The MGMT\$OH_COMP_DEP_RULE view contains information about a dependency relationship between components installed in an Oracle home.

Table 19–139 MGMT\$OH_COMP_DEP_RULE

Column	Description
ECM_SNAPSHOT_ID	ECM Snapshot ID of the current snapshot for the Oracle home target
TARGET_NAME	Name of the Oracle home target
TARGET_GUID	Oracle home target globally unique identifier
HOST_NAME	Name of the host on which the Oracle home is installed

Table 19–139 (Cont.) MGMT\$OH_COMP_DEP_RULE

Column	Description
HOME_LOCATION	Complete path to the Oracle home
HOME_NAME	OUI home name
COMPONENT_NAME	Component name
COMPONENT_VERSION	Component base version
DEPENDEE_NAME	Dependee component name
DEPENDEE_VERSION	Dependee component version
DEPENDEE_HOME_GUID	Oracle home dependee component globally unique identifier

19.21.8 MGMT\$OH_PATCHSET

The MGMT\$OH_PATCHSET view contains information about patchsets applied on an Oracle home directory.

Table 19–140 MGMT\$OH_PATCHSET

Column	Description
ECM_SNAPSHOT_ID	ECM Snapshot ID of the current snapshot for the Oracle home target
TARGET_NAME	Name of the Oracle home target
TARGET_GUID	Oracle home target globally unique identifier
HOST_NAME	Name of the host on which the Oracle home is installed
HOME_LOCATION	Complete path to the Oracle home
HOME_NAME	OUI home name
PATCHSET_NAME	Patchset name
PATCHSET_VERSION	Patchset version
INSTALL_TIME	Installation time of patchset
EXTERNAL_NAME	External name of the patchset
DESCRIPTION	A brief description of the patchset
INV_LOCATION	Patchset inventory location
INSTALLER_VERSION	Installer version
MIN_DEINSTALLER_VERSION	Minimum OUI version required to deinstall this patchset

19.21.9 MGMT\$OH_VERSIONED_PATCH

The MGMT\$OH_VERSIONED_PATCH view contains information about versioned patches applied on an Oracle home directory.

Table 19–141 MGMT\$OH_VERSIONED_PATCH

Column	Description
ECM_SNAPSHOT_ID	ECM Snapshot ID of the current snapshot for the Oracle home target
TARGET_NAME	Name of Oracle home target

Table 19–141 (Cont.) MGMT\$OH_VERSIONED_PATCH

Column	Description
TARGET_GUID	Oracle home target globally unique identifier
HOST_NAME	Name of the host on which the Oracle home is installed
HOME_LOCATION	Complete path to the Oracle home
HOME_NAME	OUI home name
VPATCH_NAME	Versioned patch name (should be same as the component name, on which the versioned patch is applied)
VPATCH_VERSION	Versioned patch version
BASE_COMP_VERSION	Base component version, on which the versioned patch is applied
PATCHSET_NAME	Name of the patchset this versioned patch is part of
PATCHSET_VERSION	Version of the patchset this versioned patch is part of
INSTALL_TIME	Installation time of the versioned patch
EXTERNAL_NAME	External name of the versioned patch
DESCRIPTION	A brief description of the versioned patch
LANGUAGES	Languages supported by this versioned patch
INSTALLED_LOCATION	Install location of the versioned patch
INSTALLER_VERSION	Installer version
MIN_DEINSTALLER_VERSION	Minimum OUI version required to remove this versioned patch

19.21.10 MGMT\$OH_PATCH

The MGMT\$OH_PATCH view contains information about patches applied on an Oracle home directory.

Table 19–142 MGMT\$OH_PATCH

Column	Description
ECM_SNAPSHOT_ID	ECM Snapshot ID of the current snapshot for the Oracle home target
TARGET_NAME	Name of the Oracle home target
TARGET_GUID	Oracle home target globally unique identifier
HOST_NAME	Name of the host on which the Oracle home is installed
HOME_LOCATION	Complete path to the Oracle home
HOME_NAME	OUI home name
PATCH_ID	Patch ID (may be same for more than one patches)
PATCH_UPI	Unique patch identifier (putting N/A when not available in metadata)
PATCH_LANG	Patch language
BUGS_FIXED	Comma separated list of bugs fixed by the patch
INSTALL_TIME	Installation time of the patch
IS_ROLLBACKABLE	Can the patch be rolled back? [0/1]

Table 19–142 (Cont.) MGMT\$OH_PATCH

Column	Description
IS_PSU	Is it a PSU? [0/1]
PROFILE	Profile used to install the patch (only for WebLogic)
PATCH_TYPE	Patch type
DESCRIPTION	A brief description of the patch
XML_INV_LOCATION	Patch XML inventory location
INSTALLER_VERSION	Installer version of the patch

19.21.11 MGMT\$OH_PATCHED_COMPONENT

The MGMT\$OH_PATCHED_COMPONENT view contains information about components affected by a patch applied on an Oracle home directory.

Table 19–143 MGMT\$OH_PATCHED_COMPONENT

Column	Description
ECM_SNAPSHOT_ID	ECM Snapshot ID of the current snapshot for the Oracle home target
TARGET_NAME	Name of the Oracle home target
TARGET_GUID	Oracle home target globally unique identifier
HOST_NAME	Name of the host on which the Oracle home is installed
HOME_LOCATION	Complete path to the Oracle home
HOME_NAME	OUI home name
PATCH_ID	Patch ID (may be the same for more than one patch)
PATCH_UPI	Unique Patch Identifier (putting N/A when not available in metadata)
PATCH_LANG	Patch language
COMPONENT_NAME	Affected component name
COMPONENT_VERSION	Current version of the affected component
COMPONENT_BASE_VERSION	Base version of the affected component
COMPONENT_EXTERNAL_NAME	External name of the affected component
FROM_VERSION	Version of the affected component before applying PSU
TO_VERSION	Version of the affected component after applying PSU

19.21.12 MGMT\$OH_PATCH_FIXED_BUG

The MGMT\$OH_PATCH_FIXED_BUG view contains information about bugs fixed by a patch applied on an Oracle home directory.

Table 19–144 MGMT\$OH_PATCH_FIXED_BUG

Column	Description
ECM_SNAPSHOT_ID	ECM Snapshot ID of the current snapshot for the Oracle home target

Table 19–144 (Cont.) MGMT\$OH_PATCH_FIXED_BUG

Column	Description
TARGET_NAME	Name of the Oracle home target
TARGET_GUID	Oracle home target globally unique identifier
HOST_NAME	Name of the host on which the Oracle home is installed
HOME_LOCATION	Complete path to the Oracle home
HOME_NAME	OUI home name
PATCH_ID	Patch ID (may be the same for more than one patch)
PATCH_UPI	Unique Patch Identifier (putting N/A when not available in metadata)
PATCH_LANG	Patch language
BUG_NUMBER	Bug number of a bug fixed by a patch
BUG_DESC	Bug description of a bug fixed by a patch

19.21.13 MGMT\$OH_PATCHED_FILE

The MGMT\$OH_PATCHED_FILE view contains information about the files affected by a patch applied on an Oracle home directory.

Table 19–145 MGMT\$OH_PATCHED_FILE

Column	Description
ECM_SNAPSHOT_ID	ECM Snapshot ID of the current snapshot for the Oracle home target
TARGET_NAME	Name of the Oracle home target
TARGET_GUID	Oracle home target globally unique identifier
HOST_NAME	Name of the host on which the Oracle home is installed
HOME_LOCATION	Complete path to the Oracle home
HOME_NAME	OUI home name
PATCH_ID	Patch ID (may be the same for more than one patch)
PATCH_UPI	Unique Patch Identifier (putting N/A when not available in metadata)
PATCH_LANG	Installation time of the patchset
TIMESTAMP	Patch timestamp
FILE_NAME	Name of a patched file
COMP_NAME	Name of the OUI component this file is part of
COMP_VERSION	InstaVersion of the OUI component this file is part of

19.21.14 MGMT\$OH_FILE

The MGMT\$OH_FILE view contains information about all the files affected by one or more patches applied on an Oracle home directory.

Table 19–146 MGMT\$OH_FILE

Column	Description
ECM_SNAPSHOT_ID	ECM Snapshot ID of the current snapshot for the Oracle home target
TARGET_NAME	Name of the Oracle home target
TARGET_GUID	Oracle home target globally unique identifier
HOST_NAME	Name of the host on which the Oracle home is installed
HOME_LOCATION	Complete path to the Oracle home
HOME_NAME	OUI home name
FILE_NAME	Patched file name
LAST_PATCH_ID	Patch ID of the last patch applied on the file
LAST_PATCH_UPI	UPI of the last patch applied on the file
LAST_PATCH_LANG	Language of the last patch applied on the file
LAST_PATCH_TIMESTAMP	Timestamp of the last patch applied on the file

19.21.15 MGMT\$PA_RECOM_METRIC_SOURCE

The MGMT\$PA_RECOM_METRIC_SOURCE view contains data for the patch recommendations metric source.

Table 19–147 MGMT\$PA_RECOM_METRIC_SOURCE

Column	Description
PATCH_GUID	The GUID of the patch
PATCH	The patch number
ABSTRACT	The abstract information of the patch
CLASSIFICATION	The classification information of the patch
PA_TGT_GUID	The GUID of a target which is applicable to the recommended patch
PA_TGT_NAME	The name of a target which is applicable to the recommended patch
PA_TGT_TYPE	The type id of a target, such as 'host' or 'oracle_database'
PA_TGT_TYPE_DISPLAY_NAME	The display name of a target type such as 'Host' or 'Database Instance'
HOST_NAME	The name of a host target which host this target
TARGET_GUID	The GUID of a host target which host this target. It is the target GUID of a policy violation.

19.21.16 MGMT\$OH_INV_SUMMARY

The MGMT\$OH_INV_SUMMARY view contains summary of Oracle products and corresponding target types.

Table 19–148 MGMT\$OH_INV_SUMMARY

Column	Description
ECM_SNAPSHOT_ID	ECM Snapshot ID of the current snapshot for the Oracle home target
OH_TARGET_GUID	Oracle home target globally unique identifier
HOST_NAME	Name of the host on which the Oracle home is installed
HOME_LOCATION	Complete path to the Oracle home
COMP_NAME	Component name
COMP_EXTERNAL_NAME	Component external name
COMP_VERSION	Component version
IS_PATCHED	Is this Oracle home patched?[0/1]
MAP_TARGET_TYPE	Map target type
MAP_PROPERTY_NAME	Map property name
MAP_PROPERTY_VALUE	Map property value

19.21.17 MGMT\$OH_INSTALLED_TARGETS

The MGMT\$OH_INSTALLED_TARGETS view contains summary of targets installed in the Oracle home directories.

Table 19–149 MGMT\$OH_PATCHSET

Column	Description
OH_TARGET_NAME	Name of the Oracle home target
OH_TARGET_GUID	Oracle home target globally unique identifier
HOST_NAME	Name of the host on which the Oracle home is installed
HOME_LOCATION	Complete path to the Oracle home
HOME_TYPE	OUI home type
INST_TARGET_NAME	Installed target name
INST_TARGET_TYPE	Installed target type

19.22 Oracle WebLogic Server Views

This section provides a description of each Oracle WebLogic server view and its columns.

19.22.1 MGMT\$WEBLOGIC_APPLICATIONS

The MGMT\$WEBLOGIC_APPLICATIONS view displays general information about the Application configuration.

Table 19–150 MGMT\$WEBLOGIC_APPLICATIONS

Column	Description
NAME	The name of the application
PATH	The fully resolved location of the application source files on the admin server

Table 19–150 (Cont.) MGMT\$WEBLOGIC_APPLICATIONS

Column	Description
LOADORDER	A numeric value that indicates when the unit is deployed, relative to other DeployableUnits on a server, during startup
TWOPHASE	A boolean value indicating if the application is deployed using the two-phase deployment protocol
TYPE	Type of the module. The string value must match those defined by JSR 88: <i>Java EE Application Deployment</i> such as EAR and WAR

19.22.2 MGMT\$WEBLOGIC_EJBCOMPONENT

The MGMT\$WEBLOGIC_EJBCOMPONENT view displays general information about the EJB modules.

Table 19–151 MGMT\$WEBLOGIC_EJBCOMPONENT

Column	Description
NAME	Name of the EJB component
APPLICATION	Name of the application that includes the component
DEPLOYMENTORDER	Priority that the server uses when it deploys an item. The priority is relative to the other deployable items of same type
KEEPGENERATED	Indicates whether KeepGenerated is enabled and whether EJB source files will be kept. Possible values: <ul style="list-style-type: none"> True: KeepGenerated is enabled and EJB source files are stored. False: KeepGenerated is not enabled and EJB source files are not stored

19.22.3 MGMT\$WEBLOGIC_FILESTORE

Each row of the MGMT\$WEBLOGIC_FILESTORE view represents configuration data of the file store configured for the WebLogic server.

Table 19–152 MGMT\$WEBLOGIC_FILESTORE

Column	Description
ECM_SNAPSHOT_ID	The GUID of the snapshot
NAME	Name of the file store
DIRECTORY	The path name to the file system directory where the file store maintains its data files
SYNCHRONOUSWRITEPOLICY	The disk write policy that determines how the file store writes data to disk
MAXFILESIZE	The maximum file size, in bytes

19.22.4 MGMT\$WEBLOGIC_JDBCdatasource

The MGMT\$WEBLOGIC_JDBCdatasource view displays general information about Java Database Connectivity (JDBC) Data Sources that provides database connectivity through a pool of JDBC connections.

Table 19–153 MGMT\$WEBLOGIC_JDBCdatasource

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target
CM_TARGET_TYPE	The type of target: weblogic_j2eeserver
CM_TARGET_NAME	The name of the target in Enterprise Manager
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
NAME	Name of JDBC Data Source
JNDI_NAME	The Java Naming and Directory Interface (JNDI) path to where this Data Source is bound. By default, the JNDI name is the name of the data source
ROW_PREFETCH_ENABLED	This value is set to TRUE if row prefetching is enabled. Otherwise, the value is set to FALSE.
ROW_PREFETCH_SIZE	The number of result set rows to prefetch for a client if row prefetching is enabled
ENABLE_TWO_PHASE_COMMIT	This value is set to TRUE if two phase commit is enabled. Otherwise, this value is set to FALSE.
URL	The URL of the database to which to connect. The format of the URL varies depending on the JDBC driver.
DRIVER_NAME	The full package name of the JDBC driver class used to create physical database connections in the connection pool in the data source.
CAPACITY_INCREMENT	The number of connections created when new connections are added to the connection pool.
INITIAL_CAPACITY	The number of physical connections to create when creating the connection pool in the data source.
MAX_CAPACITY	The maximum number of physical connections that the connection pool can contain.
CONNECTION_RESERVE_TIMEOUT	The number of seconds after which a call to reserve a connection from the connection pool will time out.
INACTIVE_CONNECTION_TIMEOUT	The number of inactive seconds on a reserved connection before Oracle WebLogic Server reclaims the connection and releases it back to the connection pool.
STATEMENT_CACHE_SIZE	The number of prepared and callable statements stored in the connection cache
HOST	Database host
PORT	Database port
SID	Database system identifier (SID)
SERVICE_NAME	Database service name
PROTOCOL	The communications protocol
ENABLE_JAVA_NET_FASTPATH	Enables the Oracle JDBC JavaNet Fastpath to reduce data copies and fragmentation
OPT_UTF8_CONVERSION	Enables the Oracle JDBC optimize UTF-8 conversion option
STATEMENT_CACHE_TYPE	Statement Cache type parameter from the JDBC connection pool parameters

Table 19–153 (Cont.) MGMT\$WEBLOGIC_JDBCdatasource

Column	Description
PINNEDTOTHREAD	Pinned to thread parameters from the JDBC connection pool parameters

19.22.5 MGMT\$WEBLOGIC_JDBCMULTIDS

The MGMT\$WEBLOGIC_JDBCMULTIDS view displays general information about the JDBC Multi Data Sources.

Table 19–154 MGMT\$WEBLOGIC_JDBCMULTIDS

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target
CM_TARGET_TYPE	The type of target: weblogic_j2eeserver
CM_TARGET_NAME	The name of target in Enterprise Manager
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
NAME	JDBC multi data source name
JNDINAME	The JNDI path to where this data source is bound
ALGORITHMTYPE	The algorithm determines the connection request processing for the multi data source
DATASOURCELIST	The list of data sources to which multi data source routes connection requests. The order of data sources in the list determines the failover order.

19.22.6 MGMT\$WEBLOGIC_JMSCONNFACTORY

The MGMT\$WEBLOGIC_JMSCONNFACTORY view displays general information about the JMS Connection Factory.

Table 19–155 MGMT\$WEBLOGIC_JMSCONNFACTORY

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target
CM_TARGET_TYPE	The type of target: weblogic_j2eeserver
CM_TARGET_NAME	The name of the target in Enterprise Manager
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
NAME	Name of the Java Message Service (JMS) connection factory
MODULENAME	Name of the JMS module
JNDINAME	The global JNDI name used to look up a connection factory within a clustered JNDI namespace
TXNTIMEOUTINSECS	The timeout value (in seconds) for all transactions on connections created with this connection factory
ACKNOWLEDGEPOLICY	Acknowledge policy for non-transacted sessions that use the CLIENT_ACKNOWLEDGE mode

Table 19–155 (Cont.) MGMT\$WEBLOGIC_JMSCONNFACTORY

Column	Description
MESSAGESMAXIMUM	The maximum number of messages that can exist for an asynchronous session and that have not yet been passed to the message listener
SENDTIMEOUT	The maximum length of time, in milliseconds, that a sender will wait when there is not enough available space (no quota) on a destination to accommodate the message being sent

19.22.7 MGMT\$WEBLOGIC_JMSQUEUE

The MGMT\$WEBLOGIC_JMSQUEUE view displays general information about the JMS queue.

Table 19–156 MGMT\$WEBLOGIC_JMSQUEUE

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target
CM_TARGET_TYPE	The type of target: weblogic_j2eeserver
CM_TARGET_NAME	The name of the target in Enterprise Manager
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
NAME	Name of the JMS queue
MODULENAME	Name of the JMS module
JNDI NAME	The global JNDI name used to look up a JMS queue within a clustered JNDI namespace
MAXIMUMMESSAGE SIZE	The maximum size of a message that is accepted from producers on this destination
BYTESMAXIMUM	The maximum size of a message that is accepted from producers on this destination
MESSAGESMAXIMUM	The total number of messages that can be stored in a destination that uses this quota
BYTESPAGINGENABLED	Bytes paging enabled for this JMS server
MESSAGESPAGINGENABLED	Messages paging enabled for this JMS server
STOREENABLED	Store enabled for this JMS server
TARGET	The JMS server to which the JMS queue is targeted

19.22.8 MGMT\$WEBLOGIC_JMSSERVER

The MGMT\$WEBLOGIC_JMSSERVER view displays general information about the JMS server. Each row represents configuration data of the JMS server configured for the WebLogic server.

Table 19–157 MGMT\$WEBLOGIC_JMSSERVER

Column	Description
ECM_SNAPSHOT_ID	GUID of the snapshot
NAME	Name of the JMS server

Table 19–157 (Cont.) MGMT\$WEBLOGIC_JMSSERVER

Column	Description
BYTESMAXIMUM	The maximum number of bytes that can be stored in this JMS server
MESSAGESMAXIMUM	The maximum number of messages that can be stored in this JMS server
MESSAGEBUFFERSIZE	The amount of memory (in bytes) that this JMS server can use to store message bodies before it writes them to disk
MAXIMUMMESSAGE SIZE	The maximum number of bytes allowed in individual messages on this JMS server
PERSISTENTSTORE	The file or database in which this JMS server stores persistent messages
STOREENABLED	Persistent store enabled status

19.22.9 MGMT\$WEBLOGIC_JMSTOPIC

The MGMT\$WEBLOGIC_JMSTOPIC view displays general information about the JMS topic.

Table 19–158 MGMT\$WEBLOGIC_JMSTOPIC

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target
CM_TARGET_TYPE	The type of target: weblogic_j2eeserver
CM_TARGET_NAME	The name of the target in Enterprise Manager
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
NAME	Name of the JMS topic
MODULENAME	Name of the JMS module
JNDINAME	The global JNDI name used to look up a JMS topic within a clustered JNDI namespace
MAXIMUMMESSAGE SIZE	The maximum size of a message that is accepted from producers on this destination
BYTESMAXIMUM	The total number of bytes that can be stored in a destination that uses this quota
MESSAGESMAXIMUM	The total number of messages that can be stored in a destination that uses this quota
MULTICASTPORT	The IP port that this topic uses to transmit messages to multicast consumers
TARGET	The JMS server to which the JMS topic is targeted

19.22.10 MGMT\$WEBLOGIC_JOLTCONNPOOL

The MGMT\$WEBLOGIC_JOLTCONNPOOL view displays general information about the Jolt Connection Pool.

Table 19–159 MGMT\$WEBLOGIC_JOLTCONNPOOL

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target
CM_TARGET_TYPE	The type of target: weblogic_j2eeserver
CM_TARGET_NAME	The name of the target in Enterprise Manager
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
NAME	Name of the Jolt connection pool
PRIMARYADDRESSES	The list of addresses for the primary Jolt Server Listeners (JSLs)
FAILOVERADDRESSES	The list of Jolt Server Listeners (JSL) addresses that is used if the connection pool cannot establish connections to the Primary Addresses, or if the primary connections fail
MINIMUMPOOLSIZE	The minimum number of connections to be added to this Jolt connection pool when the WebLogic Server starts
MAXIMUMPOOLSIZE	The maximum number of connections that can be made from this Jolt connection pool
RECVTIMEOUT	The number of seconds the client waits to receive a response before timing out

19.22.11 MGMT\$WEBLOGIC_JVMSYSPROPS

Each row in the MGMT\$WEBLOGIC_JVMSYSPROPS view represents configuration data of JVM System Properties that are configured for the WebLogic server.

Table 19–160 MGMT\$WEBLOGIC_JVMSYSPROPS

Column	Description
ECM_SNAPSHOT_ID	GUID of the snapshot
KEY	JVM system properties name
VALUE	Value for Operating System (OS) user, name, version, architecture, Java home, and JVM version

19.22.12 MGMT\$WEBLOGIC_MACHINE

The MGMT\$WEBLOGIC_MACHINE view displays general information about the systems.

Table 19–161 MGMT\$WEBLOGIC_MACHINE

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target
CM_TARGET_TYPE	The type of target: weblogic_j2eeserver
CM_TARGET_NAME	The name of the target in Enterprise Manager
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
MACHINENAME	The name of the system

Table 19–161 (Cont.) MGMT\$WEBLOGIC_MACHINE

Column	Description
MACHINETYPE	The type of the system
POSTBINDGID	The UNIX group ID (GID) that a server running on this system will run under after it has carried out all privileged startup actions. Otherwise, the server will continue to run under the group under which it was started. (Requires that you enable Post-Bind GID.)
POSTBINDGIDENABLED	Specifies whether a server running on this system binds to a UNIX Group ID (GID) after it has carried out all privileged startup actions
POSTBINDUID	The UNIX user ID (UID) that a server running on this system will run under after it has carried out all privileged startup actions. Otherwise, the server will continue to run under the account under which it was started. (Requires that you enable Post-Bind UID.)
POSTBINDUIDENABLED	Specifies whether a server running on this system binds to a UNIX User ID (UID) after it has carried out all privileged startup actions

19.22.13 MGMT\$WEBLOGIC_NETWORK_CHANNELS

The MGMT\$WEBLOGIC_NETWORK_CHANNELS view displays general information about the Network Channels.

Table 19–162 MGMT\$WEBLOGIC_NETWORK_CHANNELS

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target
CM_TARGET_TYPE	The type of target: weblogic_j2eeserver
CM_TARGET_NAME	The name of the target in Enterprise Manager
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
NAME	Name of the network channel
LISTEN_ADDRESS	The IP address or DNS name this network channel uses to listen for incoming connections
LISTEN_PORT	The default TCP port this network channel uses to listen for regular (non-SSL) incoming connections
ENABLED	Specifies whether this channel should be started
SDP_ENABLED	Specifies if Socket Direct Protocol (SDP) is enabled on this channel
OUTBOUND_ENABLED	Specifies whether new server-to-server connections can consider this network channel when initiating a connection
CUSTOM_PROPERTIES	Custom properties for the network channel

19.22.14 MGMT\$WEBLOGIC_NODEMANAGER

The MGMT\$WEBLOGIC_NODEMANAGER view displays general information about the Node Manager.

Table 19–163 MGMT\$WEBLOGIC_NODEMANAGER

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target
CM_TARGET_TYPE	The type of target: weblogic_j2eeserver
CM_TARGET_NAME	The name of the target in Enterprise Manager
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
LISTENADDRESS	The host name or IP address where Node Manager listens for connection requests
MACHINENAME	The name of the Node manager system
NMTYPE	Node manager type
LISTENPORT	The port number where Node Manager listens for connection requests
NODEMANAGERUSERNAME	The Node manager user name
STARTSCRIPTENABLED	Value of the StartScriptEnabled property in the nodemanager.properties
NODEMANAGERHOME	Home directory path of the Node Manager

19.22.15 MGMT\$WEBLOGIC_RACONFIG

The MGMT\$WEBLOGIC_RACONFIG view displays general information about the Resource Adapter.

Table 19–164 MGMT\$WEBLOGIC_RACONFIG

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target
CM_TARGET_TYPE	The type of target: weblogic_j2eeserver
CM_TARGET_NAME	The name of the target in Enterprise Manager
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
NAME	Name of the resource adapter
VERSION	Version of the resource adapter
VENDORNAME	Vendor name
EISTYPE	Enterprise Information Systems (EIS) type
RAVERSION	Resource adapter version
ENABLEACCESS	Enable access outside application

19.22.16 MGMT\$WEBLOGIC_RAOUTBOUNDCONFIG

The MGMT\$WEBLOGIC_RAOUTBOUNDCONFIG view displays general information about the Resource Adapter Outbound configuration.

Table 19–165 MGMT\$WEBLOGIC_RAOUTBOUNDCONFIG

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target
CM_TARGET_TYPE	The type of target: weblogic_j2eeserver
CM_TARGET_NAME	The name of the target in Enterprise Manager
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
NAME	Name of the resource adapter
CONNFACINTERFACE	Connection factory interface
MANAGEDCONNFACCLASS	Managed connection factory class
JNDI_NAME	JNDI Name
TRANSACTIONSUPPORT	Specifies the level of transaction support for a particular Connection Factory. It provides the ability to override the transaction-support value specified in the ra.xml deployment descriptor that is intended to be the default value for all Connection Factories of the resource adapter.
INITIALCAPACITY	Specifies the initial number of ManagedConnections, which WebLogic Server attempts to create during deployment
MAXCAPACITY	Specifies the maximum number of ManagedConnections, which WebLogic Server will allow. Requests for newly allocated ManagedConnections beyond this limit results in a ResourceAllocationException being returned to the caller.
CAPACITYINCREMENT	Specifies the maximum number of additional ManagedConnections that WebLogic Server attempts to create during resizing of the maintained connection pool.
SHRINKENABLED	Specifies whether unused ManagedConnections will be destroyed and removed from the connection pool as a means to control system resources
SHRINKFREQ	Specifies the amount of time (in seconds) the Connection Pool Management waits between attempts to destroy unused ManagedConnections
HIGHNOWAITER	Specifies the maximum number of threads that can concurrently block waiting to reserve a connection from the pool
HIGHNOUNAVAILABLE	Specifies the maximum number of ManagedConnections in the pool that can be made unavailable to the application for purposes such as refreshing the connection
CONNCREATIONRETRYREQ	The number of seconds between when the connection pool retries to establish connections to the database
CONNRESERVETIMEOUT	The number of seconds after which a call to reserve a connection from the connection pool will timeout
TESTFREQUENCY	The number of seconds between when WebLogic Server tests unused database connections
TESTCONNONCREATE	Specifies whether WebLogic Server tests a connection after creating it but before adding it to the list of connections available in the pool
TESTCONNONRELEASE	Specifies whether WebLogic Server tests a connection before returning it to this JDBC connection pool

Table 19–165 (Cont.) MGMT\$WEBLOGIC_RAOUTBOUNDCONFIG

Column	Description
TESTCONNONRESERVE	Specifies whether WebLogic Server tests a connection before giving it to the client
PROFILEHARVESTFREQ	Specifies how frequently the profile for the connection pool is being harvested
IGNOREINUSECONNENABLED	When the connection pool is being shut down, this element is used to specify whether it is acceptable to ignore connections that are in use at that time
MATCHCONNSUPPORTED	Indicates whether the resource adapter supports the <code>ManagedConnectionFactory.matchManagedConnections()</code> method. If the resource adapter does not support this method (always returns null for this method), then WebLogic Server bypasses this method call during a connection request.

19.22.17 MGMT\$WEBLOGIC_RESOURCECONFIG

The MGMT\$WEBLOGIC_RESOURCECONFIG view displays general information about the Resource configuration.

Table 19–166 MGMT\$WEBLOGIC_RESOURCECONFIG

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target
CM_TARGET_TYPE	The type of target: weblogic_j2eeserver
CM_TARGET_NAME	The name of the target in Enterprise Manager
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
STARTHEAP	Start heap value
MAXHEAP	Maximum heap value

19.22.18 MGMT\$WEBLOGIC_SERVER

The MGMT\$WEBLOGIC_SERVER view displays the information about the various ports of Oracle WebLogic Server.

Table 19–167 MGMT\$WEBLOGIC_SERVER

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target
CM_TARGET_TYPE	The type of target: weblogic_j2eeserver
CM_TARGET_NAME	The name of target in Enterprise Manager
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
LISTENPORT	The default TCP port that this server uses to listen for regular (non-SSL) incoming connections
ADMINISTRATIONPORT	The common secure administration port for this WebLogic Server domain

Table 19–167 (Cont.) MGMT\$WEBLOGIC_SERVER

Column	Description
NATIVEIOENABLED	Specifies whether native input or output is enabled for the server
MAXOPENSOCKETCOUNT	The maximum number of open sockets allowed in the server at a given point of time
STUCKTHREADMAXTIME	The number of seconds that a thread must be continually working before this server considers the thread stuck
STUCKTHREADTIMERINTERVAL	The number of seconds after which WebLogic Server periodically scans threads to see if they have been continually working for the configured maximum length of time
ACCEPTBACKLOG	The number of backlogged, new TCP connection requests that should be allowed for this server's regular and SSL ports
LOGINTIMEOUT	The login timeout for this server's default regular (non-SSL) listen port. This is the maximum amount of time allowed for a new connection to establish.
MANAGEDSERVERINDENABLED	Specifies whether this Managed Server can be started when the Administration Server is unavailable.
LOWMEMGCTHRESHOLD	The threshold level (in percent) that this server uses for logging low memory conditions and changing the server health state to Warning
LOWMEMGRANULARITYLEVEL	The granularity level (in percent) that this server uses for logging low memory conditions and changing the server health state to Warning
LOWMEMORYSAMPLESIZE	The number of times this server samples free memory during the time period specified by LowMemoryTimeInterval
LOWMEMTIMEINTERVAL	The amount of time (in seconds) that defines the interval over which this server determines average free memory values
SSLLISTENPORT	The TCP/IP port at which this server listens for SSL connection requests
SSLLOGINTIMEOUT	SSL Login time out
CLUSTERNAME	The cluster, or group of WebLogic Server instances, to which this server belongs
CLUSTERWEIGHT	The proportion of the load that this server will bear, relative to other servers in a cluster
JAVAVMVENDOR	Java Virtual Machine (VM) vendor
JAVAVERSION	Java version installed on this server
MACHINENAME	Name of the system where this server is installed
DOMAINHOME	Path of the WebLogic domain which contains the WebLogic server target
MAXPOSTSIZE	Maximum post size
JSSE_ENABLED	Java Secure Socket Extension (JSSE) enabled
SCATTERED_READS_ENABLED	Specifies whether scattered reads over NIO Socket channels is enabled
GATHERED_WRITES_ENABLED	Specifies whether gathered writes over NIO socket channels. is enabled
REPLICATION_PORTS	Listen ports used by replication channels when the WebLogic server is running on Exalogic systems

Table 19–167 (Cont.) MGMT\$WEBLOGIC_SERVER

Column	Description
BINARY_HOST	The host on which the WebLogic is installed in cases where WebLogic binaries installed on one host are used through mounts in WebLogic instances running on other hosts
BINARY_WEBLOGICHOME	The WebLogic home where the WebLogic binaries are installed on the binary host
LISTENADDRESS	The listen address on which the server is listening to on a non-secure port as configured
SSLLISTENADDRESS	The listen address on which the server is listening to on a secure port as configured

19.22.19 MGMT\$WEBLOGIC_STARTSHUTCLASSES

The MGMT\$WEBLOGIC_STARTSHUTCLASSES view displays general information about the Startup and Shutdown classes.

Table 19–168 MGMT\$WEBLOGIC_STARTSHUTCLASSES

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target
CM_TARGET_TYPE	The type of target: weblogic_j2eeserver
CM_TARGET_NAME	The name of the target in Enterprise Manager
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
NAME	Name of the startup or shutdown class
TYPE	Type of class – startup or shutdown
CLASSNAME	The fully qualified name of a class to load and run. The class must be on the server's class path.
DEPLOYMENTORDER	A priority that the server uses to determine when it deploys an item. The priority is relative to other deployable items of the same type.
ARGUMENTS	Arguments that a server uses to initialize a class

19.22.20 MGMT\$WEBLOGIC_VIRTUALHOST

The MGMT\$WEBLOGIC_VIRTUALHOST view displays general information about the Virtual Hosts configuration.

Table 19–169 MGMT\$WEBLOGIC_VIRTUALHOST

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target
CM_TARGET_TYPE	The type of target: weblogic_j2eeserver
CM_TARGET_NAME	The name of the target in Enterprise Manager
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot

Table 19–169 (Cont.) MGMT\$WEBLOGIC_VIRTUALHOST

Column	Description
NAME	Name of virtual host
DEPLOYMENTORDER	A priority that the server uses to determine when it deploys an item. The priority is relative to other deployable items of the same type.
FRONTENDHOST	The name of the host to which all redirected URLs will be sent. If specified, WebLogic Server will use this value rather than the one in the HOST header.
FRONTENDHTTPPORT	The name of the HTTP port to which all redirected URLs will be sent. If specified, WebLogic Server will use this value rather than the one in the HOST header.
FRONTENDHTTPSPORT	The name of the secure HTTP port to which all redirected URLs will be sent. If specified, WebLogic Server will use this value rather than the one in the HOST header.
VIRTUALHOSTNAMES	The list of host names, separated by line breaks, for which this virtual host will serve requests
NETWORKACCESSPOINT	The dedicated server channel name (NetworkAccessPoint) for which this virtual host will serve HTTP request
LOGFILENAME	The name of the file that stores HTTP requests. If the path name is not absolute, the path is assumed to be relative to the root directory of the system on which this server is running.
LOGGINGENABLED	Indicates whether this server logs HTTP requests
MAXPOSTSIZE	The maximum post size this server allows for reading HTTP POST data in a servlet request. A value less than 0 indicates an unlimited size.

19.22.21 MGMT\$WEBLOGIC_WEBAPPCOMPONENT

The MGMT\$WEBLOGIC_WEBAPPCOMPONENT displays general information about the web modules.

Table 19–170 MGMT\$WEBLOGIC_WEBAPPCOMPONENT

Column	Description
NAME	Name of the web module
APPLICATION	Name of the application that includes the component
DEPLOYMENTORDER	Priority that the server uses when it deploys an item. The priority is relative to the other deployable items of same type
CONTEXTPATH	Context path of the web module
SESSIONTIMEOUTSECS	Session timeout in seconds

19.22.22 MGMT\$WEBLOGIC_WORKMANAGER

The MGMT\$WEBLOGIC_WORKMANAGER view displays general information about the Work Manager configuration.

Table 19–171 MGMT\$WEBLOGIC_WORKMANAGER

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target

Table 19–171 (Cont.) MGMT\$WEBLOGIC_WORKMANAGER

Column	Description
CM_TARGET_TYPE	The type of target: weblogic_j2eeserver
CM_TARGET_NAME	The name of the target in Enterprise Manager
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
NAME	Name of the work manager
IGNORESTUCKTHREADS	Specifies whether this Work Manager ignores “stuck” threads
MINIMUMTHREADS	Minimum number of concurrent threads executing requests that share this constraint
MAXIMUMTHREADS	Maximum number of concurrent threads that can execute requests sharing this constraint
REQUESTCLASSTYPE	Type of request class
REQUESTCLASSNAME	Request class name
CAPACITYCONSTRAINT	Total number of requests that can be enqueued

19.22.23 MGMT\$WEBLOGIC_WSCONFIG

The MGMT\$WEBLOGIC_WSCONFIG view displays general information about the web service configuration.

Table 19–172 MGMT\$WEBLOGIC_WSCONFIG

Column	Description
NAME	The name of the web service configuration
SERVICENAME	Name of the web service. Corresponds to the name attribute of the service element in the WSDL that describes the web service
APPNAME	The name of the application
IMPLEMENTATIONTYPE	Implementation type of the service Possible values: <ul style="list-style-type: none"> ■ JAX-WS 2.0 ■ JAX-RPC 1.1
URI	URI of this web service. The value corresponds to the final part of the endpoint address in the WSDL that describes the web services

19.22.24 MGMT\$WEBLOGIC_WSPORTCONFIG

The MGMT\$WEBLOGIC_WSPORTCONFIG view displays general information about the web services port configuration.

Table 19–173 MGMT\$WEBLOGIC_WSPORTCONFIG

Column	Description
NAME	Name of the web services port configuration
SERVICENAME	Name of the web service. Corresponds to the name attribute of the service element in the WSDL that describes the web service

Table 19–173 (Cont.) MGMT\$WEBLOGIC_WSPORTCONFIG

Column	Description
APPNAME	The name of the application
TRANSPORTPROTOCOL	Transport protocol used to invoke this web service, such as HTTP, HTTPS, or JMS

19.23 Oracle WebLogic Domain Views

This section provides a description of each Oracle WebLogic domain view and its columns

19.23.1 MGMT\$WEBLOGIC_DOMAIN

The MGMT\$WEBLOGIC_DOMAIN view displays general information about the WebLogic Domain configuration.

Table 19–174 MGMT\$WEBLOGIC_DOMAIN

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target
CM_TARGET_TYPE	The type of target: weblogic_domain
CM_TARGET_NAME	The name of the target in Enterprise Manager
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
ADMINISTRATIONPORTENABLED	WebLogic administration port
PRODUCTIONMODEENABLED	WebLogic server production mode status
EXALOGIC_OPT_ENABLED	Exalogic optimizations enabled status
NAME	Name of the WebLogic domain
DOMAINVERSION	Version of the WebLogic domain

19.23.2 MGMT\$WEBLOGIC_OPSSSYSPROP

The MGMT\$WEBLOGIC_OPSSSYSPROP view displays general information about the Oracle Platform Security Services (OPSS) System Properties.

Table 19–175 MGMT\$WEBLOGIC_OPSSSYSPROP

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target
CM_TARGET_TYPE	The type of target: weblogic_domain
CM_TARGET_NAME	The name of target in Enterprise Manager
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot

Table 19–175 (Cont.) MGMT\$WEBLOGIC_OPSSSYSPROP

Column	Description
COMBINER_LAZYEVAL	Enables or disables the evaluation of a subject protection domain when check permission is triggered. Default value: FALSE
COMBINER_OPTIMIZE	Enables or disables the caching of a subject protection domain. Default value: FALSE
AUTHORIZATION	Enables or disables the delegation of calls to JDK API AccessController.checkPermission, which reduces runtime and debugging overhead. Default value: FALSE
HYBRID_MODE	Enables or disables the hybrid mode. The hybrid mode is used to facilitate the transition from the Sun <code>java.security.Policy</code> to the OPSS Java PolicyProvider. Default value: TRUE

19.23.3 MGMT\$WEBLOGIC_OAMCONFIG

Each row in the MGMT\$WEBLOGIC_OAMCONFIG view represents configuration data of Oracle Access Manager (OAM) configured for the WebLogic domain.

Table 19–176 MGMT\$WEBLOGIC_OAMCONFIG

Column	Description
ECM_SNAPSHOT_ID	GUID of the snapshot
PORT	Provides the value of the port where OAM is deployed
HOSTNAME	Provides the name of the host where OAM is deployed

19.24 Oracle WebLogic Cluster Views

This section provides a description of each Oracle WebLogic cluster view and its columns.

19.24.1 MGMT\$WEBLOGIC_CLUSTER

The MGMT\$WEBLOGIC_CLUSTER view displays general information about the WebLogic Cluster configuration.

Table 19–177 MGMT\$WEBLOGIC_CLUSTER

Column	Description
CM_TARGET_GUID	The unique global identifier (GUID) for the target
CM_TARGET_TYPE	The type of target: <code>weblogic_cluster</code>
CM_TARGET_NAME	The name of target in Enterprise Manager
LAST_COLLECTION_TIMESTAMP	The date and time when the metrics were collected
ECM_SNAPSHOT_ID	GUID of the snapshot
SESS_LAZY_DESER_ENABLED	Session lazy deserialization enabled
CLUSTER_ADDRESS	Cluster address
CLUSTER_BROADCAST_CHANNEL	Cluster broadcast channel
DEFAULT_LOAD_ALGO	Default load algorithm

Table 19–177 (Cont.) MGMT\$WEBLOGIC_CLUSTER

Column	Description
CLUSTER_MESSAGING_MODE	Cluster messaging mode
CLUSTER_TYPE	Cluster type
REPLICATION_CHANNEL	Replication channel

19.25 Security Views

This section provides a description of each security view and its columns.

19.25.1 MGMT\$ESA_ALL_PRIVS_REPORT

The MGMT\$ESA_ALL_PRIVS_REPORT view displays a table containing users and roles that have the 'GRANT ANY PRIVILEGE' privilege in database security reports.

Table 19–178 MGMT\$ESA_ALL_PRIVS_REPORT

Column	Description
TARGET_GUID	The GUID of the target for which the report has the data
TARGET_NAME	The name of the target for which the report has the data
PRINCIPAL	The user or roles that have been granted this privilege (that is, GRANT ANY PRIVILEGE->DBA->SYS)
OBJECT_NAME	The name of the user that been granted the privilege (GRANT ANY PRIVILEGE)

19.25.2 MGMT\$ESA_ANY_DICT_REPORT

The MGMT\$ESA_ANY_DICT_REPORT view displays a table and a chart containing users and roles with access to any dictionary in database security reports.

Table 19–179 MGMT\$ESA_ANY_DICT_REPORT

Column	Description
TARGET_GUID	The GUID of the target for which the report has the data
TARGET_NAME	The name of the target for which the report has the data
PRINCIPAL	The user or roles that been granted this privilege. For example, SELECT ANY DICTIONARY->SCHEMA_OWNER_ROLE->SYS
OBJECT_NAME	The user that has been granted any of the ANY DICTIONARY privileges. For example, SELECT ANY DICTIONARY, ANALYZE ANY DICTIONARY, and so on.

19.25.3 MGMT\$ESA_ANY_PRIV_REPORT

The MGMT\$ESA_ANY_PRIV_REPORT view displays a table and a chart containing users with 'ANY' in some privilege granted to them in database security reports.

Table 19–180 MGMT\$ESA_ANY_PRIV_REPORT

COLUMN	Description
TARGET_GUID	The GUID of the target for which the report has the data

Table 19–180 (Cont.) MGMT\$ESA_ANY_PRIV_REPORT

COLUMN	Description
TARGET_NAME	The name of the target for which the report has the data
PRINCIPAL	The user or role which has been granted a privilege having 'ANY'. For example, BACKUP ANY TABLE->EXP_FULL_DATABASE->DATAPUMP_EXP_FULL_DATABASE->SYS
OBJECT_NAME	The user that has been granted one of the ANY privileges. For example, ALTER ANY MATERIALIZED VIEW, ALTER ANY INDEX, BACKUP ANY TABLE, and so on.

19.25.4 MGMT\$ESA_AUDIT_SYSTEM_REPORT

The MGMT\$ESA_AUDIT_SYSTEM_REPORT view displays a table containing users and roles with the 'AUDIT SYSTEM' privilege in database security reports.

Table 19–181 MGMT\$ESA_AUDIT_SYSTEM_REPORT

Column	Description
TARGET_GUID	The GUID of the target for which the report has the data
TARGET_NAME	The name of the target for which the report has the data
PRINCIPAL	The user or role which has been granted 'AUDIT SYSTEM' privilege. For example, AUDIT SYSTEM->SYS, AUDIT SYSTEM->IMP_FULL_DATABASE->DATAPUMP_IMP_FULL_DATABASE->DBA->SYSTEM, and so on.
OBJECT_NAME	The user that has been granted 'ALTER SYSTEM' privilege

19.25.5 MGMT\$ESA_BECOME_USER_REPORT

The MGMT\$ESA_BECOME_USER_REPORT view displays a table containing users and roles with the 'BECOME USER' privilege in database security reports.

Table 19–182 MGMT\$ESA_BECOME_USER_REPORT

Column	Description
TARGET_GUID	The GUID of the target for which the report has the data
TARGET_NAME	The name of the target for which the report has the data
PRINCIPAL	The user or role which has been granted 'BECOME USER' privilege. For example, BECOME USER->SYS, BECOME USER->DBA->SYSTEM, BECOME USER->IMP_FULL_DATABASE->DATAPUMP_IMP_FULL_DATABASE->DBA->BAM, and so on.
OBJECT_NAME	The user that has been granted the 'BECOME USER' privilege

19.25.6 MGMT\$ESA_CATALOG_REPORT

The MGMT\$ESA_CATALOG_REPORT view displays a table and a chart containing all the users that have a role such as '%CATALOG%' in database security reports.

Table 19–183 MGMT\$ESA_CATALOG_REPORT

Column	Description
TARGET_GUID	The GUID of the target for which the report has the data
TARGET_NAME	The name of the target for which the report has the data

Table 19–183 (Cont.) MGMT\$ESA_CATALOG_REPORT

Column	Description
PRINCIPAL	The user or role which has been granted a role like '%CATALOG%'. For example, RECOVERY_CATALOG_OWNER->SYS, EXECUTE_CATALOG_ROLE->TBLO_ROLE->CRM, and so on.
OBJECT_NAME	User that has been granted one of the 'CATALOG' privileges. For example, SELECT_CATALOG_ROLE, EXECUTE_CATALOG_ROLE, DELETE_CATALOG_ROLE, and so on.

19.25.7 MGMT\$ESA_CONN_PRIV_REPORT

The MGMT\$ESA_CONN_PRIV_REPORT view displays a table and a chart containing users and roles with the CONNECT or RESOURCE role in database security reports.

Table 19–184 MGMT\$ESA_CONN_PRIV_REPORT

Column	Description
TARGET_GUID	The GUID of the target for which the report has the data
TARGET_NAME	The name of the target for which the report has the data
PRINCIPAL	The user or role which has been granted the CONNECT or RESOURCE role
OBJECT_NAME	The role if granted directly, or the role through it has been granted

19.25.8 MGMT\$ESA_CREATE_PRIV_REPORT

The MGMT\$ESA_CREATE_PRIV_REPORT view displays a table and a chart containing users and roles with the CREATE privilege in database security reports.

Table 19–185 MGMT\$ESA_CREATE_PRIV_REPORT

Column	Description
TARGET_GUID	The GUID of the target for which the report has the data
TARGET_NAME	The name of the target for which the report has the data
PRINCIPAL	The user or role which has been granted the privilege to create an object in the database. For example, CREATE ANY CONTEXT->SYS, CREATE ANY INDEX->OLAP_DBA->OLAPSYS, and so on.
OBJECT_NAME	User that has been granted one of the 'CREATE' privileges

19.25.9 MGMT\$ESA_DBA_GROUP_REPORT

The MGMT\$ESA_DBA_GROUP_REPORT view displays a table containing members of the operating system user group DBA in database security reports.

Table 19–186 MGMT\$ESA_DBA_GROUP_REPORT

Column	Description
TARGET_GUID	The GUID of the target for which the report has the data
TARGET_NAME	The name of the target for which the report has the data
PRINCIPAL	The operating system user which is in the user group DBA.

Table 19–186 (Cont.) MGMT\$ESA_DBA_GROUP_REPORT

Column	Description
OBJECT_NAME	DBA Group

19.25.10 MGMT\$ESA_DBA_ROLE_REPORT

The MGMT\$ESA_DBA_ROLE_REPORT view displays a table containing users and roles with the DBA role granted to them in database security reports.

Table 19–187 MGMT\$ESA_DBA_ROLE_REPORT

Column	Description
TARGET_GUID	The GUID of the target for which the report has the data
TARGET_NAME	The name of the target for which the report has the data
PRINCIPAL	The user or role which has been granted the DBA role
OBJECT_NAME	User that has been granted the DBA role

19.25.11 MGMT\$ESA_DIRECT_PRIV_REPORT

The MGMT\$ESA_DIRECT_PRIV_REPORT view displays a table and a chart containing privileges granted directly in database security reports.

Table 19–188 MGMT\$ESA_DIRECT_PRIV_REPORT

Column	Description
TARGET_GUID	The GUID of the target for which the report has the data
TARGET_NAME	The name of the target for which the report has the data
PRINCIPAL	User which has been granted a privilege directly i.e. not via a role
OBJECT_NAME	The privilege that has been granted directly. For example, ALTER SESSION, SELECT ANY DICTIONARY, and so on.

19.25.12 MGMT\$ESA_EXMPT_ACCESS_REPORT

The MGMT\$ESA_EXMPT_ACCESS_REPORT view displays a table containing users and roles with the EXEMPT ACCESS POLICY privilege in database security reports.

Table 19–189 MGMT\$ESA_EXMPT_ACCESS_REPORT

Column	Description
TARGET_GUID	The GUID of the target for which the report has the data
TARGET_NAME	The name of the target for which the report has the data
PRINCIPAL	The user or role which has been granted the 'EXEMPT ACCESS POLICY' privilege
OBJECT_NAME	User that has been granted one of the 'EXEMPT ACCESS POLICY' privilege

19.25.13 MGMT\$ESA_KEY_OBJECTS_REPORT

The MGMT\$ESA_KEY_OBJECTS_REPORT view displays a table and a chart containing users and roles with access to key objects in database security reports.

Table 19–190 MGMT\$ESA_KEY_OBJECTS_REPORT

Column	Description
TARGET_GUID	The GUID of the target for which the report has the data
TARGET_NAME	The name of the target for which the report has the data
USER	The user which has access to key objects
OBJECT_NAME	The key object to which that user has access. For example, View DBA_USERS, Table SOURCE\$, Table USER\$
PRIVILEGE	The privilege on the key object that has been granted to the user. For example, SELECT, DELETE, and so on.

19.25.14 MGMT\$ESA_OH_OWNERSHIP_REPORT

The MGMT\$ESA_OH_OWNERSHIP_REPORT view displays a table containing file ownership by Oracle home in database security reports.

Table 19–191 MGMT\$ESA_OH_OWNERSHIP_REPORT

Column	Description
TARGET_GUID	The GUID of the target for which the report has the data
TARGET_NAME	The name of the target for which the report has the data
PRINCIPAL	The file whose owner is not the ORACLE HOME owner
OBJECT_NAME	The owner of the file

19.25.15 MGMT\$ESA_OH_PERMISSION_REPORT

The MGMT\$ESA_OH_PERMISSION_REPORT view displays a table containing file permissions by Oracle home in database security reports.

Table 19–192 MGMT\$ESA_OH_PERMISSION_REPORT

Column	Description
TARGET_GUID	The GUID of the target for which the report has the data
TARGET_NAME	The name of the target for which the report has the data
PRINCIPAL	The file that has an insecure permission
OBJECT_NAME	The permission of the file

19.25.16 MGMT\$ESA_POWER_PRIV_REPORT

The MGMT\$ESA_POWER_PRIV_REPORT view displays a table and a chart containing all the users and roles with ALTER SESSION, ALTER SYSTEM, CREATE PROCEDURE or CREATE LIBRARY privileges in database security reports.

Table 19–193 MGMT\$ESA_POWER_PRIV_REPORT

Column	Description
TARGET_GUID	The GUID of the target for which the report has the data
TARGET_NAME	The name of the target for which the report has the data
PRINCIPAL	The user that has powerful privileges
OBJECT_NAME	The powerful privilege held by the user

19.25.17 MGMT\$ESA_PUB_PRIV_REPORT

The MGMT\$ESA_PUB_PRIV_REPORT view displays a table and a chart containing privileges granted to PUBLIC in database security reports.

Table 19–194 MGMT\$ESA_PUB_PRIV_REPORT

Column	Description
TARGET_GUID	The GUID of the target for which the report has the data
TARGET_NAME	The name of the target for which the report has the data
PRINCIPAL	The object on which some privilege has been granted to PUBLIC
OBJECT_NAME	The privilege on the object which has been granted to PUBLIC. For example, SELECT, EXECUTE, and so on.

19.25.18 MGMT\$ESA_SYS_PUB_PKG_REPORT

The MGMT\$ESA_SYS_PUB_PKG_REPORT view displays a table containing system packages with public execute privileges in database security reports.

Table 19–195 MGMT\$ESA_SYS_PUB_PKG_REPORT

Column	Description
TARGET_GUID	The GUID of the target for which the report has the data
TARGET_NAME	The name of the target for which the report has the data
PRINCIPAL	PUBLIC
OBJECT_NAME	The package owned by SYS on which PUBLIC has execute privileges

19.25.19 MGMT\$ESA_TABSP_OWNERS_REPORT

The MGMT\$ESA_TABSP_OWNERS_REPORT view displays a table containing tablespaces and their owners in database security reports.

Table 19–196 MGMT\$ESA_TABSP_OWNERS_REPORT

Column	Description
TARGET_GUID	The GUID of the target for which the report has the data
TARGET_NAME	The name of the target for which the report has the data
PRINCIPAL	The tablespace
OBJECT_NAME	The owner of the tablespace

19.25.20 MGMT\$ESA_TRC_AUD_PERM_REPORT

The MGMT\$ESA_TRC_AUD_PERM_REPORT view displays a table containing trace and audit files permissions in database security reports.

Table 19–197 MGMT\$ESA_TRC_AUD_PERM_REPORT

Column	Description
TARGET_GUID	The GUID of the target for which the report has the data
TARGET_NAME	The name of the target for which the report has the data

Table 19–197 (Cont.) MGMT\$ESA_TRC_AUD_PERM_REPORT

Column	Description
PRINCIPAL	The file path
OBJECT_NAME	The purpose of the file. For example, audit file destination, background dump destination, core dump destination, user dump destination, and so on.
PERMISSION	Permission of the file

19.25.21 MGMT\$ESA_WITH_ADMIN_REPORT

The MGMT\$ESA_WITH_ADMIN_REPORT view displays a table and a chart containing users and roles having some privileges granted to them with the WITH ADMIN option in database security reports.

Table 19–198 MGMT\$ESA_WITH_ADMIN_REPORT

Column	Description
TARGET_GUID	The GUID of the target for which the report has the data
TARGET_NAME	The name of the target for which the report has the data
PRINCIPAL	The user or role which has been granted a privilege 'WITH ADMIN' option
OBJECT_NAME	The privilege which has been granted 'WITH ADMIN' option

19.25.22 MGMT\$ESA_WITH_GRANT_REPORT

The MGMT\$ESA_WITH_GRANT_REPORT view displays a table and a chart containing users and roles having some privileges granted to them with 'WITH GRANT' option in database security reports.

Table 19–199 MGMT\$ESA_WITH_GRANT_REPORT

Column	Description
TARGET_GUID	The GUID of the target for which the report has the data
TARGET_NAME	The name of the target for which the report has the data
PRINCIPAL	The user or role which has been granted a privilege 'WITH GRANT' option
OBJECT_NAME	The privilege which has been granted 'WITH GRANT' option

19.25.23 MGMT\$ESM_COLLECTION_LATEST

The MGMT\$ESM_COLLECTION_LATEST view contains properties relating to security for database targets.

Table 19–200 MGMT\$ESM_COLLECTION_LATEST

Column	Description
TARGET_GUID	The GUID of the database target
PROPERTY	Name of the attribute
VALUE	Value of the attribute
VALUE2	Used to capture additional values of the attribute

19.25.24 MGMT\$ESM_FILE_SYSTEM_LATEST

The MGMT\$ESM_FILE_SYSTEM_LATEST view contains the file system type for the Windows host targets.

Table 19–201 MGMT\$ESM_FILE_SYSTEM_LATEST

Column	Description
TARGET_GUID	The GUID of the Windows host target
FILE_SYSTEM	The type of file system

19.25.25 MGMT\$ESM_PORTS_LATEST

The MGMT\$ESM_PORTS_LATEST view contains the open ports for the host target.

Table 19–202 MGMT\$ESM_PORTS_LATEST

Column	Description
TARGET_GUID	The GUID of the host target
PORT	The value of the open port (listening mode)

19.25.26 MGMT\$ESM_SERVICE_LATEST

The MGMT\$ESM_SERVICE_LATEST view contains the insecure services running on the host targets.

Table 19–203 MGMT\$ESM_SERVICE_LATEST

Column	Description
TARGET_GUID	The GUID of the host target
SERVICE	The port value for the service

19.25.27 MGMT\$ESM_STACK_LATEST

The MGMT\$ESM_STACK_LATEST view contains executable stack status host targets.

Table 19–204 MGMT\$ESM_STACK_LATEST

Column	Description
TARGET_GUID	The GUID of the host target
EXE_STACK	The status of the executable stack

19.26 Service Tag Views

This section provides a description of each service tag view and its columns.

19.26.1 MGMT\$SERVICETAG_INSTANCES

The MGMT\$SERVICETAG_INSTANCES view provides information about the product ID and instance of the product installed on the server, such as Oracle Solaris Cluster.

Table 19–205 MGMT\$SERVICETAG_INSTANCES

Column	Description
PRODUCT_URN	Denotes the product ID for the product installed on the server

Table 19–205 (Cont.) MGMT\$SERVICETAG_INSTANCES

Column	Description
INSTANCE_URN	Uniquely identifies the instance of the product installed on that server

19.26.2 MGMT\$SERVICETAG_REGISTRY

The MGMT\$SERVICETAG_REGISTRY view uniquely identifies the instance of a Service Tag registry on a server and the version of Service Tag software that created it.

Table 19–206 MGMT\$SERVICETAG_REGISTRY

Column	Description
AGENT_ID	Uniquely identifies the instance of a Service Tag registry on the server
AGENT_VERSION	Uniquely identifies the version of the Service Tag software
REGISTRY_VERSION	Identifies the version of the registry. This is the registry that Oracle Configuration Manager (OCM) harvests to collect the data for this metric.

19.27 Storage Reporting Views

This section provides a description of each storage reporting view and its columns.

19.27.1 MGMT\$STORAGE_REPORT_DATA

The MGMT\$STORAGE_REPORT_DATA view displays the Storage Data metric attributes which are common across all instrumented Storage Entities.

Table 19–207 MGMT\$STORAGE_REPORT_DATA

Column	Description
TARGET_NAME	Target Name in Enterprise Manager
TARGET_TYPE	Target Type in Enterprise Manager
KEY_VALUE	Unique Key Value for the Storage Entity
GLOBAL_UNIQUE_ID	A globally unique persistent identifier for a storage entity. All instances of a shared Storage Entity will have the same global_unique_identifier
NAME	Name of the storage entity
STORAGE_LAYER	Storage layer of the storage entity. Sample Usage: - OS_DISK - VOLUME_MANAGER - LOCAL_FILESYSTEM - NFS
ENTITY_TYPE	Indicates the type of Entity. Value is vendor specific. Example: Plex, Sub Disk, Diskgroup, Volume group, Metadevice and so on.
RAWSIZEB	Total space in bytes
SIZEB	Size in bytes

Table 19-207 (Cont.) MGMT\$STORAGE_REPORT_DATA

Column	Description
USEDDB	Used size in bytes
FREEB	Free size in bytes

19.27.2 MGMT\$STORAGE_REPORT_KEYS

The MGMT\$STORAGE_REPORT_KEYS view displays the relationship between instrumented Storage Entities.

Table 19-208 MGMT\$STORAGE_REPORT_KEYS

Column	Description
TARGET_NAME	Target Name in Enterprise Manager
TARGET_TYPE	Target type in Enterprise Manager
KEY_VALUE	Unique KEY_VALUE for the storage entity
PARENT_KEY_VALUE	Key value for the parent storage entity.

19.27.3 MGMT\$STORAGE_REPORT_PATHS

The MGMT\$STORAGE_REPORT_PATHS view displays the OS paths for all instrumented storage entities.

Table 19-209 MGMT\$STORAGE_REPORT_PATHS

Column	Description
TARGET_NAME	Target name in Enterprise Manager
TARGET_TYPE	Target type in Enterprise Manager
KEY_VALUE	Unique key value for the storage entity
NAME	Name of the storage entity
PATH	OS path to the storage entity
FILE_TYPE	Type of file Examples: - _BLOCKSPECIAL - _DIRECTORY - _REGULAR
STORAGE_LAYER	Storage layer of the storage entity. Sample Usage: - OS_DISK - VOLUME_MANAGER - LOCAL_FILESYSTEM - NFS
ENTITY_TYPE	Indicates the type of entity. Value is vendor-specific. Examples: Plex, Sub Disk, Diskgroup, Volume group, Metadevice, and so on.

19.27.4 MGMT\$STORAGE_REPORT_ISSUES

The MGMT\$STORAGE_REPORT_ISSUES view displays the consistency issues encountered when analyzing the instrumented storage metrics.

Table 19–210 MGMT\$STORAGE_REPORT_ISSUES

Column	Description
TARGET_NAME	Target name in Enterprise Manager
TARGET_TYPE	Target type in Enterprise Manager
TYPE	Type of inconsistency. Value can be: - MAPPING_ISSUE - MAPPING_WARNING
MESSAGE_COUNT	Count of the number of messages

19.27.5 MGMT\$STORAGE_REPORT_DISK

The MGMT\$STORAGE_REPORT_DISK view displays Additional Storage Data Metric Attributes for all physical disk device storage entities.

Table 19–211 MGMT\$STORAGE_REPORT_DISK

Column	Description
TARGET_NAME	Target name in Enterprise Manager
TARGET_TYPE	Target type in Enterprise Manager
ENTITY_TYPE	Indicates the type of disk device such as disk or disk partition
USED_PATH	The OS path to the disk or partition. If the disk or partition is allocated, then this is the path that is in use.
FILE_TYPE	Type of file Examples: - _BLOCKSPECIAL - _REGULAR
SIZEB	Size in bytes
USEDDB	Used size in bytes
FREEB	Free size in bytes
VENDOR	Name of the disk vendor; detected through SCSI enquiry
PRODUCT	Product family from the vendor; detected through SCSI enquiry

19.27.6 MGMT\$STORAGE_REPORT_VOLUME

The MGMT\$STORAGE_REPORT_VOLUME view displays Additional Storage Data Metric attributes for all volume manager storage entities.

Table 19–212 MGMT\$STORAGE_REPORT_VOLUME

Column	Description
TARGET_NAME	Target name in Enterprise Manager
TARGET_TYPE	Target type in Enterprise Manager
VENDOR	Vendor name of the volume or software raid manager

Table 19-212 (Cont.) MGMT\$STORAGE_REPORT_VOLUME

Column	Description
PRODUCT	Vendor name of the volume or software raid manager
TYPE	Indicates the type of volume entity. It can be vendor specific. For example, in the case of Veritas Volume Manager for Volume, Plex, VM Disk, Diskgroup, Sub Disk, Metadevice, Metadevice Partition, Array, Raiddevice, Submirror, Diskset, Slice, raid-disk, spare-disk, Hot spare, and so on.
DISK_GROUP	Disk group or volume group name
NAME	The name of the entity in the volume manager namespace
USED_PATH	The OS path to the device. If the device is allocated, then this is the path that is in use.
FILE_TYPE	Type of file Examples: - _BLOCKSPECIAL - _REGULAR
RAWSIZEB	In bytes. For a 2-way mirrored Veritas Volume. It is the sum of the size of each plex.
SIZEB	Size in bytes
USEDDB	Used size in bytes
FREEDB	Free size in bytes
CONFIGURATION	A string describing the configuration of the volume.

19.27.7 MGMT\$STORAGE_REPORT_LOCALFS

The MGMT\$STORAGE_REPORT_LOCALFS view displays Additional Storage Data Metric attributes for all local file system storage entities.

Table 19-213 MGMT\$STORAGE_REPORT_LOCALFS

Column	Description
TARGET_NAME	Target name in Enterprise Manager
TARGET_TYPE	Target type in Enterprise Manager
FILESYSTEM_TYPE	The type of file system
FILESYSTEM	The file system path on the operating system
MOUNTPPOINT	The mount point path on the operating system
SIZEB	Number
USEDDB	Used size in bytes
FREEDB	Free size in bytes

19.27.8 MGMT\$STORAGE_REPORT_NFS

The MGMT\$STORAGE_REPORT_NFS view displays Additional Storage Data Metric attributes for all network file systems.

Table 19–214 MGMT\$STORAGE_REPORT_NFS

Column	Description
TARGET_NAME	Target name in Enterprise Manager
TARGET_TYPE	Target type in Enterprise Manager
FILESYSTEM	The file system name as seen on the operating system. For NFS file systems the file system name should be in the format nfs_server:/filesystem_name
MOUNTPPOINT	The mountpoint path on the operating system
SIZEB	Size in bytes
USEDDB	Used size in bytes,
FREEB	Free size in bytes
NFS_SERVER	The server name for the NFS Server
NFS_SERVER_IP_ADDRESS	The IP address of the NFS Server
NFS_VENDOR	The NFS Server vendor
MOUNT_PRIVILEGE	This is the mount privilege of the file system Possible values: <ul style="list-style-type: none"> ■ Read ■ Write

19.28 Target Views

This section provides a description of each target view and its columns.

19.28.1 MGMT\$AGENTS_MONITORING_TARGETS

The MGMT\$AGENTS_MONITORING_TARGETS view shows the available Management Agents for targets.

Table 19–215 MGMT\$AGENTS_MONITORING_TARGETS

Column	Description
TARGET_NAME	Name of the target
TARGET_TYPE	Type of the target
TARGET_GUID	Unique ID of the target
AGENT_NAME	Name of the Management Agent monitoring the target
AGENT_TYPE	Type of Management Agent
AGENT_GUID	Unique ID of the Management Agent
AGENT_IS_MASTER	Specifies whether the Management Agent is a master or slave. Possible values: <ul style="list-style-type: none"> ■ 1: Master ■ 0: Slave
EMD_URL	URL of the emd monitoring the target
HOST_NAME	Host on which the target or Management Agent resides

Table 19–215 (Cont.) MGMT\$AGENTS_MONITORING_TARGETS

Column	Description
MONITORING	Indicates how the target is monitored. Possible values: <ul style="list-style-type: none"> MGMT_GLOBAL.G_MON_MODE_DEFAULT: Single Management Agent, vanilla monitoring mode MGMT_GLOBAL.G_MON_MODE_AGENT_MEDIATED: Multi Management Agent, agent-mediated monitoring

19.28.2 MGMT\$EM_ECM_MOS_PROPERTIES

The MGMT\$EM_ECM_MOS_PROPERTIES view exposes target information useful for My Oracle Support and Patching. It returns values from the EM_ECM_MOS_PROPERTIES table.

Table 19–216 MGMT\$EM_ECM_MOS_PROPERTIES

Column	Description
TARGET_GUID	The GUID of the target
HOST_TARGET_GUID	The GUID of the host that is hosting this target. The value is NULL for host targets.
ORACLE_HOME_TARGET_GUID	The GUID of the Oracle home target where this target is installed
PLATFORM_ID	The platform ID of the host that is hosting this target
PLATFORM_VERSION_ID	The platform version ID of the host that is hosting this target

19.28.3 MGMT\$EM_ECM_TARGET_FRESHNESS

The MGMT\$EM_ECM_TARGET_FRESHNESS view exposes the newest and oldest configuration snapshot information on a per target basis. It returns the collection timestamp and snapshot GUID for the most recently collected (newest) and least recently (oldest) collected snapshot for the target. If the target is a system, then the values will be valid across all snapshots of all member targets of the system.

Table 19–217 MGMT\$EM_ECM_TARGET_FRESHNESS

Column	Description
TARGET_GUID	The GUID of the target
NEWEST_SNAPSHOT_GUID	The snapshot GUID of the most recent collection
NEWEST_SNAPSHOT_TIMESTAMP	The date and time of the most recent collection
OLDEST_SNAPSHOT_GUID	The snapshot GUID of the oldest collection
OLDEST_SNAPSHOT_TIMESTAMP	The date and time of the oldest collection

19.28.4 MGMT\$MANAGEABLE_ENTITIES

The MGMT\$MANAGEABLE_ENTITIES view contains a list of all the Manageable Entities in Enterprise Manager.

Table 19–218 MGMT\$MANAGEABLE_ENTITIES

Column	Description
ENTITY_GUID	<p>Unique ID corresponding to the Manageable Entity.</p> <ul style="list-style-type: none"> For systems, services, groups, and so on, this ID contains the TARGET_GUID For target components, this ID contains the COMPONENT_GUID
ENTITY_TYPE	<p>Type of the Manageable Entity.</p> <ul style="list-style-type: none"> For systems, services, groups, and so on, this column contains the target type For target components, this column contains the component type
ENTITY_NAME	<p>Name of the Manageable Entity.</p> <ul style="list-style-type: none"> For systems, services, groups, and so on, this column contains the entity name For target components, this column contains the component name
PARENT_ME_GUID	<p>Parent entity containing the target component. For non-target components, this column is equal to ENTITY_GUID.</p>
PARENT_ME_TYPE	<p>Type of the parent target. (Valid for target components only)</p>
PARENT_ME_NAME	<p>Name of the parent target. (Valid for target components only)</p>
MANAGE_STATUS	<p>Manage status of the entity.</p> <p>Possible values:</p> <ul style="list-style-type: none"> 0: Ignored 1: Not managed yet 2: Managed 3: Managed target component
PROMOTE_STATUS	<p>Promote status of the target</p> <ul style="list-style-type: none"> 0: Cannot promote (existence only entity) 1: Eligible for promotion 2: Promotion in progress 3: Promoted
DYNAMIC_PROPERTY_STATUS	<p>Status of the dynamic properties</p> <p>Possible values:</p> <ul style="list-style-type: none"> 0: Dynamic properties have not been uploaded by the Management Agent 1: Dynamic properties are uploaded by the Management Agent
TYPE_META_VER	<p>Metadata version number</p>
CATEGORY_PROP_N	<p>Up to five category properties can be used to distinguish different metric definitions based on different configurations.</p> <p>For example, OS version, DB version, RAC configuration, and so on.</p>
TIMEZONE_REGION	<p>Name of the time zone region that the target operates in.</p> <p>Note: Component will be in same time zone as the target</p>
DISPLAY_NAME	<p>User-friendly name for the Manageable Entity</p>

Table 19–218 (Cont.) MGMT\$MANAGEABLE_ENTITIES

Column	Description
OWNER	Enterprise Manager administrator that owns the target
HOST_NAME	Host on which the target resides
EMD_URL	The URL for the EMD where the target is being collected
BROKEN_REASON	Code representing a reason why a target is broken Possible values: <ul style="list-style-type: none"> 0: Not broken 1: Missing required properties 2: Metadata not found 4: Error computing dynamic properties 8: Dynamic property missing in the result 16: Target name not specified 32: Target could not be saved 64: Errors in target test metrics
BROKEN_STR	String associated with the broken reason
MONITORING_MODE	Indicates how the target is monitored. Possible values: <ul style="list-style-type: none"> G_MON_MODE_DEFAULT(0): Single Management Agent, vanilla monitoring mode G_MON_MODE_OMS_MEDIATED(1): Multi Management Agent, OMS mediated monitoring G_MON_MODE_AGENT_MEDIATED(2): Multi Management Agent, agent mediated monitoring
IS_PROPAGATING	Specifies whether the Manageable Entity is privilege propagating
DISCOVERED_NAME	Name with which the Manageable Entity was discovered
ORG_ID	Organization ID used by the Oracle Configuration Manager (OCM) Harvester. This ID is used to uniquely identify customers. Note: A large customer might have multiple organization IDs to represent different lines of business within the company.
ORACLE_HOME	Oracle home of the target
ORACLE_CONFIG_HOME	This value is used by the OCM collector and CCR to indicate targets that might share a single Oracle home but have their configuration state kept in separate directories
LOAD_TIMESTAMP	The date and time when the Manageable Entity record was first loaded (in the time zone of the Management Repository)

19.29 VT Target Views

This section provides a description of each vt target view and its columns.

19.29.1 MGMT\$VT_VM_CONFIG

The MGMT\$VT_VM_CONFIG retrieves the Oracle VM Guest configuration.

Table 19–219 MGMT\$VT_VM_CONFIG

Column	Description
ECM_SNAPSHOT_ID	The ECM Snapshot ID corresponding to the latest ECM snapshot uploaded
TARGET_GUID	GUID of the target instance
BOOT_TYPE	Boot Type cdrom/disk/network
CPU_COUNT	Number of CPUs for this virtual machine
CPU_PRIORITY	Priority value for scheduling. Range 1-100
DOMAIN_TYPE	Domain type (hvm or pvm) of the virtual machine
HA_FLAG	Current state of the High Availability flag
MEM_REQD_MB	Amount of memory required to run this virtual machine
MOUSE_TYPE	The type of mouse device to be used with this virtual machine
REBOOT_LIMIT	Reboot limit
ALLOC_MEM_MB	Allocated memory in MB for this virtual machine
ALLOC_DISK_MB	Allocated disk space in MB for this virtual machine

19.29.2 MGMT\$VT_VM_SW_CFG

The MGMT\$VT_VM_SW_CFG retrieves the software related configuration for the Oracle VM Guest target

Table 19–220 MGMT\$VT_VM_SW_CFG

Column	Description
ECM_SNAPSHOT_ID	The ECM Snapshot ID corresponding to the latest ECM snapshot uploaded
TARGET_GUID	GUID of the target instance
OS_NAME	Name of the operating system
KERNEL_VER	Kernel version
USE_VS_TOOLS	Are VS tools used for this virtual machine?
VS_TOOLS_OS	Operating system as returned by VS tools
VS_TOOLS_VER	VS tools version
GUEST_UUID	UUID of the virtual machine in OVM Manager
VS_UUID	UUID of the server the virtual machine is running on
VSP_UUID	UUID of the server pool the virtual machine is contained in
OVM UUID	UUID of the OVM Manager the virtual machine is registered with. Has data only if the virtual machine is not part of a server pool.
OVM_DISPLAY_NAME	The current display name of the virtual machine in OVM. Need not be the same as OEM's target name.

19.29.3 MGMT\$VT_VM_VNIC

The MGMT\$VT_VM_VNIC retrieves Virtual Network Interface Card information for the virtual machine.

Table 19–221 MGMT\$VT_VM_VNIC

Column	Description
ECM_SNAPSHOT_ID	The ECM Snapshot ID corresponding to the latest ECM snapshot uploaded
TARGET_GUID	GUID of the target instance
MAC_ADDRESS	MAC address
ETH_NET	Ethernet network the VNIC is connected to
IP_ADDRESS	IP address. Each VNIC can have multiple IP addresses
IP_ADDRESS_TYPE	Address type - static/dhcp

19.29.4 MGMT\$VT_VM_VNIC_QOS

The MGMT\$VT_VM_VNIC_QOS displays the quality of service associated with the VNIC.

Table 19–222 MGMT\$VT_VM_VNIC_QOS

Column	Description
ECM_SNAPSHOT_ID	The ECM Snapshot ID corresponding to the latest ECM snapshot uploaded
TARGET_GUID	GUID of the target instance
MAC_ADDRESS	MAC address
QOS_NAME	Quality of service name

19.29.5 MGMT\$VT_VM_EM_CFG

The MGMT\$VT_VM_EM_CFG displays the OEM specified configuration for this virtual machine

Table 19–223 MGMT\$VT_VM_EM_CFG

Column	Description
ECM_SNAPSHOT_ID	The ECM Snapshot ID corresponding to the latest ECM snapshot uploaded
TARGET_GUID	GUID of the target instance
VM_TYPE	Type of virtual machine. For example, small/medium/large. These values are defined in OEM

19.29.6 MGMT\$VT_VM_VDISKS

The MGMT\$VT_VM_VDISKS displays virtual disk information for the virtual machine.

Table 19–224 MGMT\$VT_VM_VDISKS

Column	Description
ECM_SNAPSHOT_ID	The ECM Snapshot ID corresponding to the latest ECM snapshot uploaded
TARGET_GUID	GUID of the target instance
UUID	The UUID of the virtual disk in OVM Manager

Table 19–224 (Cont.) MGMT\$VT_VM_VDISKS

Column	Description
NAME	Display name of the virtual disk in OVM Manager
SIZE_MB	Size in MB
TYPE	Type of virtual disk
SHAREABLE	Is the virtual disk shareable or not?
REPO	The repository that the virtual disk comes from
DISK_MODE	Mode. For example, read-only or read-write.

19.29.7 MGMT\$VT_VM_VDISK_QOS

The MGMT\$VT_VM_VDISK_QOS displays virtual disk quality of service information.

Table 19–225 MGMT\$VT_VM_VDISK_QOS

Column	Description
ECM_SNAPSHOT_ID	The ECM Snapshot ID corresponding to the latest ECM snapshot uploaded
TARGET_GUID	GUID of the target instance
UUID	The UUID of the virtual disk in OVM Manager
QOS_NAME	Quality of service associated with the virtual disk

19.29.8 MGMT\$VT_EXA_CTRL_VSERVER_TAGS

The MGMT\$VT_EXA_CTRL_VSERVER_TAGS displays Exalogic Control VServer Tags.

Table 19–226 MGMT\$VT_EXA_CTRL_VSERVER_TAGS

Column	Description
ECM_SNAPSHOT_ID	The ECM Snapshot ID corresponding to the latest ECM snapshot uploaded
ID	The Identifier for the Tag
TYPE	Type of VServer
DESCRIPTION	Description for the VServer
SOFTWARE_NAME	Software Name on the VServer
URL	URL of the Software
VERSION	Version

19.29.9 MGMT\$VT_VSP_CONFIG

The MGMT\$VT_VSP_CONFIG view displays the server pool configuration.

Table 19–227 MGMT\$VT_VSP_CONFIG

Column	Description
ECM_SNAPSHOT_ID	The ECM Snapshot ID corresponding to the latest ECM snapshot uploaded
TARGET_GUID	GUID of the target instance

Table 19-227 (Cont.) MGMT\$VT_VSP_CONFIG

Column	Description
MASTER_VIP	Master Virtual IP Address
IS_DRS_ENABLED	Is Distributed Resource Scheduler enabled for this server pool?
IS_DPM_ENABLED	Is Distributed Power Management enabled for this server pool?
TOTAL_MEMORY_MB	Total memory across all servers in the server pool in MB
TOTAL_DISKSPACE_MB	Total storage across all servers in the server pool in MB
TOTAL_PHYSICAL_CPUS	Total number of physical CPUs across all servers in the server pool
TOTAL_VCPUS	Total number of VCPUs allocated across all virtual machines in the server pool
TOTAL_NICS	Total number of network interface cards across all servers in the server pool
CLSTR_ENABLED	Is the pool cluster enabled?
VM_MIGRATE_SECURE	Is secure virtual machine migration enabled?
POOL_FILESYSTEM	Server Pool filesystem information
VSP_UUID	The UUID of the server pool in OVM Manager
ZONE_UUID	The UUID of the parent zone in OVM Manager
OVM UUID	The UUID of the OVM Manager managing the server pool. Only collected if the server pool is not part of a zone.
OVM_DISPLAY_NAME	The display name of the server pool in OVM Manager. Need not be the same as OEM's target name

19.29.10 MGMT\$VT_VS_HW_CFG

The MGMT\$VT_VS_HW_CFG view displays the Oracle VM Server hardware configuration.

Table 19-228 MGMT\$VT_VS_HW_CFG

Column	Description
ECM_SNAPSHOT_ID	The ECM Snapshot ID corresponding to the latest ECM snapshot uploaded
TARGET_GUID	GUID of the target instance
VENDOR	Vendor
PROCESSOR_SPEED	CPU speed
PROCESSOR_TYPE	CPU type
NO_EXECUTE_FLAG	Is the no execute flag set?
NUM_POP_PROC SOCKS	Sockets filled
NUM_PROC SOCKS	Sockets per NUMA node
NUM_THREADS_PER_CORE	Threads per core
BIOS_RELEASE_DATE	BIOS release date
BIOS_VENDOR	BIOS vendor
BIOS_VERSION	BIOS version

Table 19–228 (Cont.) MGMT\$VT_VS_HW_CFG

Column	Description
MEMORY_MB	Server memory in MB
AVAILABLE_MEM_MB	Available memory on the server in MB
MEM_OVERHEAD_MB	Control domain memory in MB
USABLE_MEM_MB	Memory usable by virtual machines in MB
SWAP_SPACE_MB	Swap space on the server in MB
ADDR_LENGTH_BITS	Address length - for example, 32 bit and 64 bit
LOCAL_DISK_GB	Local disk space in GB
ENABLED_PROCESSORS	Number of enabled processors
ENABLED_CORES	Number of enabled cores
NUM_VCPUS	Number of VCPUs allocated on the server
CD_CPU_COUNT	Number of CPUS for the control domain
NUM_CORES	Number of cores
NUM_CORES_PER SOCK	Number of cores per socket

19.29.11 MGMT\$VT_VS_HYPERVISOR

The MGMT\$VT_VS_HYPERVISOR view displays Hypervisor related information.

Table 19–229 MGMT\$VT_VS_HYPERVISOR

Column	Description
ECM_SNAPSHOT_ID	The ECM Snapshot ID corresponding to the latest ECM snapshot uploaded
TARGET_GUID	GUID of the target instance
VERSION	Version
HVM_CAPABLE	Is the hypervisor HVM capable?
CAPABILITY	Capability supported by the hypervisor
TYPE	Type - for example, LDOM, VBOX, XEN

19.29.12 MGMT\$VT_VS_PROCESSORS

The MGMT\$VT_VS_PROCESSORS view displays information about processors on the server.

Table 19–230 MGMT\$VT_VS_PROCESSORS

Column	Description
ECM_SNAPSHOT_ID	The ECM Snapshot ID corresponding to the latest ECM snapshot uploaded
TARGET_GUID	GUID of the target instance
PROC_ID	Processor ID
CPU_FAMILY	CPU family the processor belongs to
CPU_MODEL	CPU Model
FAMILY	Family

Table 19–230 (Cont.) MGMT\$VT_VS_PROCESSORS

Column	Description
FLAGS	CPU flags
L1CACHE_KB	Size of L1 cache in KB
L2CACHE_KB	Size of L2 cache in KB
L3CACHE_KB	Size of L3 cache in KB
MANUFACTURER	Name of the manufacturer
MODEL_NAME	Model name
VENDOR_ID	Vendor ID

19.29.13 MGMT\$VT_VS_SW_CFG

The MGMT\$VT_VS_SW_CFG view displays

Table 19–231 MGMT\$VT_VS_SW_CFG

Column	Description
ECM_SNAPSHOT_ID	The ECM Snapshot ID corresponding to the latest ECM snapshot uploaded
TARGET_GUID	GUID of the target instance
AGENT_PORT	OVS agent port
AGENT_VERSION	OVS agent version
OVM_VERSION	OVM version
CONSOLE_SHELL_FLAG	Is the console shell flag set?
DISABLE_USB_FLAG	Is the disable USB flag set?
HALT_ON_ERROR_FLAG	Is the halt on error flag set?
KERNEL_RELEASE	Kernel release
KERNEL_VERSION	Kernel version
ASSOC_CLUSTER	Associated cluster
PROTECTED_FLAG	Is the protected flag set?
PYTHON_BIND_VER	Python binding version
RPM_VER	RPM version
VS_UUID	UUID of the server in OVM Manager
VSP_UUID	UUID of the parent server pool in OVM Manager
OVM UUID	UUID of the OVM Manager where the server is registered. Only contains data if the server is not part of a pool.
OVM_DISPLAY_NAME	The display name of the server in OVM Manager. Need not correspond to OEM target name.
CD_HOST_OS_NAME	Control domain host OS name
CD_HOST_OS_TYPE	Control domain host OS type
CD_HOST_OS_MAJOR_VER	Control domain host OS major version
CD_HOST_OS_MINOR_VER	Control domain host OS minor version

19.29.14 MGMT\$VT_VS_ATTRIBUTES

The MGMT\$VT_VS_ATTRIBUTES view displays miscellaneous server configuration information.

Table 19–232 MGMT\$VT_VS_ATTRIBUTES

Column	Description
ECM_SNAPSHOT_ID	The ECM Snapshot ID corresponding to the latest ECM snapshot uploaded
TARGET_GUID	GUID of the target instance
IS_MASTER	Is the server running as the master in the parent server pool?
UPGRADE_REQUIRED	Is upgrade required on the server?

19.29.15 MGMT\$VT_VS_FS_MOUNTS

The MGMT\$VT_VS_FS_MOUNTS view displays filesystem mounts on the server.

Table 19–233 MGMT\$VT_VS_FS_MOUNTS

Column	Description
ECM_SNAPSHOT_ID	The ECM Snapshot ID corresponding to the latest ECM snapshot uploaded
TARGET_GUID	GUID of the target instance
FILESERVER_UUID	UUID of the file server the file system is being served from
MOUNT_POINT	Filesystem mount point
REMOTE_PATH	The remote path of the filesystem
OPTIONS	Mount options
PHYSICAL_SIZE	Physical size of the filesystem

19.29.16 MGMT\$VT_VS_NET_DEVICE

The MGMT\$VT_VS_NET_DEVICE view displays information about network devices on the server.

Table 19–234 MGMT\$VT_VS_NET_DEVICE

Column	Description
ECM_SNAPSHOT_ID	The ECM Snapshot ID corresponding to the latest ECM snapshot uploaded
TARGET_GUID	GUID of the target instance
MAC_ADDRESS	MAC address
IP_ADDRESS	IP address
INTERFACE_NAME	Name of the interface
MTU	Maximum transmission unit
BANDWIDTH_MBPS	Bandwidth in MB/S
NETMASK	Netmask
ADDRESS_TYPE	Address type

19.29.17 MGMT\$VT_VS_FILESERVERS

The MGMT\$VT_VS_FILESERVERS view displays information about file servers associated with the server.

Table 19–235 MGMT\$VT_VS_FILESERVERS

Column	Description
ECM_SNAPSHOT_ID	The ECM Snapshot ID corresponding to the latest ECM snapshot uploaded
TARGET_GUID	GUID of the target instance
FILESERVER_UUID	UUID of the fileserver
FILESERVER	Name of the fileserver
FSNAME	Name of the filesystem (for example, nfs)
FSTYPE	Type (for example, networkfs)

19.29.18 MGMT\$VT_VS_REPOS

The MGMT\$VT_VS_REPOS view displays storage repositories visible to the server.

Table 19–236 MGMT\$VT_VS_REPOS

Column	Description
ECM_SNAPSHOT_ID	The ECM Snapshot ID corresponding to the latest ECM snapshot uploaded
TARGET_GUID	GUID of the target instance
FILESERVER_UUID	UUID of the fileserver
MOUNT_POINT	mountpoints on the virtual server from the fileserver
REPOSITORY_UUID	UUID of the repository
REPOSITORY	Name of the repository

19.29.19 MGMT\$VT_VS_ABILITIES

The MGMT\$VT_VS_ABILITIES view displays

Table 19–237 MGMT\$VT_VS_ABILITIES

Column	Description
ECM_SNAPSHOT_ID	The ECM Snapshot ID corresponding to the latest ECM snapshot uploaded
TARGET_GUID	GUID of the target instance
ABILITY	Name of the ability
ENABLED	Is the above ability enabled or not?

19.29.20 MGMT\$VT_ZONE_CONFIG

The MGMT\$VT_ZONE_CONFIG view displays the Oracle VM Zone configuration.

Table 19–238 *MGMT\$VT_ZONE_CONFIG*

Column	Description
ECM_SNAPSHOT_ID	The ECM Snapshot ID corresponding to the latest ECM snapshot uploaded
TARGET_GUID	GUID of the target instance
TOTAL_MEMORY_MB	Total memory across servers in the zone in MB
TOTAL_DISKSPACE_MB	Total diskspace across servers in the zone in MB
TOTAL_PHYSICAL_CPUS	Total number of physical CPUs across servers in the zone
TOTAL_NICS	Total number of NICS across servers in the zone
ZONE_UUID	UUID of the zone in OVM Manager
OVMM_UUID	UUID of the OVM Manager where this zone is registered
OVMM_DISPLAY_NAME	OVMM display name

19.30 Examples

This section provides examples of how to use the Management Repository views. It includes the following:

- [How do I create a derived associations rule which establishes associations between a host and oracle_vm_guest on which it is deployed?](#)
- [How do I return all targets under a blackout?](#)
- [How do I view a list of all the compliance rules?](#)
- [How do I view the monitoring compliance rules only?](#)
- [How do I view all the repository compliance rules for a specific author?](#)
- [How do I view a list of all the compliance standards?](#)
- [How do I view all compliance standards owned by a specific user](#)
- [How do I view a list of all the compliance standard groups?](#)
- [How do I view all compliance standard groups in production?](#)
- [How do I query results for compliance standards with no included standards](#)
- [How do I obtain the results for compliance standards with included standards?](#)
- [How do I obtain the results for compliance standard rules in a compliance standard for a target?](#)
- [How do I obtain the results for compliance standard groups?](#)
- [How do I obtain association information for compliance standards and targets?](#)
- [How do I obtain the violation ID for an active violation of a compliance rule?](#)
- [How do I obtain the violation column information?](#)
- [How do I access the compliance rule violation context definition-related metadata?](#)
- [How do I find all bundles that are in violation?](#)
- [How do I find all observations \(all states\) for all bundles in violation?](#)
- [How do I get a list of all the actions that occurred on all targets during a specific time range?](#)

- How do I get a list of all actions that occurred on a single target during a specific time range?
- How do I get a list of all the file changes that occurred on a single target during a specific time range?
- How do I get a list of all unauthorized actions that occurred during a specific time range?
- How do I get a list of all occurrences of sudo?
- How do I get the number of targets for a metric?
- How do I get the number of Management Agents for a version?
- How do I get the listener port for each database?
- How do I get the number of databases for each category version?
- How do I get the number of databases for each category version and CPU count?
- How do I get the number of databases for each category version and OS platform?
- How do I get database metrics with outstanding severities?
- How do I get a list of all disabled metrics on Management Agents?
- How do I get the number of down targets?
- How do I get the availability information for the Enterprise Manager website?
- How do I return the current thresholds for the alertlog metric?
- How do I get the number of alertlog severities for the database in the last 24 hours?
- How do I get the current CPU utilization of a host?
- How do I get a list of all the collected user-defined metrics (UDMs)?
- How do I get the first byte response for the Enterprise Manager website at a specific time?
- How do I obtain the average number of connections for a listener for a specific period?
- How do I retrieve information from MGMT\$OS_SUMMARY for a specific host from the Management Repository?

How do I create a derived associations rule which establishes associations between a host and oracle_vm_guest on which it is deployed?

The following example is an example of a rule between an oracle_vm_guest target type and a host. The rule relies on a published EDK [MGMT\\$HW_NIC](#) view for the host and an ECM-generated CM\$VT_VM_VNIC view. While the rule resides in the plug-in that defines the oracle_vm_guest target type, then it can reference the CM\$ view for the snapshot type belonging to that target type and any EDK-accessible view (such as MGMT\$ views) from the host target type, which might reside in a different plug-in.

```
<Rule name="host_deployed_on_oracle_vm_guest">
  <query>
    select 'deployed_on' as assoc_type,
           host.target_guid as source_me_guid,
           guest.cm_target_guid as dest_me_guid
    from   mgmt$hw_nic host,
           cm$vt_vm_vnic guest
```

```

        where guest.mac_address = host.mac_address_std
    </query>
    <trigger>
        <targetType>host</targetType>
        <snapshotType>ll_host_config</snapshotType>
        <table>MGMT$HW_NIC</table>
        <idColumn>source</idColumn>
    </trigger>
    <trigger>
        <targetType>oracle_vm_guest</targetType>
        <snapshotType>ovm_guest_config</snapshotType>
        <table>CM$VT_VM_VNIC</table>
        <idColumn>destination</idColumn>
    </trigger>
</Rule>

```

How do I return all targets under a blackout?

To return all targets under blackout, enter the following query:

```

SELECT target_name, target_type, start_time, end_time
FROM   mgmt$blackout_history
WHERE  sysdate BETWEEN start_time AND NVL(end_time,sysdate+1/60*60*24);

```

How do I view a list of all the compliance rules?

To view a list of all the compliance rules, enter the following query:

```

select * from mgmt$compliance_standard_rule;

```

How do I view the monitoring compliance rules only?

To view the monitoring compliance rules only, enter the following query:

```

select * from mgmt$compliance_standard_rule where RULE_TYPE='Monitoring';

```

How do I view all the repository compliance rules for a specific author?

To view all the repository compliance rules where the author is John Smith;, enter the following query:

```

select * from mgmt$compliance_standard_rule where RULE_TYPE='Repository' AND
AUTHOR='John Smith';

```

How do I view a list of all the compliance standards?

The following queries provide examples about how to use this view:

To view a list of all the compliance standards, enter the following query:

```

select * from mgmt$compliance_standard;

```

How do I view all compliance standards owned by a specific user

To view all compliance standards owned by John Smith, enter the following query:

```

select * from mgmt$compliance_standard where OWNER='John Smith';

```

How do I view a list of all the compliance standard groups?

To view a list of all the compliance standard groups, enter the following query:

```
select * from mgmt$compliance_standard_group;
```

How do I view all compliance standard groups in production?

To view all compliance standard groups in production, enter the following query:

```
select * from mgmt$compliance_standard_group where LIFECYCLE_STATUS='Production';
```

How do I query results for compliance standards with no included standards

To query results for compliance standards with no included standards, enter the following query, where ? represents the values for each attribute:

```
select * from mgmt$cs_eval_summary where cs_guid = ? and target_guid = ?;
```

Note: To obtain CS_GUID, query the [MGMT\\$COMPLIANCE_STANDARD](#) view on compliance standard attributes such as name or target type. For example:

```
select CS_NAME from mgmt$compliance_standard;
```

How do I obtain the results for compliance standards with included standards?

The following queries provide examples of how to use this view:

To obtain the results for compliance standards with included standards, enter the following query, where ? represents the value of each attribute:

```
select * from mgmt$composite_cs_eval_summary where root_cs_guid = ? and root_target_guid = ?;
```

This query returns values for all the following possible results for the root compliance standard or root target:

- ROOT_GUID
- RQS_GUID
- CS_GUID
- ROOT_TARGET_GUID

How do I obtain the results for compliance standard rules in a compliance standard for a target?

To obtain the results for compliance standard rules in a compliance standard for a target, enter the following query where ? represents the values for ROOT_CS_GUID and ROOT_TARGET_GUID:

```
select * from mgmt$cs_rule_eval_summary where root_cs_guid = ? and root_target_guid = ?;
```

How do I obtain the results for compliance standard groups?

To obtain the results for compliance standard groups, enter the following query where ? is the value for CS_GUID:

```
select * from mgmt$cs_group_eval_summary where cs_guid = ?;
```

Note: To obtain CS_GUID, query the [MGMT\\$COMPLIANCE_STANDARD](#) view on compliance standard attributes such as name or target type. For example:

```
select CS_NAME from mgmt$compliance_standard;
```

How do I obtain association information for compliance standards and targets?

To obtain association information for compliance standards and targets, enter the following query where ? represents the value for each of the attributes:

```
select * from mgmt$cs_target_assoc where cs_guid = ? and target_guid = ?;
```

Note: To obtain CS_GUID, query the [MGMT\\$COMPLIANCE_STANDARD](#) view on compliance standard attributes such as name or target type. For example:

```
select CS_NAME from mgmt$compliance_standard;
```

How do I obtain the violation ID for an active violation of a compliance rule?

To obtain the violation GUID, enter the following query, where ? represents the value for RULE_GUID:

```
select * from mgmt$csr_current_violation where rule_guid=?;
```

Note: To obtain the RULE_GUID, query the [MGMT\\$CS_RULE_EVAL_SUMMARY](#) view.

For more information, see ["How do I obtain the results for compliance standard rules in a compliance standard for a target?"](#).

How do I obtain the violation column information?

To obtain the additional columns defined in a compliance rule to be collected for a violation, enter the following query where ? represents the value for VIOLATION_GUID:

```
select * from mgmt$csr_violation_context where violation_guid=?;
```

Note: To obtain the VIOLATION_GUID, query the [MGMT\\$CSR_CURRENT_VIOLATION](#) view. For example:

```
select * from mgmt$csr_current_violation where rule_guid=?;
```

For more information, see ["How do I obtain the violation ID for an active violation of a compliance rule?"](#).

How do I access the compliance rule violation context definition-related metadata?

To access the compliance rule violation context definition-related metadata, enter the following query, where ? represents the value for RULE_GUID:

```
select * from mgmt$em_rule_viol_ctxt_def where rule_guid=?;
```

Note: To obtain the RULE_GUID, query the [MGMT\\$CS_RULE_EVAL_SUMMARY](#) view.

For more information, see ["How do I obtain the results for compliance standard rules in a compliance standard for a target?"](#).

How do I find all bundles that are in violation?

To find all bundles that are in violation (that is, at least one unauthorized observation in the bundle), enter the following query:

```
select * from mgmt$ccc_all_obs_bundles where bundle_in_violation = 'true';
```

How do I find all observations (all states) for all bundles in violation?

To find all observations (all states) for all bundles in violation, enter the following query:

```
select * from mgmt$ccc_all_observations o, mgmt$ccc_all_obs_bundles b where  
o.bundle_id=b.bundle_id and b.bundle_in_violation='true';
```

How do I get a list of all the actions that occurred on all targets during a specific time range?

To get a list of all the actions that occurred on all targets during a specific time range, enter the following query:

```
select * from mgmt$ccc_all_observations where action_time between hh:mm and hh:mm;
```

How do I get a list of all actions that occurred on a single target during a specific time range?

To get a list of all actions that occurred on a single target during a specific time range, enter the following query:

```
select * from mgmt$ccc_all_observations where action_time between hh:mm and hh:mm  
and target = target_name;
```

How do I get a list of all the file changes that occurred on a single target during a specific time range?

To get a list of all the file changes that occurred on a single target during a specific time range, enter the following query:

```
select * from mgmt$ccc_all_observations where action_time between hh:mm and hh:mm  
and target = target_name and entity_type = 'OS File';
```

Note: You can replace 'OS File' with any entity type from the Cloud Control console, such as 'OS Process' or 'OS User'.

How do I get a list of all unauthorized actions that occurred during a specific time range?

To get a list of all unauthorized actions that occurred during a specific time range, enter the following query:

```
select * from mgmt$ccc_all_observations where action_time between hh:mm and hh:mm
and target = target_name and audit_status='Unauthorized';
```

How do I get a list of all occurrences of sudo?

To get a list of all occurrences of sudo, enter the following query:

```
select * from mgmt$ccc_all_observations where action_time between hh:mm and hh:mm
and target = target_name and action = 'osuser_sudo_suc';
```

All possible actions can be seen in the EM_CCC_META_OBSTYPE table.

How do I get the number of targets for a metric?

To return the number of targets for a metric, enter the following query:

```
SELECT  metric_name, COUNT(DISTINCT target_name)
FROM    mgmt$target_type
WHERE   target_type = 'oracle_database'
GROUP BY metric_name;
```

How do I get the number of Management Agents for a version?

To return the number of Management Agents for a version, enter the following query:

```
SELECT  property_value, COUNT(*)
FROM    mgmt$target_properties
WHERE   target_type = 'oracle_emd'
        AND property_name = 'Version'
GROUP BY property_value;
```

How do I get the listener port for each database?

To return the listener port for each database, enter the following query:

```
SELECT target_name, property_value
FROM    mgmt$target_properties
WHERE   target_type = 'oracle_database'
        AND property_name = 'Port';
```

How do I get the number of databases for each category version?

To return the number of databases for each category version, enter the following query:

```
SELECT  property_value, COUNT(*)
FROM    mgmt$target_properties
WHERE   target_type = 'oracle_database'
        AND property_name = 'VersionCategory'
GROUP BY property_value;
```

How do I get the number of databases for each category version and CPU count?

To return the number of databases for each category version and CPU count, enter the following query:

```
SELECT  p1.property_value "Version", p2.property_value "CPU Count", COUNT(*)
        "Total"
FROM    mgmt$target_properties p1, mgmt$target_properties p2
WHERE   p1.target_type = 'oracle_database'
        AND p1.target_guid = p2.target_guid
        AND p1.property_name = 'VersionCategory'
        AND p2.property_name = 'CPUCount'
GROUP BY p1.property_value, p2.property_value
ORDER BY p1.property_value, p2.property_value;
```

How do I get the number of databases for each category version and OS platform?

To return the number of databases for each category version and OS platform, enter the following query:

```
SELECT  p3.property_value "Platform", p1.property_value "Version", COUNT(*)
        "Total"
FROM    mgmt$target_properties p1, mgmt$target_properties p2, mgmt$target_
        properties p3
WHERE   p1.target_type = 'oracle_database'
        AND p1.target_guid = p2.target_guid
        AND p3.target_name = p2.property_value
        AND p3.target_type = 'host'
        AND p1.property_name = 'VersionCategory'
        AND p2.property_name = 'MachineName'
        AND p3.property_name = 'OS'
GROUP BY p3.property_value, p1.property_value
ORDER BY p3.property_value, p1.property_value;
```

How do I get database metrics with outstanding severities?

To return database metrics with outstanding severities, enter the following query:

```
SELECT  target_name, metric_name, COUNT(*),
        TO_CHAR(MAX(collection_timestamp), 'DD-MON-YYYY HH24:MI:SS')
FROM    mgmt$alert_current
WHERE   target_type = 'oracle_database'
GROUP BY target_name, metric_name;
```

How do I get a list of all disabled metrics on Management Agents?

To return a list of all disabled metrics on Management Agents, enter the following query:

```
SELECT collection_name, COUNT(*) nr_agents
FROM    mgmt$target_metric_collections
WHERE   is_enabled = 0
GROUP BY collection_name
ORDER BY collection_name
;
```

How do I get the number of down targets?

To return the number of down targets, enter the following query:

```
SELECT COUNT(*)
FROM   mgmt$availability_current
WHERE  availability_status='Target Down';
```

How do I get the availability information for the Enterprise Manager website?

To return the availability information for the Enterprise Manager website, enter the following query:

```
SELECT status, ROUND(duration,2) duration,
       ROUND((RATIO_TO_REPORT(duration) OVER ())*100,2) AS total
FROM   (SELECT NVL(availability_status,'-unknown-') status,
              SUM(NVL(end_timestamp,SYSDATE)-start_timestamp) duration
        FROM   mgmt$availability_history
        WHERE  target_name = 'Enterprise Manager'
              AND target_type = 'website'
        GROUP BY availability_status);
```

How do I return the current thresholds for the alertlog metric?

To return the current thresholds for the alertlog metric, enter the following query:

```
SELECT target_name, metric_column, warning_operator, warning_threshold, critical_
operator, critical_threshold
FROM   mgmt$metric_collection
WHERE  target_type = 'oracle_database'
      AND metric_name = 'alertLog'
ORDER BY target_name, metric_column;
```

How do I get the number of alertlog severities for the database in the last 24 hours?

To view the number of alertlog severities for the database in the last 24 hours, enter the following query:

```
SELECT target_name, COUNT (*)
FROM   mgmt$alert_history
WHERE  target_type = 'oracle_database'
      AND metric_name = 'alertlog'
      AND collection__timestamp > SYSDATE-1
GROUP BY target_name;
```

How do I get the current CPU utilization of a host?

To return the current CPU utilization of a host, enter the following query:

```
SELECT column_label, value
FROM   mgmt$metric_current
WHERE  metric_name = 'Load'
      AND metric_column = 'cpuUtil'
      AND target_name = 'my.acme.com';
```

How do I get a list of all the collected user-defined metrics (UDMs)?

To return a list of all the collected user-defined metrics (UDMs), enter the following query:

```
SELECT key_value udm_name, target_name, target_type, collection_timestamp, value
FROM   sysman.mgmt$metric_current
WHERE  metric_label = 'User Defined Metrics'
ORDER BY udm_name, target_type, target_name, collection_timestamp DESC;
```

How do I get the first byte response for the Enterprise Manager website at a specific time?

To return the first byte response for the Enterprise Manager Web site at 11.00 am yesterday, enter the following query:

```
SELECT target_name, AVG(average)
FROM   mgmt$metric_hourly
WHERE  target_name = 'EM Website'
      AND metric_name = 'http_response'
      AND metric_column = 'avg_first_byte_time'
      AND rollup_timestamp = TO_DATE(TO_CHAR(TRUNC(sysdate-1), 'DD-MON-YYYY') || '
11:00:00', 'DD-MON-YYYY HH24:MI:SS')
GROUP BY target_name;
```

How do I obtain the average number of connections for a listener for a specific period?

To return the average number of connections for a listener for the last seven days, enter the following query:

```
SELECT target_name, average
FROM   mgmt$metric_daily
WHERE  target_type = 'oracle_listener'
      AND metric_name = 'Load'
      AND metric_column = 'estConns'
      AND rollup_timestamp = TRUNC(sysdate-7);
```

How do I retrieve information from MGMT\$OS_SUMMARY for a specific host from the Management Repository?

The following query retrieves information from MGMT\$OS_SUMMARY for a specific host from the Management Repository:

```
select * from MGMT$OS_SUMMARY
where target_name = 'target_name' and target_type = 'host';
```

If you know a host name, you can use a similar query to access any of the views to retrieve information at the metric level for a particular host.

Aggregate queries can be written to provide counts of OS as follows:

```
SQL:
select name, base_version, count(*)
from mgmt$os_summary
group by name, base_version
;
```

Using Receivelets

This chapter contains the following sections:

- [About Receivelets](#)
- [SNMP Receivelet](#)

20.1 About Receivelets

A receivelet is a library that allows Enterprise Manager to receive external notifications sent by third-party network elements. These are notifications that are asynchronously sent and without any requests from the Management Agent.

Usually, the Management Agent data retrieval mechanism is based on a polling model, that is, modular libraries, called fetchlets. Fetchlets collect values of various metrics from their managed targets on a regular basis. The Management Agent then compares the gathered data with user-defined thresholds and generates events when the thresholds were met.

Receivelets are a more efficient way of dealing with these metrics. It depends on the ability of the managed target to detect the condition for its own events, and then communicate with Enterprise Manager only when an event occurs. When this communication happens, the Management Agent uses receivelets to receive the information.

You can use receivelets as a quicker way to get alerts on data that will be eventually collected via fetchlets. You can also use receivelets as a way to send both alerts and data, or just alerts for cases where there is no real data chart associated with the alert.

A receivelet is not a substitute for a fetchlet, but it is another way of collecting data. It is more for immediacy of notification compared to periodic polling that the fetchlet offers. Therefore, if you can fetch data, then use fetchlets to get that data. However, if your server is capable of sending you events or data at a cost lower than that associated with fetchlets, then use receivelets.

A receivelet may be tightly coupled to a particular type of managed target, or may be useful to a broad range of potential targets.

The SNMP receivelet is offered with Enterprise Manager as described in the following section.

20.2 SNMP Receivelet

An SNMP receivelet allows you to receive SNMP trap notifications from third-party network elements, and translate them into a form compatible with Oracle Management Service (OMS).

While monitoring third-party entities in your managed environment, if the third-party entity wants to send a notification to Enterprise Manager, then the SNMP agent of that third-party entity sends a notification to the Management Agent. These notifications are in the form of SNMP traps that get triggered asynchronously and without any requests from the Management Agent.

Since these traps are based on SNMP, the Management Agent uses SNMP receivelets to receive and translate these SNMP traps into a form compatible with OMS.

When the SNMP traps are received, the SNMP receivelet extracts information pertaining to those object identifiers (OIDs) that are defined in the `<PushDescriptor>` section of the target type metadata file only. For more information about the target type metadata file, see [Section 3.3, "Creating the Target Type Metadata File"](#).

When a customer wants to manage a network element using the SNMP receivelet, they must configure the element's SNMP agent to send traps to the responsible Management Agent's SNMP receivelet. When the SNMP receivelet receives such traps, it translates them to an Enterprise Manager format (such as an event or datapoint based on the *push descriptor* information), and stores that information (in XML files) in the upload directory. The Upload Manager checks for such new files in the upload directory, and then uploads those files to OMS. Then Enterprise Manager accesses OMS to extract the collected information and displays it to the user.

Receiving SNMP Traps

To receive SNMP traps, you have to make some configuration settings at the Management Agent side and at SNMP target agent side.

This enables the SNMP targets to send SNMP traps to the SNMP receivelet. When the SNMP traps are received, the SNMP receivelet uses the Push Descriptor properties, such as `MatchAgentAddr`, `MatchEnterprise`, and so on, to identify the target and metric for which the traps belongs. Then the SNMP receivelet uses the Push Descriptor properties, such as `Eventmetric-column` or `Eventmetric-columnOID`, `SeverityCode`, and so on, to convert the traps into an event. When this happens, the SNMP receivelet uploads the converted event to the Management Repository and it is now available in the Cloud Control console.

Configuration Settings Required at the Management Agent Side

By default, the SNMP receivelet listens over UDP on the same port as that of the Management Agent. However, if you want to use a different listening port for the SNMP receivelet, then add the `SnmpRecvletListenNIC(=8002)` entry to the `emd.properties` file.

Configuration Settings Required at SNMP Target Agent Side

Configure the SNMP target agent to send its traps to the SNMP target agent's host name and port.

Input Parameters

Table 20–1 *SNMP Receivelet Input Parameters*

Parameter	Type	Description	Use
MatchEnterprise	String	<p>Note: For Push Descriptors that intend to match SNMPv1 traps only.</p> <p>OID used to define the trap being sent.</p> <p>Note: If MatchEnterprise is present, then you must include MatchGenericTrap and MatchSpecificTrap also.</p>	Required (SNMPv1)
MatchGenericTrap	String	Code for a generic SNMP trap.	Required (SNMPv1)
MatchSpecificTrap	String	Trap defined in a MIB (not one of the generic traps), the ID assigned in that MIB.	Required (SNMPv1)
MatchTrapOID	String	<p>Note: For Push Descriptors that intend to match SNMPv2 or SNMPv3 traps.</p> <p>For an SNMPv2 or SNMPv3 trap, this is the OID assigned to the NOTIFICATION-TYPE in the MIB definition of the trap.</p>	Required (SNMPv2, SNMPv3)
MatchAgentAddr	String	IP address of the generating SNMP agent, as sent in the trap.	Required
MatchVarBind	String	If this parameter is present, its value is an OID. If a trap is received that matches the other Match* parameters, it still must have this OID on its received varbind list. Otherwise, it will not generate an Enterprise Manager event or datapoint.	Optional
UseCredential	String	Specifies credential use. If it is set to TRUE, then the receivelet will accept SNMPv3 traps using a set of target SNMPv3Creds sent to the receivelet as a CredentialRef. Also, if UseCredential is set, then the string property "hostname" and the numeric property "PORT" must be specified; their values will be used for engine-id discovery.	Optional
Eventmetric-column	String	Specifies that, on receiving this trap, the receivelet must generate a severity on this metric column. The value of <i>metric-column</i> must be the value of this parameter. (This case is useful where the expected values of the Enterprise Manager metric are not the same as the triggering SNMP variable.)	Required, if events have to be generated. However, if Eventmetric-columnOID is provided, then this is not required.
Eventmetric-columnOID	String	Specifies that, on receiving this trap, the receivelet must generate a severity on this metric column. The value of the metric column should be taken from the varbind in the trap with an OID equal to the value of this parameter.	Required, if events have to be generated. However, if Eventmetric-column is provided, then this is not required.

Table 20–1 (Cont.) SNMP Receivelet Input Parameters

Parameter	Type	Description	Use
SeverityCode	String	Specifies the level at which the severity must be generated. The value of this parameter must be 'CRITICAL', 'WARNING', or 'CLEAR'. Note: SeverityCode or SeverityCodeOID must be present if events are to be generated when the trap is received. However, if the Push Descriptor intends to generate a metric datapoint (and specifies one or more <i>Datametric-columnOID</i> properties), then it must <i>not</i> have a SeverityCode or SeverityCodeOID property.	Required. However, if SeverityCodeOID is provided, or the Push Descriptor is generating a metric datapoint, then this is not required.
SeverityCodeOID	String	Specifies the level at which the severity should be generated. If the varbind in the trap with OID equal to the value of this parameter is one of the strings 'CRITICAL', 'WARNING', or 'CLEAR', then the severity must be generated at that level; otherwise, no severity must be generated. (This parameter is only used if you are designing a trap exclusively for use with Enterprise Manager, but can be useful in this case.) Note: SeverityCode or SeverityCodeOID must be present if events are to be generated when the trap is received. However, if the Push Descriptor intends to generate a metric datapoint (and specifies one or more <i>Datametric-columnOID</i> properties), then it must <i>not</i> have a SeverityCode or SeverityCodeOID property.	Required. However, if SeverityCode is provided, or the PushDescriptor is generating a metric datapoint, then this is not required.
<i>Datametric-columnOID</i>	String	Specifies that, on receiving this trap, the receivelet must generate a datapoint on the metric, for which the value of this metric column should be taken from the varbind in the trap with OID equal to the value of this parameter. (An SNMP Push Descriptor can have many <i>Data*</i> parameters, in which case a single row will be returned, with all specified columns populated from the appropriate varbind in the trap. An SNMP Push Descriptor <i>cannot</i> have both a <i>Data*</i> parameter and a <i>Severity*</i> parameter, nor can it have multiple <i>Severity*</i> parameters.)	Required, if datapoints have to be generated.
<i>Keymetric-columnOID</i>	String	Severity or datapoint generated by this Push Descriptor should contain a key-value for this metric column. The key-value should be taken from the varbind in the trap with OID equal to the value of this parameter. For every key-column in the metric definition, there must be a <i>Key*</i> parameter in the Push Descriptor.	Optional
<i>Contextmetric-columnOID</i>	String	If the Push Descriptor generates a severity, then the severity must contain a value for this metric column in its event context. The value should be taken from the varbind in the trap with OID equal to the value of this parameter. If the Push Descriptor generates a datapoint, then this parameter is ignored.	Optional. This can be used for severities only.

Example

[Example 20–1](#) shows how a trap from a vendor-specific router looks like.

Example 20–1 Trap from a Vendor-Specific Router

```
ascendLinkDown      TRAP-TYPE
  ENTERPRISE        ascend
  VARIABLES          { ifIndex, ifAdminStatus, ifOperStatus, ifType,
                      ifName }
  DESCRIPTION        "This trap is in addition to the generic linkDown
                      trap defined in RFC1215. This trap provides
                      additional information such as ifAdminStatus,
                      ifOperStatus, ifName, slotIfSlotIndex, slotIfItemIndex.
                      This is an Alarm class trap and it can
                      be enabled/disabled via alarmEnabled and/or
                      ascendLinkDownTrapEnabled in trap profile.
                      This trap is sent only if rfc1215 linkDown trap is generated."

 ::= 50
```

[Example 20–2](#) show how the trap will be received by the Management Agent. Note that <x> in this example is the value of *ifIndex* that identifies the particular interface that's having problems.

Example 20–2 Trap Received by the Management Agent

```
Message:
  version: 0
  community: 'public'
  Trap-PDU:
    enterprise: enterprises.ascend (1.3.6.1.4.1.529)
    agent-addr: 138.2.204.10
    generic-trap: 6
    specific-trap: 50
    time-stamp: <timestamp from router's sysUptime>
    variable-bindings:
      Name: ifIndex.<x> (1.3.6.1.2.1.2.2.1.1.<x>)
      Type: INTEGER
      Value: <x>

      Name: ifAdminStatus.<x> (1.3.6.1.2.1.2.2.1.7.<x>)
      Type: INTEGER
      Value: up (1)

      Name: ifOperStatus.<x> (1.3.6.1.2.1.2.2.1.8.<x>)
      Type: INTEGER
      Value: down (2)

      Name: ifType.<x> (1.3.6.1.2.1.2.2.1.3.<x>)
      Type: INTEGER
      Value: iso88023-csmacd (7)

      Name: ifName.<x> (1.3.6.1.2.1.2.2.1.31.<x>)
      Type: DisplayString
      Value: 'eth0'
```

[Example 20–3](#) shows how the metric can be defined in the target type metadata file.

Example 20–3 Metric Defined in the Target Type Metadata File

```
<Metric NAME="interfaces" TYPE="TABLE">
```

```

<TableDescriptor>
  <ColumnDescriptor NAME="name" TYPE="STRING" IS_KEY="TRUE" />
  <ColumnDescriptor NAME="type" TYPE="NUMBER" IS_KEY="FALSE" />
  <ColumnDescriptor NAME="status" TYPE="NUMBER" IS_KEY="FALSE" />
  <ColumnDescriptor NAME="configured_status" TYPE="NUMBER" IS_KEY="FALSE" />
</TableDescriptor>
</Metric>

```

To Receive SNMPV1 Trap

[Example 20–4](#) shows how the push descriptor can be defined in the target type metadata file to trigger a severity.

Example 20–4 Push Descriptor in the Target Type Metadata File For Triggering a Severity

```

<PushDescriptor RECVLET_ID="SNMPTrap">
  <Property NAME="MatchEnterprise" SCOPE="GLOBAL">1.3.6.1.4.1.529</Property>
  <Property NAME="MatchGenericTrap" SCOPE="GLOBAL">6</Property>
  <Property NAME="MatchSpecificTrap" SCOPE="GLOBAL">50</Property>
  <Property NAME="MatchAgentAddr" SCOPE="INSTANCE">AdminAddress</Property>
  <Property NAME="SeverityStatusOID"
SCOPE="GLOBAL">1.3.6.1.2.1.2.2.1.8</Property>
  <Property NAME="KeyNameOID" SCOPE="GLOBAL">1.3.6.1.2.1.2.2.1.31</Property>
  <Property NAME="ContextTypeOID" SCOPE="GLOBAL">1.3.6.1.2.1.2.2.1.3</Property>
  <Property NAME="ContextConfigured_statusOID"
SCOPE="GLOBAL">1.3.6.1.2.1.2.2.1.7</Property>
  <Property NAME="SeverityCode" SCOPE="GLOBAL">CRITICAL</Property>
  <CredentialRef NAME="monCreds">monCredentials</CredentialRef>
</PushDescriptor>

```

[Example 20–5](#) shows how the push descriptor can be defined in the target type metadata file to trigger a datapoint, which would specify the reporting of data on the same trap, with *ifName* as the key-column and the other three as data columns.

Example 20–5 Push Descriptor in the Target Type Metadata File For Triggering a Datapoint

```

<PushDescriptor RECVLET_ID="SNMPTrap">
  <Property NAME="MatchEnterprise" SCOPE="GLOBAL">1.3.6.1.4.1.529</Property>
  <Property NAME="MatchGenericTrap" SCOPE="GLOBAL">6</Property>
  <Property NAME="MatchSpecificTrap" SCOPE="GLOBAL">50</Property>
  <Property NAME="MatchAgentAddr" SCOPE="INSTANCE">AdminAddress</Property>
  <Property NAME="KeyNameOID" SCOPE="GLOBAL">1.3.6.1.2.1.2.2.1.31</Property>
  <Property NAME="DataStatusOID" SCOPE="GLOBAL">1.3.6.1.2.1.2.2.1.8</Property>
  <Property NAME="DataTypeOID" SCOPE="GLOBAL">1.3.6.1.2.1.2.2.1.3</Property>
  <Property NAME="DataConfigured_statusOID"
SCOPE="GLOBAL">1.3.6.1.2.1.2.2.1.7</Property>
  <CredentialRef NAME="monCreds">monCredentials</CredentialRef>
</PushDescriptor>

```

To Receive SNMPV2 Notifications

[Example 20–6](#) shows how the push descriptor can be defined in the target type metadata file to trigger a severity:

Example 20–6 Push Descriptor in the Target Type Metadata File For Triggering a Severity

```

<PushDescriptor RECVLET_ID="SNMPTrap">
  <Property NAME="MatchTrapOID" SCOPE="GLOBAL">1.3.6.1.4.1.529.50</Property>
  <Property NAME="MatchAgentAddr" SCOPE="INSTANCE">AdminAddress</Property>

```

```

<Property NAME="SeverityStatusOID"
SCOPE="GLOBAL">1.3.6.1.2.1.2.2.1.8</Property>
<Property NAME="KeyNameOID" SCOPE="GLOBAL">1.3.6.1.2.1.2.2.1.31</Property>
<Property NAME="ContextTypeOID" SCOPE="GLOBAL">1.3.6.1.2.1.2.2.1.3</Property>
<Property NAME="ContextConfigured_statusOID"
SCOPE="GLOBAL">1.3.6.1.2.1.2.2.1.7</Property>
<Property NAME="SeverityCode" SCOPE="GLOBAL">CRITICAL</Property>
<CredentialRef NAME="monCreds">monCredentials</CredentialRef>
</PushDescriptor>

```

[Example 20–7](#) shows how the push descriptor can be defined in the target type metadata file to trigger a datapoint, which would specify the reporting of data on the same trap, with *ifName* as the key-column and the other three as data columns.

Example 20–7 Push Descriptor in the Target Type Metadata File For Triggering a Datapoint

```

<PushDescriptor RECVLET_ID="SNMPTrap">
  <Property NAME="MatchEnterprise" SCOPE="GLOBAL">1.3.6.1.4.1.529</Property>
  <Property NAME="MatchGenericTrap" SCOPE="GLOBAL">6</Property>
  <Property NAME="MatchSpecificTrap" SCOPE="GLOBAL">50</Property>
  <Property NAME="MatchAgentAddr" SCOPE="INSTANCE">AdminAddress</Property>
  <Property NAME="KeyNameOID" SCOPE="GLOBAL">1.3.6.1.2.1.2.2.1.31</Property>
  <Property NAME="DataStatusOID" SCOPE="GLOBAL">1.3.6.1.2.1.2.2.1.8</Property>
  <Property NAME="DataTypeOID" SCOPE="GLOBAL">1.3.6.1.2.1.2.2.1.3</Property>
  <Property NAME="DataConfigured_statusOID"
SCOPE="GLOBAL">1.3.6.1.2.1.2.2.1.7</Property>
  <CredentialRef NAME="monCreds">monCredentials</CredentialRef>
</PushDescriptor>

```

To Receive SNMPV3 Notifications

[Example 20–8](#) shows how the push descriptor can be defined in the target type metadata file to trigger a severity.

Example 20–8 Push Descriptor in the Target Type Metadata File For Triggering a Severity

```

<PushDescriptor RECVLET_ID="SNMPTrap">
  <Property NAME="MatchTrapOID" SCOPE="GLOBAL">1.3.6.1.4.1.529.50</Property>
  <Property NAME="MatchAgentAddr" SCOPE="INSTANCE">AdminAddress</Property>
  <Property NAME="SeverityStatusOID"
SCOPE="GLOBAL">1.3.6.1.2.1.2.2.1.8</Property>
  <Property NAME="KeyNameOID" SCOPE="GLOBAL">1.3.6.1.2.1.2.2.1.31</Property>
  <Property NAME="ContextTypeOID" SCOPE="GLOBAL">1.3.6.1.2.1.2.2.1.3</Property>
  <Property NAME="ContextConfigured_statusOID"
SCOPE="GLOBAL">1.3.6.1.2.1.2.2.1.7</Property>
  <Property NAME="SeverityCode" SCOPE="GLOBAL">CRITICAL</Property>
  <Property NAME="hostname" SCOPE="INSTANCE">snmpHost</Property>
  <Property NAME="PORT" SCOPE="INSTANCE">snmpPort</Property>
  <CredentialRef NAME="monCreds">monCredentials</CredentialRef>
</PushDescriptor>

```

[Example 20–9](#) shows how the push descriptor can be defined in the target type metadata file to trigger a datapoint, which would specify the reporting of data on the same trap, with *ifName* as the key-column and the other three as data columns.

Example 20–9 Push Descriptor in the Target Type Metadata File For Triggering a Datapoint

```

<PushDescriptor RECVLET_ID="SNMPTrap">
  <Property NAME="MatchEnterprise" SCOPE="GLOBAL">1.3.6.1.4.1.529</Property>

```

```

<Property NAME="MatchGenericTrap" SCOPE="GLOBAL">6</Property>
<Property NAME="MatchSpecificTrap" SCOPE="GLOBAL">50</Property>
<Property NAME="MatchAgentAddr" SCOPE="INSTANCE">AdminAddress</Property>
<Property NAME="KeyNameOID" SCOPE="GLOBAL">1.3.6.1.2.1.2.2.1.31</Property>
<Property NAME="DataStatusOID" SCOPE="GLOBAL">1.3.6.1.2.1.2.2.1.8</Property>
<Property NAME="DataTypeOID" SCOPE="GLOBAL">1.3.6.1.2.1.2.2.1.3</Property>
<Property NAME="DataConfigured_statusOID"
SCOPE="GLOBAL">1.3.6.1.2.1.2.2.1.7</Property>
<Property NAME="hostname" SCOPE="INSTANCE">snmpHost</Property>
<Property NAME="PORT" SCOPE="INSTANCE">snmpPort</Property>
<CredentialRef NAME="monCreds">monCredentials</CredentialRef>
</PushDescriptor>

```

Example 20–10 shows how the monCredentials is defined in target type metadata file

Example 20–10 monCredentials in the Target Type Metadata File

```

<CredentialInfo>
  <CredentialSet NAME="monCredentials" USAGE="MONITORING">
    <AllowedCredType TYPE="SNMPV1Creds" />
    <AllowedCredType TYPE="SNMPV3Creds" />
  </CredentialSet>
</CredentialInfo>

```

SNMPV1Creds or SNMPV3Creds values for the respective targets must be set from the Cloud Control console by selecting **Setup**, then **Security**, and then **Monitoring Credentials**.

To receive an SNMPV1 trap or SNMPV2 Notification, the user must choose SNMPV1Creds. Choosing SNMPV1Creds will ask for “Community String” parameter value. Appropriate community string values need to be set by the user.

To receive an SNMPV3 Notification, the user must choose SNMPV3Creds. Choosing SNMPV3Creds will ask for “UserName”, “Auth Password”, “Auth Protocol” and “Privacy Password” parameter values. The user must set the required values according to the secLevel they want to use.

Example 20–11 SNMPV1 Trap Received by the Management Agent

```

Message:
version: 0
community: 'public'
Trap-PDU:
enterprise: enterprises.ascend (1.3.6.1.4.1.529)
agent-addr: 138.2.204.10
generic-trap: 6
specific-trap: 50
time-stamp: <timestamp from router's sysUptime>
variable-bindings:
Name: ifIndex.<x> (1.3.6.1.2.1.2.2.1.1.<x>)
Type: INTEGER
Value: <x>
Name: ifAdminStatus.<x> (1.3.6.1.2.1.2.2.1.7.<x>)
Type: INTEGER
Value: up (1)
Name: ifOperStatus.<x> (1.3.6.1.2.1.2.2.1.8.<x>)
Type: INTEGER
Value: down (2)
Name: ifType.<x> (1.3.6.1.2.1.2.2.1.3.<x>)
Type: INTEGER
Value: iso88023-csmacd (7)

```

```
Name: ifName.<x> (1.3.6.1.2.1.2.2.1.31.<x>)
Type: DisplayString
Value: 'eth0'
```

Example 20–12 SNMV2/SNMPV3 Notification Received by the Management Agent

```
Message: Received from address 138.2.204.10 (ip address from UDP layer)
version: 1 (or 3 i.e 1 for snmpv2 and 3 for snmpv3 notification)
Security params : (community if SNMPV2, SNMPV3 security params if SNMPV3)
Trap-PDU:
variable-bindings:
Name: sysUpTime.0 (.1.3.6.1.2.1.1.3.0)
Type: INTEGER
Value: 43053404
Name: snmpTrapOID.0 (1.3.6.1.6.3.1.1.4.1.0)
Type: Object Identifier
Value: 1.3.6.1.4.1.529.50
Name: ifAdminStatus.<x> (1.3.6.1.2.1.2.2.1.7.<x>)
Type: INTEGER
Value: up (1)
Name: ifOperStatus.<x> (1.3.6.1.2.1.2.2.1.8.<x>)
Type: INTEGER
Value: down (2)
Name: ifType.<x> (1.3.6.1.2.1.2.2.1.3.<x>)
Type: INTEGER
Value: iso88023-csmacd (7)
Name: ifName.<x> (1.3.6.1.2.1.2.2.1.31.<x>)
Type: DisplayString
Value: 'eth0'
```

Notes

- The target type metadata file can have multiple metrics with push descriptor definitions. Also, a single metric can have multiple push descriptors attached.
- For event format, ensure that your push descriptor defines only one Event* parameter indicating one metric column as described in [Table 20–1](#). The event push descriptor can have one or more Context* parameters to indicate additional metric column values to send as part of an AlertContext.
- For datapoint format, ensure that your push descriptor defines one or more Data* properties as described in [Table 20–1](#) and demonstrated in [Example 20–5](#) and [Example 20–9](#).

Using Fetchlets

This chapter contains the following sections:

- [About Fetchlets](#)
- [OS Command Fetchlets](#)
- [SQL Fetchlet](#)
- [SNMP Fetchlet](#)
- [HTTP Data Fetchlets](#)
- [URLXML Fetchlet](#)
- [URL Timing Fetchlet](#)
- [Dynamic Monitoring Service \(DMS\) Fetchlet](#)
- [JDBC Fetchlet](#)
- [WBEM Fetchlet](#)
- [JMX Fetchlet](#)
- [Web Services Fetchlet](#)
- [WS-Management Fetchlet](#)
- [REST Fetchlet](#)

21.1 About Fetchlets

Enterprise Manager data retrieval is handled through predefined "fetchlets." A fetchlet is a parametrized data access mechanism that takes arguments (for example, a script, a SQL statement, a target instance's properties) as input and returns formatted data. Each fetchlet handles a specific type of data access. The fetchlets supplied with Enterprise Manager provide data retrieval capability for the most common data access methods, such as SQL, SNMP (Simple Network Management Protocol), HTTP, and DMS (Dynamic Monitoring Service). To handle more complex data access requirements, Enterprise Manager also provides an OS command fetchlet that allows developers to implement custom metric collection methods.

The following sections describe the fetchlets supplied with Enterprise Manager:

21.2 OS Command Fetchlets

The operating system (OS) command fetchlets allow you to obtain metric data by executing OS commands (either individually or from scripts) that return a standard out (stdout) data stream.

Three OS command fetchlets are available:

- OS Fetchlet (raw)
- OSLines Fetchlet (split into lines)
- OSLineToken Fetchlet (tokenized lines)

21.2.1 OS Fetchlet

The OS fetchlet executes a given OS command and returns the command's output in a single cell table.

Input Parameters

Table 21–1 OS Fetchlet Input Parameters

Parameter	Type	Description	Use
command	string	Operating system command to be executed.	Required
ENV <i>name</i>	string	OSFetchlet parameters starting with "ENV" appear in the execution environment for the command as <i>name</i> environment variables	Zero or more
errStartsWith	String	When defined, this property allows you to define a custom prefix for error messages. If this property is not defined, the OSFetchlet defaults to "em_error=" as the message prefix.	Optional
script	string	Specifies the script to be executed if <i>command</i> property only provides an interpreter. For example, <i>command</i> might consist of "perl." <i>script</i> is then used to specify the particular perl script to be run. Although scripts can be specified directly from the <i>command</i> parameter, using the <i>script</i> parameter adds to stylistic clarity and readability when defining a target type metadata file.	Optional
args	string	A property that is taken as one or more arguments to the command and script properties. Although args can be specified directly from the command parameter, using the args parameter adds to stylistic clarity and readability when defining a target type metadata file.	Optional
separateErrorStream	boolean	If an error occurs while executing a <i>command</i> , this property instructs the fetchlet whether to return both the stdout and stderr to the user as an error message. When set to TRUE, only stderr output is sent to the user as an error message when there is a <i>command</i> error. When set to FALSE (default value), both the stdout and the stderr are sent to the user as an error message upon <i>command</i> failure.	Optional. (TRUE/FALSE)

Table 21–1 (Cont.) OS Fetchlet Input Parameters

Parameter	Type	Description	Use
em_metric_timeout	integer	Metric timeout period (in seconds). After the timeout period has finished, the Management Agent returns a timeout exception and terminates any child processes that may have been created. The Management Agent <i>does not</i> terminate any of the grandchild processes. Specify "-1" for no timeout period.	Optional

Example

You want to obtain metric data by executing the UNIX *echo* command.

To run the command from the shell environment, enter:

```
echo Line 1|some more|even more\nLine 2\n\nLine 4|a little more|\n|Line 5\n|Line 6|\n|Line 7|again|\nLine 8|the|end
```

The *echo* command produces the following standard output:

```
Line 1|some more|even more
Line 2
```

```
Line 4|a little more|
|Line 5
|Line 6|
|Line 7|again|
Line 8|the|end
```

Using the OS fetchlet with the given example command.

The fetchlet returns the following 1 x 1 table:

Figure 21–1 Table Returned by the OS Fetchlet

```
Line 1|some more|even more
Line 2

Line 4|a little more|
|Line 5
|Line 6|
|Line 7|again|
Line 8|the|end
```

The raw output of the OS command is returned. Any standard error output is appended to the standard output.

Error Handling

Any problems encountered launching the *command* (For example, the *command* program no longer exists) results in an `oracle.sysman.emSDK.emd.fetchlet.MetricSourceException` wrapping a `java.io.IOException`. If the command exits with a non-zero exit value, the fetchlet throws an `oracle.sysman.emSDK.emd.fetchlet.MetricSourceException` wrapping an `oracle.sysman.emd.fetchlets.CommandFailedException`.

Notes

Commands are *not* executed as if they are being run in a shell. This means that common shell symbols do not work, including piping, output redirection, and backgrounding.

Commands cannot read from standard input.

The fetchlet blocks and waits for the command to finish.

21.2.2 OSLines Fetchlet (split into lines)

The OS Lines fetchlet executes a given OS command and tokenizes the OS command's output. The output is tokenized by lines. The fetchlet returns the tokens in a single column table. The *n*th row in the table represents the *n*th line in the output of the OS command.

To get the raw, untokenized output of an OS command, use the OS fetchlet. To get the output of an OS command tokenized by lines and each line tokenized by a given delimiter, see the OS Line Token fetchlet.

Input Parameters

Table 21–2 OSLines Fetchlet Input Parameters

Parameter	Type	Description	Use
command	string	Operating system command to be executed.	Required
startsWith	string	Only lines starting with this string are included in the result.	Optional; default = "" (all lines are included)
ENVname	string	Parameters starting with "ENV" appear in the execution environment for the command as <i>name</i> environment variables	Zero or more of these
errStartsWith	string	When defined, this property allows you to define a custom prefix for error messages. If this property is not defined, the OSFetchlet defaults to "em_error=" as the message prefix.	Optional
script	string	Specifies the script to be executed if <i>command</i> property only provides an interpreter. For example, <i>command</i> might consist of "perl." <i>script</i> is then used to specify the particular perl script to be run. Although scripts can be specified directly from the <i>command</i> parameter, using the <i>script</i> parameter adds to stylistic clarity and readability when defining a target type metadata file.	Optional
args	string	A property that is taken as one or more arguments to the <i>command</i> and <i>script</i> properties. Although args can be specified directly from the <i>command</i> parameter, using the <i>args</i> parameter adds to stylistic clarity and readability when defining a target type metadata file.	Optional

Table 21–2 (Cont.) OSLines Fetchlet Input Parameters

Parameter	Type	Description	Use
separateErrorStream	boolean	If an error occurs while executing a command, this property instructs the fetchlet to whether to return both the stdout and stderr to the user as an error message. When set to TRUE, only stderr output is sent to the user as an error message when there is a command error. When set to FALSE (default value), both the stdout and the stderr are sent to the user as an error message upon command failure.	Optional. (TRUE/FALSE)
em_metric_timeout	integer	Metric timeout period (in seconds). After the timeout period has finished, the Management Agent returns a timeout exception and terminates any child processes that may have been created. The Management Agent DOES NOT kill any of the grandchild processes. Specify "-1" for no timeout period.	Optional

Example

Take the following UNIX command:

```
echo Line 1|some more|even more\nLine 2\n\nLine 4|a little more|\n|Line 5\n|Line 6|\n|Line 7|again|\nLine 8|the|end
```

It produces the following output:

```
Line 1|some more|even more
Line 2
```

```
Line 4|a little more|
|Line 5
|Line 6|
|Line 7|again|
Line 8|the|end
```

Running OSLinesFetchlet with the given example command produces the following single column table.

Figure 21–2 Table Returned by the OS LINES Fetchlet

Line 1 some more even more
Line 2
Line 4 a little more
Line 5
Line 6
Line 7 again
Line 8 the end

Note that without content, "\n" results in a blank line inserted between Line 2 and Line 4.

Notes: Commands are *not* executed as if they are being run in a shell. This means that common shell symbols do not work, including piping, output redirection, and backgrounding.

Commands cannot read from standard input.

The fetchlet blocks and waits for the command to finish.

The standard output of the command is captured and the standard error is captured and appended to the standard output.

Lines are tokenized using "\n".

21.2.3 OSLineToken Fetchlet (tokenized lines)

The OS Line Token fetchlet executes a given OS command and tokenizes the output of the OS command. The output is tokenized first by lines, and then each line is tokenized by a given delimiter set. The fetchlet returns the tokens in a table. The n th row in the table represents the n th line in the output of the OS command. The n th column in the table represents the n th token in a line as determined by the given delimiter set.

To get the raw, untokenized output of an OS command, see [Section 21.2.1, "OS Fetchlet"](#).

Input Parameters

Table 21–3 OSLineToken Fetchlet Input Parameters

Parameter	Type	Description	Use
command	String	Operating system command to be executed.	Required
delimiter	String	Set of characters that act as delimiters to tokenize the lines	Optional; default = "" (just breaks output into lines)
startsWith	String	Only lines starting with this string are included in the result	Optional; default = "" (all lines are included)
ENV $name$	String	Parameters starting with "ENV" appear in the execution environment for the command as $name$ environment variables	Zero or more of these
errStartsWith	String	When defined, this property allows you to define a custom prefix for error messages. If this property is not defined, the OSFetchlet defaults to "em_error=" as the message prefix.	Optional
script	String	Specifies the script to be executed if command property only provides an interpreter. For example, <i>command</i> might consist of "perl." The script is then used to specify the particular perl script to be run. Although scripts can be specified directly from the command parameter, using the script parameter adds to stylistic clarity and readability when defining a target type metadata file.	Optional

Table 21-3 (Cont.) OSLineToken Fetchlet Input Parameters

Parameter	Type	Description	Use
args	String	A property that is taken as one or more arguments to the <i>command</i> and <i>script</i> properties. Although args can be specified directly from the <i>command</i> parameter, using the args parameter adds to stylistic clarity and readability when defining a target type metadata file.	Optional
separateErrorStream	Boolean	If an error occurs while executing a <i>command</i> , this property instructs the fetchlet to whether to return both the stdout and stderr to the user as an error message. When set to TRUE, only stderr output is sent to the user as an error message when there is a <i>command</i> error. When set to FALSE (default value), both the stdout and the stderr are sent to the user as an error message upon <i>command</i> failure.	Optional. (TRUE/FALSE)
em_metric_timeout	Integer	Metric timeout period (in seconds). After the timeout period has finished, the Management Agent returns a timeout exception and terminates any child processes that may have been created. The Management Agent DOES NOT kill any of the grandchild processes. Specify "-1" for no timeout period.	Optional

Example

Take the following UNIX command:

```
echo Line 1|some more|even more\nLine 2\n\nLine 4|a little more|\n|Line 5\n|Line 6|\n|Line 7|again|\nLine 8|the|end
```

It produces the following output:

```
Line 1|some more|even more
Line 2
```

```
Line 4|a little more|
|Line 5
|Line 6|
|Line 7|again|
Line 8|the|end
```

Running OSLineTokenFetchlet with the given example command and a single character "|" for the delimiter generates the following table:

Figure 21–3 Table Returned by the OS Token Lines Fetchlet

Line 1	some more	even more
Line 2		
Line 4	a little more	
Line 5		
Line 6		
Line 7	again	
Line 8	the	end

Error Handling

Any problem launching the command (unable to find the command program) results in an `oracle.sysman.emSDK.emd.fetchlet.MetricSourceException` wrapping a `java.io.IOException`.

If the command exits with a non-zero exit value, the fetchlet throws a `oracle.sysman.emSDK.emd.fetchlet.MetricSourceException` wrapping a `oracle.sysman.emd.fetchlets.CommandFailedException`.

Notes

Commands are *not* executed as if they are being run in a shell. This means that common shell symbols do not work, including piping, output redirection, and backgrounding.

The fetchlet promptly closes the input stream to the running command.

The fetchlet blocks and waits for the command to finish.

Lines are tokenized using `"\n"`.

The delimiter can be a single character or a set of characters. For example, it can be `"|+_"`, if the line should be broken up by pipes, pluses, and underscores. If two or more delimiters are together in the output text, such as `"| |"` or `"+|+"`, then it is as if there are empty string tokens between them. These empty strings get columns in the result table. It is *not* considered that there are empty strings preceding a delimiter that starts a line or following a delimiter that ends a line.

In order to express non-printable characters in the delimiter set (such as tabs) in XML, use `"&#xHH;"` where H is the hexadecimal identifier for the character.

21.2.4 Invoke an OS Fetchlet as a Specific User for Metric Collection

Depending on requirements, your plug-in may need to utilize the OS fetchlet to invoke a pre-existing script on the Management Agent monitoring a target to collect data for a specific metric as a specific user; that is, as a user other than the default "oracle" user.

Enterprise Manager supports the use of Privilege Delegation Providers (sudo and powerbroker) to invoke metric collections as a specific user. Enabling PDP for a plug-in requires credential setup on both the plug-in and on hosts where the target(s) being monitored are deployed.

In your plug-in, you must set the credential reference in the metric definition in the target metadata file. In the example, example Credentialref line has "your_cred". This value refers to monitoring credential set name.

Example 21–1 Example of Credential Reference in Target Metadata

```

<TargetMetadata TYPE="my_type" NAME="my_target_name">
  ...
  <Metric NAME="my_special_metric" TYPE="TABLE">
    <TableDescriptor>
      <ColumnDescriptor NAME="test" TYPE="STRING"/>
    </TableDescriptor>
    <QueryDescriptor FETCHLET_ID="OS">
      <Property NAME="command" SCOPE="GLOBAL">%perlBin%/perl</Property>
      <Property NAME="script" SCOPE="GLOBAL">%scriptsDir%/your_
        script.pl</Property>
      <CredentialRef NAME="OSCreds">your_cred</CredentialRef>
    </QueryDescriptor>
  </Metric>
  ...
</TargetMetadata>

```

On the Management Agent monitoring the target, a referenced credential type must be created that points to host:HostCreds, and allow the monitoring credential set be of the new type that you add. See "Sudo and PowerBroker Support" in the *Enterprise Manager Cloud Control Administrator's Guide* for details on privilege delegation setup using Enterprise Manager Cloud Control.

The credential data will be persisted to the target metadata file (target.xml) for the Management Agent monitoring the target.

The below example defines the referenced credential type MyHostCreds in target.xml, which is of the credential type host:HostCreds.

Example 21–2 Credential Type Definition in Management Agent

```

<Target TYPE="<removed>" NAME="<removed>" DISPLAY_NAME="<removed>" ON_HOST="" EMD_
URL="https://<removed>/emd/main/" TIMEZONE_REGION="" IDENTIFIER="TARGET_
GUID=<removed>">
  ...
  <CredentialType NAME="MyHostCreds"><CredentialTypeRef REF_NAME="HostRef" REF_
TYPE="HostCreds" REF_TARGETTYPE="host" ASSOCIATION="host"><CredentialTypeRefColumn
NAME="HostUserName" REF_TYPECOLUMN="HostUserName"/><CredentialTypeRefColumn
NAME="HostPassword" REF_
TYPECOLUMN="HostPassword"/></CredentialTypeRef></CredentialType>
  ...
</Target>

```

When monitoring credentials are updated Cloud Control (via Setup->Security->Monitoring Credentials), the data shown above will be updated on the Management Agent automatically.

The next example defines the HostMonCredSet monitoring credential set, which is of type MyHostCreds (and therefore type host:HostCreds)

Example 21–3

```

<CredentialSet NAME="HostMonCredSet" CREDENTIAL_TYPE="MyHostCreds"
USAGE="MONITORING"><AllowedCredType TYPE="MyHostCreds"/>
</CredentialSet>

```

21.3 SQL Fetchlet

The SQL fetchlet executes a given SQL statement on a given database as a given user and returns the table result.

Input Parameters

Table 21–4 SQL Fetchlet Input Parameters

Parameter	Type	Description	Use
Connection Information			
MachineName	String	Database host	Required
Port	Integer	Database port	Required
SID	String	Database SID	Required unless ServiceName is specified
ServiceName	String	Database ServiceName	Required unless SID is specified
OracleHome	String	Database's Oracle Home (used in conjunction with OidRepSchemaName).	Required when OidRepSchemaName is used.
Credential Information			
UserName	String	user name	Required
password	String	user password	Optional; default is ""
Role	String	Role used when connecting to the database (e.g., SYSDBA)	Optional; allowed choices are SYSDBA, SYSOPER, and NORMAL (default)
General			
STATEMENT	String	SQL statement or PL/SQL block	Required unless FILENAME is specified.
FILENAME	String	Full path of the file containing the SQL statement or PL/SQL block	Required unless STATEMENT is specified.
NUMROWS	Integer	Maximum number of rows to output.	Required
Bind Parameters			
SQLINPARAM<position>	String	Value of input parameter at position <position>	Zero or more, one for each input parameter.
SQLOUTPARAM<position>	Integer	Position of output parameter, if it exists	There can be exactly one output parameter, if it exists.

Table 21–4 (Cont.) SQL Fetchlet Input Parameters

Parameter	Type	Description	Use
SQLOUTPARAMTYPE	String	Type of the output parameter, if it exists.	There can be exactly one output parameter type, if it exists.
transpose	TRUE/ FALSE	If TRUE, the result is transposed: rows to columns and columns to rows.	

Notes

The SQL statement or PL/SQL block can be specified either through the STATEMENT property, or via a file whose name is provided through the FILENAME property.

The SQL fetchlet supports input parameters. Input and output parameters are indicated in the SQL/PLSQL text in the usual way, by using ":<number>". Input parameters can be used to bind values to both SQL queries and PL/SQL blocks.

Input parameter values are specified as properties of the form SQLINPARAM<position>. There can be any number of input parameters. The input parameters need to be scalar: input parameters of type ARRAY and STRUCT are not supported.

The SQL fetchlet supports the execution of anonymous PL/SQL blocks (which may call other functions or procedures) to retrieve data. When executing a block of PL/SQL, data is returned to the fetchlet by means of an OUT parameter. There can be exactly one out parameter. It must be of type SQL_CURSOR (a PL/SQL REF CURSOR), or it must be a named type that represents an array of objects. In the latter case, each field of the object represents one column of the table; and each object instance in the array represents one complete row in the table. The OUT parameter position and type are specified by means of the properties SQLOUTPARAMPOS and SQLOUTPARAMTYPE. If an OUT parameter is specified, then the fetchlet assumes it is executing PL/SQL and treats the STATEMENT property as an anonymous PL/SQL block.

Note: When using a SQLOUTPARAMTYPE of type 'ARRAY', you must identify the array as follows:

- If you create the array type specified in the SQLOUTPARAMTYPE from SQL*Plus or any utility *without* using double quotation marks to surround the array name, then you must specify the array name using all upper-case letters in the target metadata file for this property. The reason for this because the RDBMS automatically changes the array name to all upper-case.
 - If you create the array type specified in the SQLOUTPARAMTYPE from SQL*Plus or any utility using double-quotation marks to surround the array name, then the RDBMS retains the case specified. For this reason, users must specify the array name using the same case used in the target metadata file.
-

If no OUT parameter is specified, the fetchlet assumes that it is executing a SQL query.

Note that all input parameters to the SQL fetchlet are strings. This means that all other datatypes will have to be converted to strings. This is straightforward for datatypes such as numbers, but not, for example, dates and timestamps. You can pass an absolute date or timestamp by passing a character representation of the value (using a `DateFormat` class). There is no way currently to pass in a date function, such as `SYSDATE` or `SYSDATE+1`. In such case, you could embed the date argument directly in the SQL, for example:

```
begin func1(:1, :2, SYSDATE); end;
```

The other caveat is passing null arguments to a procedure. Consider the following SQL:

```
STATEMENT=begin func1(:1,:2); end;  
SQLINPARAM1=null  
SQLOUTPARAMPOS=2  
SQLOUTPARAMTYPE=fooret
```

Assume that the first argument is intended to be a `varchar2`. By parameterizing it and passing 'null' as the first argument, what we are really doing is passing the **string** 'null' to the argument, and not a null value. If you intend to pass a null value, do the following:

```
STATEMENT=begin func1(null, :1); end;  
SQLOUTPARAMPOS=1  
SQLOUTPARAMTYPE=fooret
```

Examples

The following properties execute a query (get all users) with no parameters:

Example 21–4 Query With No Parameters

```
MachineName=skini-pc  
Port=1521  
SID=o817  
UserName=scott  
password=tiger  
STATEMENT=select * from all_users;  
NUMROWS=30
```

The following properties execute a query (get the first few objects of a specified type owned by a specified user) with input parameters:

Example 21–5 Query With Input Parameters

```
MachineName=skini-pc  
Port=1521  
SID=o817  
UserName=scott  
password=tiger  
STATEMENT=select * from all_objects where owner=:1 and object_type=:2 and  
rownum<:3ttt  
NUMROWS=30  
SQLINPARAM1=SYSTEM  
SQLINPARAM2=INDEX  
SQLINPARAM3=10
```

[Example 21–6](#) executes a PL/SQL procedure that returns a cursor and has input parameters:

Example 21–6 PL/SQL Procedure With Input Parameters

```

achineName=skini-pc
Port=1521
SID=o817
UserName=scott
password=tiger
STATEMENT=begin :1 := skini_junk.func1(:2); end;
NUMROWS=30
SQLINPARAM2=SYSTEM
SQLOUTPARAMPOS=1
SQLOUTPARAMTYPE=sql_cursor

```

[Example 21–7](#) specifies a PL/SQL procedure that returns an array of strings:

Example 21–7 PL/SQL Procedure Returning an Array of Strings

```

MachineName=skini-pc
Port=1521
SID=o817
UserName=scott
password=tiger
STATEMENT=begin skini_junk.newproc(:1,:2); end;
NUMROWS=30
SQLINPARAM1=SYSTEM
SQLOUTPARAMPOS=2
SQLOUTPARAMTYPE=my_string_array

```

[Example 21–8](#) specifies a PL/SQL package that returns an array of structures:

Example 21–8 PL/SQL Package Returning an Array of Structures

```

MachineName=skini-pc
Port=1521
SID=o817
UserName=scott
password=tiger
STATEMENT=begin :1 := skini_junk.func2(:2,:3,:4,:5,:6); end;
NUMROWS=30
SQLINPARAM2=somename
SQLINPARAM3=someplace
SQLINPARAM4=someanimal
SQLINPARAM5=something
SQLINPARAM6=22
SQLOUTPARAMPOS=1
SQLOUTPARAMTYPE=my_struct_array

```

[Example 21–9](#) provides the PL/SQL used in the previous examples for reference.

Example 21–9 PL/SQL Used in Examples

```

create or replace type my_type as Object (
    name varchar2(128),
    place varchar2(128),
    animal integer,
    thing number,
    thing2 number);
/
create or replace type my_struct_array as table of my_type;
/

```

```
create or replace type my_string_array as table of varchar2(3000);
/

create or replace type my_int_array as table of integer;
/

create or replace package skini_junk as

type Jcr is ref cursor;

function func1(username in varchar2) return Jcr;
function func2(name varchar2, place varchar2, animal integer,
               thing number, thing2 number) return my_struct_array;
procedure newproc(name varchar2, outArray OUT my_string_array);
procedure newproc2(numrows in varchar2, outArray OUT my_int_array);

end skini_junk;
/

create or replace package body skini_junk as

function func1(username in varchar2) return Jcr is
cr Jcr;
begin
    open cr for select object_name, object_type, status from all_objects where
                owner=upper(username);

    return cr;
end;

function func2(name varchar2, place varchar2, animal integer,
               thing number, thing2 number) return my_struct_array IS
ret my_struct_array := my_struct_array();

begin
    ret.extend(50);

    for i in 1..50 loop
        ret(i) := my_type(name || i,
                          place || i,
                          animal+i,
                          thing+i,
                          thing2+i);
    end loop;
    return ret;
end;

procedure newproc(name varchar2, outArray OUT my_string_array) IS
begin
    outArray := my_string_array();
    outArray.extend(100);

    for i in 1..100 loop
        outArray(i) := name || i;
    end loop;
end;

procedure newproc2(numrows in varchar2, outArray OUT my_int_array) IS
begin
```

```

        outArray := my_int_array();
        outArray.extend(numrows);
        for i in 1..numrows loop
            outArray(i) := i;
        end loop;
    end;

    end skini_junk;
/

```

21.4 SNMP Fetchlet

An *object identifier* (OID) corresponds to either a MIB variable instance or a MIB variable with multiple instances. Given a list of OIDs, the SNMP fetchlet polls an *SNMP agent* on a given host for corresponding instances.

Input Parameters

Table 21–5 *SNMP Fetchlet Input Parameters*

Parameter	Type	Description	Use
hostname	String	Host name of the SNMP agent	Required. Default is "localhost" Examples: "bigip.host.example.com" "148.87.10.5"
PORT	String	Port of the SNMP agent	Optional. Default is "161"
COMMUNITY	String	SNMP community string	Optional. Default is "public"
TIMEOUT	String	SNMP timeout.	Optional. Default is five seconds
OIDS	String	A list of substrings separated by delimiters. Each substring starts with an OID (in numerical dot notation), and can be optionally ended with *PlacementOID. (See notes for details.)	Required. Examples: "1.3.6.1.2.1.2.1.1.1.0,1.3.6.1.2.1.2.1.1.3.0,1.3.6.1.2.1.2.1.1.5.0" "1.3.6.1.2.1.2.1.2.2.1.2" 1.3.6.1.2.1.2.1.2.2.1.10 1.3.6.1.2.1.2.1.2.2.1.16" "1.3.6.1.2.1.2.2.1.3 1.3.6.1.2.1.2.2.1.5" 1.3.6.1.2.1.4.20.1.1*1.3.6.1.2.1.4.20.1.2 1.3.6.1.2.1.4.20.1.3*1.3.6.1.2.1.4.20.1.2"
DELIM	String	A delimiter to separate individual substrings in OIDS.	Optional; default is whitespace characters, (dot), *(star) and 0-9 (digits) cannot be used as delimiters
TABLE	String	Each OID in OIDS corresponds to a variable with multiple instances if this parameter is "TRUE" and to a single variable instance if it is "FALSE".	Optional; default is "FALSE"

Table 21–5 (Cont.) SNMP Fetchlet Input Parameters

Parameter	Type	Description	Use
PINGMODE	Boolean	<p>Used for defining PINGMODE Response metric</p> <p>If set to TRUE, then a successful GetResponse generates a single-row, single-column table with the value "1" in its cell. A timeout generates a single-row, single-column table with the value "0".</p> <p>This is useful for defining a Response metric for an SNMP-based target.</p>	Optional. Default is "FALSE"
IGNORE_TIMEOUT_ERR_BOOLEAN	Boolean	<p>Specifies whether to generate an error when a non-PINGMODE invocation times out while waiting for a response.</p> <p>If set to TRUE, then a non-PINGMODE invocation that times out while waiting for a response should <i>not</i> generate a metric collection error. This is reasonable behavior for targets that define a PINGMODE Response metric. If that Response metric is going to generate an availability severity when the SNMP agent stops responding, then there is no need to generate metric errors on any non-Response metrics that happen to run before the Response metric can identify the problem.</p>	Optional. Default is "TRUE"
MAX_NUM_ROWS_FETCH	Integer	<p>The maximum number of rows to be returned by a TABLE invocation.</p> <p>The configuration property "SnmptableMaxNumRowsFetch" can override the default value.</p>	Optional. Default is 1000
CONTEXT_NAME	String	<p>Along with CONTEXT_ENGINE_ID, these two properties specify a set of SNMPv3 credentials, which replace the community string used by SNMPv1 and SNMPv2c.</p> <p>Note: If these two properties are specified, then the COMMUNITY and the VERSION parameters are ignored, and the sent request is an SNMPv3 request.</p>	Optional. No default value
CONTEXT_ENGINE_ID	String	For information about this parameter, see the CONTEXT_NAME description.	Optional. No default value

Table 21–5 (Cont.) SNMP Fetchlet Input Parameters

Parameter	Type	Description	Use
VERSION	String	<p>Specifies the SNMP version.</p> <p>If the following occurs, then VERSION is set to "v2c", indicating an SNMPv2c request:</p> <ul style="list-style-type: none"> disallowV1Requests is set to TRUE or hasV2Types is set to TRUE and CONTEXT_NAME and CONTEXT_ENGINE_ID are not specified <p>If these previous conditions do not apply, then VERSION is set to "v1" indicating an SNMPv1 request.</p>	Optional. Default is "v1"
disallowV1Requests	Boolean	This parameter enables the user to specify that the Management Agent should use SNMPv2c only when sending any request to a particular target	Optional. Default is "FALSE"
hasV2Types	Boolean	This parameter is a global-scoped property for an SNMP QueryDescriptor that includes OIDs for MIB variables whose types are 64-bit integer values. These are not representable in SNMPv1. Even if other requests for the same target instance are sent using SNMPv1, the target-type owner knows that this request must be SNMPv2c.	Optional. Default is "FALSE"
USE_GET_NEXT_ONLY	Boolean	<p>If an SNMP QueryDescriptor is SNMPv2c, according to the conditions described in the VERSION description, and if TABLE is TRUE, then the multiple rows that the SNMP fetchlet returns will be fetched using the SNMPv2c GetBulk request, and not the GetNext request used in SNMPv1.</p> <p>If USE_GET_NEXT_ONLY is set to TRUE, then the SNMP fetchlet returns will be fetched using GetNext requests.</p>	Optional. Default is "FALSE"

Error Handling

MissingParameterException is thrown if either host name or OIDs is not given. FetchletException is thrown if TABLE is not equal to either TRUE or FALSE, an I/O error occurs while sending or receiving SNMP messages to or from the agent, or the agent responds with an SNMP error.

Notes

The table returned by the fetchlet contains a column for every OID in OIDS. If input OIDs correspond to single variable instances, the table will have just one row with those instances. On the other hand, if the OIDs correspond to variables with multiple instances, each column in the table will contain instances for its OID and each row will correspond to a different *subidentifier*. (A subidentifier is an OID extension that uniquely identifies a particular variable instance for some MIB variable.) OIDS must contain either all OIDs with subidentifiers or all OIDs without the subidentifiers.

For example, to request the variable instances for the three OIDs: `sysDescr`, `sysUpTime`, and `sysName`, OIDS would have to be "1.3.6.1.2.1.2.1.1.0 1.3.6.1.2.1.2.1.1.3.0 1.3.6.1.2.1.2.1.1.5.0". In this case, all OIDs contain the instance subidentifier, ".0". The return table appears as follows (the actual values may be different):

Figure 21–4 SNMP Fetchlet

Sun SNMP Agent, Ultra-4	32504340	necdc-view3
-------------------------	----------	-------------

This figure shows the output of the SNMP fetchlet

Alternatively, assume that some MIB contains the following 3 columns and 4 instances:

Figure 21–5 SNMP Fetchlet: Columns 3 and 4 Content

ifDescr (network interface description)	ifInOctets (bytes into an interface)	ifOutOctets (bytes out of an interface)
OID: 1.3.6.1.2.1.2.2.1.2	OID: 1.3.6.1.2.1.2.2.1.10	OID: 1.3.6.1.2.1.2.2.1.16
subidentifier : variable instance	subidentifier : variable instance	subidentifier : variable instance
1: wx1	1: 26150844	1: 29368527
2: wx2	2: 2763185941	2: 3023812977
3: wx3	3: 123615396	3: 2055140730
4: wx4	4: 2068257723	4: 3212899913

To construct a table with 3 columns corresponding to `ifDescr`, `ifInOctets`, and `ifOutOctets`, OIDS would be defined as follows

"1.3.6.1.2.1.2.2.1.2 1.3.6.1.2.1.2.2.1.10 1.3.6.1.2.1.2.2.1.16"

The fetchlet returns the following:

Figure 21–6 SNMP Fetchlet: ifDescr, ifInOctets, and ifOutOctets OIDS

wx1	26150844	29368527
wx2	2763185941	3023812977
wx3	123615396	2055140730
wx4	2068257723	3212899913

The rows correspond to subidentifiers 1,2,3,4 respectively.

Any OID in OIDS can be appended with another *placement* OID. The variable instances for the placement OID do not appear in the returned table. Instead, they determine the place for the variable instances of the original OID within a column. In particular, for every variable instance I with subidentifier S in the set of instances for the original OID, (a) there must exist a variable instance X with subidentifier S in the set of instances corresponding to the placement OID, and (b) X is used as the subidentifier for the instance I.

For example, consider a MIB containing the following 3 columns, each with 4 variable instances:

Figure 21–7 SNMP Fetchlet: MIB Content with 4 Variable Instances

ifDescr (network interface description)	ipAdEntNetMask (netmask)	ipAdEntIfIndex (network interface index)
OID: 1.3.6.1.2.1.2.1.2.2.1.2	OID: 1.3.6.1.2.1.2.1.4.20.1.3	OID: 1.3.6.1.2.1.2.1.4.20.1.2
subidentifier : variable instance	subidentifier : variable instance	subidentifier : variable instance
1: wx1	IP1: 255.255.255.0	IP1: 3
2: wx2	IP2: 255.255.128.0	IP2: 1
3: wx3	IP3: 255.255.255.240	IP3: 4
4: wx 4	IP4: 255.255.254.0	IP4: 2

To construct a table containing ifDescr and ipAdEntNetMask, OID of ipAdEntIfIndex would have to be used as the placement OID to "align" the columns. Thus, the OIDS input to the fetchlet would be "1.3.6.1.2.1.2.1.2.2.1.2 1.3.6.1.2.1.2.1.4.20.1.3*1.3.6.1.2.1.2.1.4.20.1.2". The fetchlet output will be as follows:

Figure 21–8 SNMP Fetchlet: Table Containing ifDescr and ipAdEntNetMask

wx1	255.255.128.0
wx2	255.255.254.0
wx3	255.255.255.0
wx 4	255.255.255.240

If OIDS were "1.3.6.1.2.1.2.1.2.2.1.2 1.3.6.1.2.1.2.1.4.20.1.3" for the previous example, the output would be as follows:

Figure 21–9 SNMP Fetchlet: Alternate OID

wx1	
wx2	
wx3	
wx 4	
	255.255.255.0
	255.255.128.0
	255.255.255.240
	255.255.254.0

21.5 HTTP Data Fetchlets

The HTTP data fetchlets obtain the contents of a URL and returns the contents of the URL as data. Three fetchlets are available:

- URL Fetchlet
- URL Lines
- URL Lines Token

21.5.1 URL Fetchlet (raw)

The URL fetchlet gets the contents of a given URL and returns the contents of the URL in a single cell table.

To get the output of a URL tokenized by lines and each line tokenized by a given delimiter, see the URL Line Token fetchlet.

Input Parameters

Table 21–6 URL Fetchlet Input Parameters

Name	Description	Use
url	URL to retrieve the contents of	required
proxyHost	proxy host through which to make the URL connection.	optional
proxyPort	proxy port through which to make the URL connection.	optional

Example

Take the following URL:

```
http://localhost/nhcities.txt
```

It has the following contents:

Line 1: Nashua, Keene,

Line 2: Concord

Line 3: , Conway, Manchester, Milford, Brookline,

Line 4:

Line 5: Hollis, Meredith

Now run the URL fetchlet with the given URL.

The fetchlet returns the following one-by-one table:

Figure 21–10 URL Fetchet Output

Nashua, Keene, Concord , Conway, Manchester, Milford, Brookline, Hollis, Meredith
--

The raw contents of the URL is returned.

Error Handling

MissingParameterException if URL parameter is missing.

FetchletException if the URL is malformed or an I/O error occurs in retrieving the content of the URL.

21.5.2 URL Lines Fetchlet (split into lines)

The URL fetchlet gets the contents of a given URL and tokenizes the contents of the URL. The output is tokenized by lines. The fetchlet returns the tokens in a single column table. The nth row in the table represents the nth line of the URL contents.

Note: To get the raw, untokenized contents of a URL, see the URL fetchlet. To get the contents of a URL tokenized by lines and each line tokenized by a given delimiter, see the URL Line Token fetchlet.

Table 21-7 URL Lines Fetchlet Input Parameters

Name	Description	Use
url	URL to retrieve the contents of	required
proxyHost	proxy host through which to make the URL connection.	optional
proxyPort	proxy port through which to make the URL connection.	optional
startsWith	only lines starting with this string are included in the result	optional; default = "" (all lines are included)

Example

Take the following URL:

`http://localhost/nhcities.txt`

It has the following contents:

Line 1: Nashua, Keene,

Line 2: Concord

Line 3: , Conway, Manchester, Milford, Brookline,

Line 4:

Line 5: Hollis, Meredith

Now run the URL fetchlet with the given URL.

The fetchlet returns the following table:

Figure 21-11 URL Lines Fetchlet Output

Nashua, Keene,
Concord
, Conway, Manchester, Milford, Brookline,
Hollis, Meredith

Error Handling

MissingParameterException if URL parameter is missing.

FetchletException if the URL is malformed or an I/O error occurs in retrieving the content of the URL.

Notes

Lines are tokenized using "\n".

21.5.3 URL Line Token Fetchlet (tokenized lines)

The URL fetchlet gets the contents of a given URL and tokenizes the contents of the URL. The output is tokenized first by lines, and then each line is tokenized by a given delimiter set. The fetchlet returns the tokens in a table. The *n*th row in the table represents the *n*th line of the URL content. The *n*th column in the table represents the *n*th token in a line as determined by the given delimiter set.

To get the raw, untokenized contents of a URL, see the URL fetchlet.

Table 21–8 URL Line Token Fetchlet Input Parameters

Name	Description	Use
url	URL to retrieve the contents of	required
delimiter	set of characters that act as delimiters to tokenize the lines	optional; default = "" (just breaks output into lines)
proxyHost	proxy host through which to make the URL connection.	optional
proxyPort	proxy port through which to make the URL connection.	optional
startsWith	only lines starting with this string are included in the result	optional; default = "" (all lines are included)

Example

Take the following URL:

`http://localhost/nhccities.txt`

It has the following contents:

Line 1: Nashua, Keene,

Line 2: Concord

Line 3: , Conway, Manchester, Milford, Brookline,

Line 4:

Line 5: Hollis, Meredith

Now run the URL fetchlet with the given URL and a single character "," for the delimiter.

The fetchlet returns the following table:

Figure 21–12 URL Token Lines Output

Nashua	Keene		
Concord			
Conway	Manchester	Milford	Brookline
Hollis	Meredith		

Error Handling

MissingParameterException if URL parameter is missing.

FetchletException if the URL is malformed or an I/O error occurs in retrieving the content of the URL.

Notes

Lines are tokenized using "\n".

The delimiter can be a single character or a set of characters. For example, it can be "|+_ ", if the line should be broken up by pipes, pluses, and underscores. If two or more delimiters are together in the output text, such as "||" or "+|+", then it is as if there are empty string tokens between them. These empty strings get columns in the result table. It is NOT considered that there are empty strings preceding a delimiter that starts a line or following a delimiter that ends a line.

In order to express non-printable characters in the delimiter set (such as tabs) in XML, use "&#xHH;" where H is the hexadecimal identifier for the character.

21.6 URLXML Fetchlet

The URL XML fetchlet obtains the XML content of a given URL, and extracts information based on a given pattern. A pattern is a list of "chunks" of XML to match against. The return table is a table with a column for each grabber (*) in the pattern in order and a row each time the pattern chunk matches in the XML content.

Input Parameters**Table 21–9 URLXML Fetchlet Input Parameters**

Name	Description	Use
url	URL to retrieve the contents of	Required.
pattern	The pattern used to extract information from XML; this is a list of XML chunks that is compared against the XML content of the URL. Each chunk contains one or more "grabbers" (*) in the text portion of the elements that define what should be flattened into text and extracted.	Required.
proxyHost	The proxy host through which to make the URL connection.	Optional.
proxyPort	The proxy port through which to make the URL connection.	Optional.

Table 21–9 (Cont.) URLXML Fetchlet Input Parameters

Name	Description	Use
ignoreDtd	If set to TRUE, specifies that the DTD file referenced by the content XML should be ignored. This is useful in cases where the DTD file cannot be accessed.	Optional.
generateKey	If set to true, a unique key will be generated for each row. The key will occupy the first column of the result, and will be numeric.	Optional.
throwConnException	If set to TRUE, a java.net.ConnectException will be thrown. Otherwise, it will be caught and an empty result set will be returned. Setting this property to FALSE provides behavior which is consistent with the DMSFetchlet.	Optional. The default value is TRUE.

Example

Take the following URL:

`http://localhost/urlxmltestfile.xml`

It has the following content:

```
<?xml version="1.0"?>
<testfile>
  <test>Simple text</test>
  <test><level>A little more complex</level></test>
  <test></test>
  <notatest></notatest>
  <test>Yet more complexity<level>Even a little more complex</level>Will it ever
stop?</test>
  <test1>must match<level>extract me!</level></test1>
  <test1>must match here<level>extract me, too!</level></test1>
</testfile>
```

Running the URL XML fetchlet with the given URL and the pattern:

```
<testfile><test>*<level>*</level></test></testfile>
```

returns the following table:

Figure 21–13 URL XML Fetchlet Output

A little more complex	A little more complex
Yet more complexityEven a little more complexWill it ever stop?	Even a little more complex

Error Handling

MissingParameterException if URL or pattern parameters are missing.

A FetchletException is generated if:

- The URL is malformed.
- An I/O error occurs in retrieving the content of the URL.
- The URL contents or pattern contains invalid XML.

Notes

Setting the proxy host and/or port changes these settings for the java.net package for the whole Java environment and is not thread-safe if the proxy settings are changing.

21.7 URL Timing Fetchlet

The URL Timing fetchlet gets the contents of a given URL, timing not only the base page source but any frames or images in the page as well.

Input Parameters

Table 21–10 URL Timing Fetchlet Input Parameters

Parameter	Description	Use
url#	URL(s) to download. "url0" is required but any number of URLs can be specified beyond url0 that following the convention: url0, url1, url2, url3.	Required.
proxy_host	Proxy host used to make a URL connection.	Optional. Specifies the proxy to be used for accessing URLs. If the proxy_host_override value is provided, then that value will be used instead.
proxy_port	Port used by the proxy host used make the URL connection.	Optional.
dont_proxy_for	Domains for which the proxy will not be used.	Optional. For example, .us.example.com, .uk.example.com
use_proxy	When used in conjunction with the proxy override input parameters, use_proxy specifies a proxy to be used in lieu of the original proxy. When set to false without the proxy override parameters set, no proxy is used.	Optional. Parameter can be set to true or false.
proxy_host_override	Alternate proxy host used to make the URL connection.	Optional. Overrides proxy_host.
proxy_port_override	Alternate proxy port used to make the URL connection.	Optional. Overrides proxy_port.
dont_proxy_override	Do not use the proxy for domains.	Optional. Parameter can be set to true or false.
internet_cert_loc	Path pointing to the location of a certificate to be used to access a secure (HTTPS) URL.	Optional.
auth_realm	Realm for the Basic Authentication log on. If the realm is not specified for the authentication, authentication does not occur and the download of the page fails with a 401 response code.	Optional.
auth_user	User name for Basic Authentication.	Optional.
auth_password	Password for Basic Authentication.	Optional.
retries	Number of times to retry the url if it initially fails.	Optional. Default = 1

Table 21–10 (Cont.) URL Timing Fetchlet Input Parameters

Parameter	Description	Use
connection_timeout	Wait time (in milliseconds) allowed to establish a connection to a server. This time also includes time required for name resolution.	Optional. Default= 60000 milliseconds (1 minute)
read_timeout	Idle time in the read waiting for the server to respond. For example, if no data is received from the server during the specified timeout period, the operation is considered failure.	Optional. Default = 12000 milliseconds (2 minutes)
timeout	Number of milliseconds after which the page download is considered a failure. This will detect if the site is functional but is extremely slow.	Optional. Default = 300000 ms (5 minutes)
status_comparator	When collating the rows to make a single row, the status_comparator parameter will indicate whether all URLs should have been a success (and) or any URLs should have been a success (or).	Optional. Default = and
cache	Indicates whether to use a cache when accessing an URL. Set the parameter to "n" to specify that no cache be used.	Optional. Default = y Note: The scope of the cache is per request. There is no persistent cache across multiple get metric requests.
output_format	Specifies the output format to be used: summary, detailed, repeat_column. For more information on output formats, see Metric Columns and Output Modes on page 21-26.	Required. summary : gives a default set of metrics in a single row for all urls detailed: gives a default set of metrics for each url. repeat_column : gives a single row of metric with timing for each of the url.
metrics	Specifies which metric columns need to be returned. For more information on metrics columns returned for each output format, see Table 21–12, "URLTIMING Fetchlet: Metric Columns"	Optional. Allows you to specify of what needs to be returned from the fetchlet and in which order. Example: status, status_description, total_response_time

Metric Columns and Output Modes

The format of information and specific metric information returned are controlled by the "output_format" and "metrics" input parameters. [Table 21–11](#) lists the format categories and the metrics (columns) returned by each. For a description of available metric columns, see [Table 21–12, "URLTIMING Fetchlet: Metric Columns"](#)

Table 21–11 URLTIMING Fetchlet: Output Formats

Output Format	Description	Metric Columns
summary	Returns a default set of metrics in a single row for all URLs If the metrics input parameter is specified, then only the columns specified will be returned.	computed_response_time, status, status_description, dns_time, connect_time, redirect_time, first_byte_time, html_time, content_time, total_response_time, rate, max_response_time, avg_response_time, avg_connect_time, avg_first_byte_time, broken_count, broken_content
detailed	Returns a default set of metrics for each url. If the metrics input parameter is specified, then only the columns specified will be returned.	url, computed_response_time, status, status_description, dns_time, connect_time, redirect_time, first_byte_time, html_time, content_time, total_response_time, rate, redirect_count, html_bytes, content_bytes, total_bytes, avg_connect_time, avg_first_byte_time, broken_count, broken_details
repeat_column	Returns a single row of metrics with timing for each of the URLs. If the metrics input parameter is specified, then those columns will be returned for each URL followed by overall status and status_description. (Note the output will always be single row).	total_response_time repeated for each URL followed by overall status and status_description.

Metric Columns

[Table 21–12](#) shows the metric columns returned by the URLTIMING fetchlet.

Table 21–12 URLTIMING Fetchlet: Metric Columns

Column Name	Description
status	The overall status of all URLs. By default AND is used to compute overall status but this can be changed using the status_comparator input parameter.
connect_time	The time to connect to the server and send the request.
first_byte_time	Time taken between sending the request and reading the first byte from the response.
total_response_time	Time taken for fetching ALL urls and associated content (gif, css, javascript, and so on).
max_response_time	Also referred as Slowest page time. This the time taken by the slowest URL.
avg_response_time	Average response time for URL. Computed as total response time / number of pages (urls).
rate	Kilo Bytes per second. This is computed by total bytes received / total time taken to receive them.

Table 21–12 (Cont.) URLTIMING Fetchlet: Metric Columns

Column Name	Description
html_time	Total time taken to download the html part of all pages. This time excludes time to fetch images and other contents. (Includes time to fetch frame html).
content_time	Time taken to download the page content (gif, javascripts, css, and so on).
redirect_time	Total time taken for all redirects occurring while fetching the set of urls specified.
redirect_count	Number of redirects.
total_bytes	Total number of bytes.
html_bytes	Total number of HTML bytes. (Includes bytes for frame html).
content_bytes	Total number of content bytes.
status_description	This is present only when the status is down. Corresponds to HTTP Status description.
request_count	Number of request made. (Includes all html as well as content requests).
broken_count	Number errors while fetching images or other content elements.
broken_details	List of images or other content elements that could not be fetched. This has format of url[broken list], url[broken list...
computed_response_time	This time approximates the time it would have taken for a client (like browser) to fetch all the pages in the transaction. This number is computed as if the contents of every page (gifs, css, and so on) were fetched using multiple threads.
avg_connect_time	Total connect_time / total number of connections made.
avg_first_byte_time	Total First Byte Time / Number of requests made (either to fetch HTML or content).
dns_time	Time to resolve host name (not implemented, always returns zero).
url	Returns the url, can only be used in 'detailed' output_format.

Example

Take the following URL:

```
url0=http://www.example.com/
```

With the input parameter output_format=summary, the fetchlet returns the following table (minus the headers on the columns):

Figure 21–14 Summary Output Format

computed response time	status	status description	dns time	connect time	redirect time	first byte time	html time	content time	total response time	rate (Kbytes per second)	max response time	avg response time	avg connect time	avg first byte time	broken count	broken content
540	1		0	548	0	149.0	1.0	7.0	705	95.16	705	705	54.80	14.90	0	

With `output_format = summary` and `metrics = total_response_time, status, status_description` the fetchlet returns the following table (minus the headers on the columns):

Figure 21-15 Summary Output Format with Specified Metric Columns

total response time	Status	Status description
705.0	1	

With `output_format = summary` and `metrics = total_response_time, status, status_description` the fetchlet returns the following table (minus the headers on the columns) and the server is giving error:

Figure 21-16 Summary Output Format with Specified Metric Columns and Internal Server Error

total response time	status	status description
	0	Internal Server Error -- http://www.example.com

Take the following URL:

```
url0=http://www.example.com/
url1=http://nedc.us.example.com/
```

With the `output_format=summary`, the fetchlet returns the following table (minus the headers on the columns). Here the numbers are time taken for fetching both the urls.

Figure 21-17 Summary Output Format for Two URLs

computed response time	status	status description	dns time	connect time	redirect time	first byte time	html time	content time	total response time	rate (Kbytes per second)	max response time	avg response time	avg connect time	avg first byte time	broken count	broken content
4344	1		0	603	0	7277	438	318.0	8636	16.92	8098	4318	22.33	269.52	0	

With the `output_format=detailed`, the fetchlet returns the following table (minus the headers on the columns):

Figure 21-18 Detailed Output for Two URLs

url	uri	computed response time	status	status description	dns time	connect time	redirect time	first byte time	html time	content time	total response time	rate	redirect count	html bytes	content bytes	total bytes	avg connect time	avg first byte time	broken count	broken content
http://www.example.com	not supported	531.0	1			548	0	149.0	1	7	705	95.16	0	18207	48883	67090	54.80	14.90	0	
http://nedc.us.example.com	not supported	4424.0	1			480	0	6227.0	442	230	7379	10.71	0	24627	54427	79054	28.24	366.29	0	

With the `output_format=repeat_column`, the fetchlet returns the following table (minus the headers on the columns):

Figure 21–19 Repeat Column Output Format

total response time(oracle.com)	total response time (necdc)	status	status description
705.0	7379	1	

Error Handling

Metric error if the URL parameter is missing, malformed, or if the metric cannot be computed.

Notes

The time required to perform a retry will be added on to the total time of the page. For example, if two retries are performed and then a success occurs, the total page time will be the time of the page that succeeded plus the time it took for the two retries to fail.

Proposed usage:

- For basic monitoring:
Use `url0=<URL to be monitored>` , `output_mode=summary` and specify `metrics=status, computed_response_time, status_description`
- For getting all columns:
Use `url0=<url to be monitored>` , `output_mode=summary`

21.8 Dynamic Monitoring Service (DMS) Fetchlet

The Dynamic Monitoring Service (DMS) fetchlet contacts an Application Server (AS) and then collects the metrics instrumented by the DMS.

The DMS allows application and system developers to measure and export customized, component-specific performance metrics. The Oracle Management Agent allows software components to import runtime performance data into Oracle Enterprise Manager Cloud Control.

The DMS fetchlet is an Oracle Management Agent plug-in module that allows the Management Agent to import the performance data that is exported by the DMS. Using the DMS fetchlet, any component that is instrumented using DMS API calls may share its performance data with Enterprise Manager Cloud Control.

21.8.1 Advantages to Using DMS for Oracle Management Agent Integration

With DMS, a component can insulate itself from the operational details of the Management Agent. A component would not need to deploy (or maintain) its own fetchlet or deploy (or maintain) a Tcl script or shell script to plug into one of the existing fetchlets. A component would not need to devise its own new way of measuring or exporting performance metrics. Performance metrics can be measured and reported in a consistent way across components. The DMS fetchlet contacts the remote DMS runtime directly with no need for forking shell scripts or Tcl scripts. Most

importantly, DMS automatically produces the long, complicated metadata document for you and thereby saves many hours of tedious and error-prone hand editing.

Input Parameters

Table 21–13 DMS Fetchlet Input Parameters

Name	Type	Description	Use
oraclehome	String	Top directory under which the monitored IAS instance is installed. It is used only for monitoring local IAS processes. For monitoring remote IAS processes, users should give it an empty value and specify property "opmnremoteport" and/or "machine" instead.	Required. Example: "/private/oracle/ias"
version	String	AS Version number of the target. It is used to distinguish the version of monitored AS instance.	Optional Example: "9.0.4"
opmnport	Integer	Oracle Process Monitoring and Notification (OPMN) port. It is used primarily for monitoring remote AS processes. It should be specified together with property "machine". If it is present and valid, property "oraclehome" and "httpport" are ignored.	Optional Example: "6200"
httpport	Integer	HTTP port is used primarily for monitoring stand-alone processes. It should be specified together with property "machine". It will be ignored, if property "opmnport" is present. If it is present and valid, property "oraclehome" is ignored.	Optional Example: "7777"
machine	String	Host name where the Internet Application Server (AS) instance runs. It should be specified together with property "opmnport". If it is not present, the local host is assumed.	Optional Example: "my-sun.us.example.com"
metric	String	Name of the table-type metric.	Required Example: "Servlets"
columnOrder	String	A list of metric column names separated by ";". The column names must be specified in same order as they appear in the target type metadata file. Do not include "name", "host", "process" and "fullname" columns.	Required Example: "processTimes;totalRequestRate"

Table 21–13 (Cont.) DMS Fetchlet Input Parameters

Name	Type	Description	Use
usecache	String	Whether to cache this metric. The legal values are "true", "false" and "refreshall" with "true" being the default. The "refreshall" value tells the DMS to delete its cache data and retrieve the most recent data from all targets.	Optional. Example: "false" Setting "usecache" to "false" will bypass DMS caching
proxyHost	String	Proxy host through which to make the HTTP connection	Optional Example: "proxy.us.example.com"
proxyPort	Integer	Proxy port through which to make the HTTP connection	Optional Example: "80"
dontProxyFor	String	Domains for which the proxy will not be used.	Optional Example: ".us.example.com" or "18.219.0"
useDefaultProxy	String	When used in conjunction with the proxy override parameters, this variable specifies a proxy other than the original one. When set to false without the proxy override parameters set, no proxy at all is used.	Optional Example: "true" or "false"
proxyHostOverride	String	proxy host through which to make the HTTP connection	Optional Example: "www-proxy.us.example.com"
proxyPortOverride	Integer	proxy port through which to make the HTTP connection	Optional Example: "80"
authrealm	String	Realm for the Basic Authentication logon. If the realm is not specified for the authentication, authentication does not occur and the download of the page fails with a 401 response code.	Optional Example: "Please input your flex account login:"
authuser	String	Username for Basic Authentication	Optional "superuser"
authpwd	String	Password for Basic Authentication	Optional Example: "welcome"

Error Handling

The DMS fetchlet throws `MissingParameterException` if any of the properties "oraclehome", "metric", "columnOrder", "opmnport", or "httpport" is missing. It throws `FetchletException` if any of the ports given is not valid.

Notes

The first four columns of the metric table returned are always column "name", "fullname", "host" and "process". Therefore, do not include them in columnOrder string. Property "machine" should be specified together with either properties "opmnport" or "httpport". In this case, the property "oraclehome" is ignored.

21.8.2 DMS Fetchlet/Oracle Management Agent Integration Instructions

DMS has been used in several components (such as Apache, JServ, OSE, and Portal) to provide a consistent performance monitoring infrastructure for Oracle 9i Application Server. The Sensors are easy to use and save most of the work related to performance measurement because they hide most of the details related to timing, counting, and categorization. Finally, DMS hides many Management Agent details from component developers and much of the Management Agent integration effort.

21.8.2.1 Integrating DMS Data with the Management Agent

As mentioned earlier, DMS allows application and system developers to measure and export customized, component-specific performance metrics. The Oracle Management Agent enables software components to import runtime performance data into Enterprise Manager Cloud Control. This section describes how to integrate DMS performance metrics with the Management Agent.

Step 1: Install AS

Step 2: Install Enterprise Manager Cloud Control

Step 3: Instrument your Component with DMS

To enable DMS metrics for Enterprise Manager Cloud Control, you must follow two additional rules:

- **Rule 1: All Nouns exported to the Management Agent must have types**
Noun types can be set either by specifying the "type" parameter in the Noun.create() methods or by using the Noun.setType(String) method. The idea is that every Noun type will be converted automatically to a Management Repository table. Every Noun of a given type will become a row in the type's corresponding Management Repository table. The metrics contained by a Noun become columns in the Management Repository table metric. Any Noun without a type will not be exported to Management Agent.
- **Rule 2: All Nouns of a given type must contain a consistent set of Sensor names**
Because the metrics contained by a Noun become columns in a management repository table, you must make sure that all Nouns of a given type contain the same Sensors. This ensures that each row of the corresponding Management Repository table has the same set of columns. DMS does not check this constraint for you.

For example, the following Java snippet shows how to create typed Nouns that contain a consistent set of Sensors. DMS will automatically convert these into a Management Repository table named "MyType":

```
/* first create the nouns*/
Noun n1 = Noun.create("/myExample/myComponent/noun1", "MyType");
Noun n2 = Noun.create("/myExample/myComponent/noun2", "MyType");

/* next, create the Sensors */
PhaseEvent pe1 = PhaseEvent.create(n1, "criticalPhase", "a critical interval");
```

```
PhaseEvent pe2 = PhaseEvent.create(n2, "criticalPhase", "a critical interval");
Event e1 = Event.create(n1, "importantEvt", "an important event");
Event e2 = Event.create(n2, "importantEvt", "an important event");

/* here is a third set that shows the use of Noun.setType(String) */
PhaseEvent pe3 = PhaseEvent.create(
    "/myExample/myComponent/noun3/criticalPhase",
    "a critical interval");
Event e3 = PhaseEvent.create(
    "/myExample/myComponent/noun3/importantEvt",
    "an important event");
Noun n3 = Noun.get("/myExample/myComponent/noun3");
n3.setType("MyType");
```

For this example, the "MyType" table will contain three rows and four columns. Besides the columns corresponding to the two Sensors, there will be a "name" column and a "path" column that will contain the DMS path name including the process name and "/myExample/myCom...".

If these Nouns/Sensors are created in several servlet engines within the AS site, then the AggreSpy will find each of the servlet engines and will aggregate all of the Nouns/Sensors into a single MyType table.

Step 4: Generate your Target Metadata Document

You can generate the Target Metadata Document using your browser. Point your browser to your AS site that you want to monitor using the following URL:

```
http://YOUR_AS_HOST:YOUR_AS_PORT/YOUR_SERVLET_PATH/AggreSpy?format=targetmetadata
```

You should use the actual host, port and servlet path of your AS installation in the above URL. The servlet path usually defaults to "servlet". The XML document you get is the Target Metadata Document for your AS site. The first comment of the XML document explains where you can obtain the Target Metadata Document and instructions telling you what needs to be done to this document.

Step 5: Install the Target Metadata Document

Follow the steps described in the first comment of the XML document. Save the XML document to a file called "oracle_dms.xml" under the "metadata" directory of your Enterprise Manager installation (OMS_ORACLE_HOME/sysman/admin/metadata/). If you want to monitor a subset of the metrics or merge the metrics with the ones in the existing "oracle_dms.xml" file, you should save this new definition to a separate file called target_name.xml. You will also need to change the Target Type entry in the generated metadata document.

Next, you should add the target instance information of your AS site to your "targets.xml" file residing under the top directory of your Enterprise Manager installation. You can find a block of XML tags in the comment you read. They look like:

```
<Target Type='oracle_dms' NAME='DMS_YOUR-IAS-HOST_YOUR-IAS-PORT' VERSION='2.0'>
  <Property NAME='host' VALUE='YOUR_IAS_HOST' />
  <Property NAME='port' VALUE='YOUR_IAS_PORT' />
  <Property NAME='dmsPath' VALUE='YOUR_SERVLET_PATH' />
</Target>
```

Copy this block and paste it to the targets.xml file between <targets> and </targets> tags.

Finally, to add the new target metadata file and target instance information from the `targets.xml` file to Enterprise Manager Cloud Control, you must run the following command:

```
>$ORACLE_HOME/bin/emctl reload
```

Step 6: View Your Metrics

You are ready to view your metrics using Enterprise Manager's Metric Browser. First, make sure that AS and your component are still running. Next, restart the Oracle Management Agent. Finally, point your browser to your Management Agent installation using the following URL:

```
http://YOUR_AGENT_HOST:YOUR_AGENT_PORT/emd/browser/main
```

The Management Agent port information can be found in the `$AGENT_HOME/sysman/config/emd.properties` file at the `EMD_URL` line.

You should use the actual host and port of your Management Agent installation in the above URL. You will find your AS site listed as the target "DMS_YOUR-AS-HOST_YOUR-AS-PORT". If you click the link, you will see a list of metric IDs. You can browse your metrics by clicking on the respective metric IDs.

21.9 JDBC Fetchlet

Call-level interfaces such as JDBC permit external access to SQL database manipulation and update commands. The Java Database Connectivity (JDBC) fetchlet allows you to execute common JDBC commands and obtain their response time for any type of database.

Input Parameters

Table 21–14 JDBC Fetchlet Input Parameters

Name	Description	Use
Transaction Name	(Standard)	Required.
Beacon Name	(Standard)	Required.
Connect String	<p>Connection string provided by the user. The Connect String must comply with the URL format specified by the vendor of the database to which the user is trying to connect.</p> <p>Examples:</p> <p>Format required by Oracle:</p> <pre>jdbc:oracle:thin:@hostname:port</pre> <p>Format required by MySQL:</p> <pre>jdbc:mysql://hostname:port</pre>	Required.

Table 21–14 (Cont.) JDBC Fetchlet Input Parameters

Name	Description	Use
Class Name String	<p>The driver class name to be used for connections.</p> <p>Example:</p> <p>oracle.jdbc.driver.OracleDriver</p> <p>You have two options for configuring the Agent to use the .jar file containing the driver:</p> <ol style="list-style-type: none"> 1. Place the .jar file in \$JAVA_HOME/jre/lib/ext. CLASSPATH does not need to be modified. 2. Place the .jar file anywhere and update CLASSPATH in emd.properties file with the path to jar. Bounce Agent. This should be scripted and be transparent to user. 	Required.
Username	User name to be used when connecting to the database.	Required.
Password	Password to be used when connecting to the database.	Required.
Role	User Role	Required.
Statement	SQL statement to be executed. Use of PL/SQL is possible by using prepareCall() API.	Required.

Table 21–15 Metric Columns Collected

Column	Description
Status	Status of the test. Status is 'down' if there is a SQLException generated by the fetchlet.
Total Time	Time required for the fetchlet to execute the test.
Connect Time	Time required for DriverManager.getConnection() to complete.
Prepare Time	Time required for conn.prepareStatement() to complete.
Execute Time	Time required for stmt.executeQuery() to complete.
Fetch Time	Time required for while(rs.next()) { rs.getRow() } to complete.
Close Time	Time required for closing resultset, statement, connection to complete.
Number of rows	Number of rows fetched.
Total time per row	
Fetch time per row	

Example

[Example 21–10](#) provides the properties passed to the JDBC fetchlet when invoked.

Example 21–10 Properties Passed to JDBC Fetchlet

```
<QueryDescriptor FETCHLET_ID="JDBC">
<Property NAME="TxnName" SCOPE="GLOBAL">TxnName</Property>
<Property NAME="BeaconName" SCOPE="GLOBAL">BeaconName</Property>
```

```

<Property NAME="connstring" SCOPE="INSTANCE">connString</Property>
<Property NAME="username" SCOPE="INSTANCE">username</Property>
<Property NAME="password" SCOPE="INSTANCE">password</Property>
<Property NAME="statement" SCOPE="GLOBAL">select * from user_tables</Property>
<Property NAME="classstring" SCOPE="GLOBAL">oracle.jdbc.none</Property>
<Property NAME="role" SCOPE="GLOBAL" OPTIONAL="TRUE">DBA</Property>
<Property NAME="useconnpool" SCOPE="GLOBAL" OPTIONAL="TRUE">FALSE</Property>
<Property NAME="GetTimingData" SCOPE="GLOBAL">TRUE</Property>
</QueryDescriptor>

```

21.10 WBEM Fetchlet

The WBEM fetchlet accesses a CIMOM and retrieves requested information using the specified CIM class. The CIM class is mapped to a Management Repository table metric. The name of the CIM class is the name of the table metric that is returned, and the properties defined for the CIM class are used to name the table columns for the metric. The properties of interest must be specified during metric definition.

The fetchlet returns the instances that have been instantiated for the CIM class as rows of the Management Repository table metric.

Input Parameters

Table 21–16 *WBEM Fetchlet Input Parameters*

Name	Type	Description	Use
hostname	String	Host name of the CIMOM	Optional; default is "localhost"
port	Integer	Port for the CIMOM	Optional; default is 5988
namespace	String	CIM Namespace	Optional; default is "root/cimv2"
username	String	User name to use for CIMOM authorization on the host where the CIMOM is running	Required
password	String	Password to use for CIMOM authorization on the host where the CIMOM is running	Required
CIMclassname	String	Name of the CIM class whose instances will be returned	Required for all operations except STATUS. STATUS operations just check whether the CIMOM is running, so a class name is not needed.
operation	String	Operation to be performed. Supported operations include COUNT, which returns a count of the number of instances in the class, VALUES, which returns the values of the specified properties for each instance of the class, or STATUS, which provides status information about the CIMOM.	Optional, default is VALUES

Table 21–16 (Cont.) WBEM Fetchlet Input Parameters

Name	Type	Description	Use
properties	String	The property names from the CIM class definition that we are interested in collecting.	Required for VALUES operation. If the operation is VALUES, we can have 1 to N of these, separated by a semicolon. If the operation is VALUES, and no properties are provided, an error is returned. Properties are handed to the EMD in the order that they are specified.

Error Handling

The following types of errors have been identified for the WBEM fetchlet.

MissingParameterException occurs when:

- No CIM Class parameters match.

Fetchlet exception occurs when:

- The class name is not found in the CIMOM namespace.
- The namespace is not found.
- The connection to the CIMOM does not have valid credentials.
- The connection to the CIMOM failed because the CIMOM was not running.
- The CIM class property does not exist
- An unsupported operation was specified
- No properties were specified.

Notes

Ports: Some CIMOM client interfaces expose the port that the CIMOM is listening on while some clients do not. To cover both cases, the port is exposed as an optional input parameter that defaults to port 5988. This is the default Pegasus CIMOM listener port. The Java API that is provided through Sun's Wbem Services does not expose the CIMOM port.

Protocols: Most CIMOMs support either an RMI or HTTP protocol for communicating with the CIMOM. The testing that has been done shows that the HTTP protocol is not as stable, and in some cases, not fully implemented in the CIMOM. Because of this, the protocol currently defaults to RMI. The actual parameters for the WBEM Services CIMOM for the protocol are: CIMClient.CIM_RMI or CIMClient.CIM_XML.

Fetchlet Operations: The WBEM APIs are very flexible at allowing clients to traverse the class hierarchies that are defined and their associations. At this point in time, the options on accessing CIM data from an EMD are restricted to counting, getting the properties of classes, and CIMOM status. These are the more important operations that need to be performed for monitoring. As additional requirements come in, we can add new operations to support them if necessary. For the prototype, only the count operation has been implemented.

Authentication: Most CIMOMs provide APIs to support authentication through a user identity mechanism. The majority of the CIMOMs have not implemented the API, so this capability is really a no-op. In any case, we've supplied the capability in the fetchlet so that as CIMOM implementations catch up with the standard, we'll have the necessary support in place.

Examples

The Wbem fetchlet supports three basic operations. At this point, the fetchlet only handles one operation at a time, so you cannot mix count, status, and value operations within a single fetchlet call. [Example 21–11](#) shows how to write the metadata for a COUNT operation:

Example 21–11 COUNT Operation Metadata

```
<Metric NAME="Load" TYPE="TABLE">
  <Display>
    <Label NLSID="wbem_cimom_load">Load</Label>
  </Display>
  <TableDescriptor>
    <ColumnDescriptor NAME="Active Clients" TYPE="NUMBER" IS_KEY="FALSE">
      <Display>
        <Label NLSID="wbem_cimom_active_clients">Active CIMOM Clients</Label>
      </Display>
    </ColumnDescriptor>
  </TableDescriptor>
  <QueryDescriptor FETCHLET_ID="Wbem">
    <Property NAME="username" SCOPE="GLOBAL">guest</Property>
    <Property NAME="password" SCOPE="GLOBAL">guest</Property>
    <Property NAME="CIMclassname" SCOPE="GLOBAL">EX_SFLProvider</Property>
    <Property NAME="operation" SCOPE="GLOBAL">COUNT</Property>
  </QueryDescriptor>
</Metric>
```

The FETCHLET_ID is identified as Wbem. Property names are passed to the fetchlet for the required parameters user name, password, and CIMclassname. The operation is identified as COUNT.

The following example shows how to implement a Response Status metric to determine whether the CIMOM is running or not. It returns a value of 1 if the connection to the CIMOM is successful, otherwise 0.

Example 21–12 Response Status Metric

```
<Metric NAME="Response" TYPE="TABLE">
  <Display>
    <Label NLSID="wbem_cimom_response">Response</Label>
  </Display>
  <TableDescriptor>
    <ColumnDescriptor NAME="Status" TYPE="NUMBER" IS_KEY="FALSE">
      <Display>
        <Label NLSID="wbem_cimom_response_status">Status</Label>
      </Display>
    </ColumnDescriptor>
  </TableDescriptor>
  <QueryDescriptor FETCHLET_ID="Wbem">
    <Property NAME="username" SCOPE="GLOBAL">guest</Property>
    <Property NAME="password" SCOPE="GLOBAL">guest</Property>
    <Property NAME="operation" SCOPE="GLOBAL">STATUS</Property>
  </QueryDescriptor>
</Metric>
```

The default operation is the VALUES operation. It is used to fetch the values of a class that is defined in the CIMOM.

In the final example, the EX_Teacher class is accessed and fetches the name column. Name is the key of the class and of the new metric being defined, so the IS_KEY property is set to true. The CIM class properties will be mapped to the Enterprise Manager columns in the order that they are specified in the properties property. In this case, there is only 1 property - Name.

Example 21-13 Single Property Fetched for a Class

```
<Metric NAME="EX_Teacher" TYPE="TABLE">
  <Display>
    <Label NLSID="wbem_EX_Teacher">EX_Teacher Class</Label>
  </Display>
  <TableDescriptor>
    <ColumnDescriptor NAME="Name" TYPE="STRING" IS_KEY="TRUE">
      <Display>
        <Label NLSID="wbem_ex_teacher_name">Name</Label>
      </Display>
    </ColumnDescriptor>
  </TableDescriptor>
  <QueryDescriptor FETCHLET_ID="Wbem">
    <Property NAME="username" SCOPE="GLOBAL">guest</Property>
    <Property NAME="password" SCOPE="GLOBAL">guest</Property>
    <Property NAME="CIMclassname" SCOPE="GLOBAL">EX_Teacher</Property>
    <Property NAME="properties" SCOPE="GLOBAL">Name</Property>
  </QueryDescriptor>
</Metric>
```

If multiple properties are fetched for a class, semi-colons should separate them. The properties should be provided in the order that the column descriptors are specified for the metric table definition.

Example 21-14 Multiple Properties Fetched for a Class

```
<Metric NAME="EX_SFLProvider" TYPE="TABLE">
  <Display>
    <Label NLSID="wbem_EX_SFLProvider">EX_SFLProvider Class</Label>
  </Display>
  <TableDescriptor>
    <ColumnDescriptor NAME="Name" TYPE="STRING" IS_KEY="TRUE">
      <Display>
        <Label NLSID="wbem_ex_sfl_name">Name</Label>
      </Display>
    </ColumnDescriptor>
    <ColumnDescriptor NAME="Win" TYPE="NUMBER" IS_KEY="FALSE">
      <Display>
        <Label NLSID="wbem_ex_sfl_win">Win</Label>
      </Display>
    </ColumnDescriptor>
    <ColumnDescriptor NAME="Lost" TYPE="NUMBER" IS_KEY="FALSE">
      <Display>
        <Label NLSID="wbem_ex_sfl_lost">Lost</Label>
      </Display>
    </ColumnDescriptor>
  </TableDescriptor>
  <QueryDescriptor FETCHLET_ID="Wbem">
    <Property NAME="username" SCOPE="GLOBAL">guest</Property>
```

```

<Property NAME="password" SCOPE="GLOBAL">guest</Property>
<Property NAME="CIMClassname" SCOPE="GLOBAL">EX_SFLProvider</Property>
<Property NAME="properties" SCOPE="GLOBAL">Name;Win;Lost</Property>
</QueryDescriptor>
</Metric>

```

21.11 JMX Fetchlet

The JMX fetchlet retrieves Java Management Extensions (JMX) attributes (or invokes a JMX operation) from an MBean and returns the result as a (table) metric. If the ObjectName specified is an ObjectName pattern, then multiple rows are returned. Each row corresponds to an MBean matching the specified ObjectName pattern.

Input Parameters

Table 21–17 JMX Fetchlet Major Input Parameters

Name	Type	Description	Use
MachineName	String	MBean server host name	Optional
Port		Port on which the MBean server is listening for new connections	Optional
UserName	String	User name for JMX connections, if required	Required
password	String	Password for JMX connections, if required	Required
protocol	String	Protocol used for the connection	Optional
service	String	Service used for connection	Optional
serviceURL	String	serviceURL used for JMX connection. This is instead of the previous MachineName, Port, protocol, and service properties. Note: For middleware targets, the serviceURL can be obtained from either the farm or managedServer association depending on whether metric needs to be collected from AdminServer or the managed server.	Required (unless MachineName and Port are specified)
Metric	String	Mbean object name (or if MetricService=true, the DMS table name)	Required
columnOrder	String	Semi colon separated list of JMX attributes for the previous MBean corresponding to the column definitions in the TableDescriptor of the metric.	Required
operation	String	Name of the JMX operation to be invoked. In this case, the columnOrder represents the values from the return object to be populated in the Metric. (Oracle recommends using jmxcli to generate this).	Optional
arguments	String	The XML representing the arguments for the JMX operation. Oracle recommends using jmxcli to generate this.	Optional

Table 21–17 (Cont.) JMX Fetchlet Major Input Parameters

Name	Type	Description	Use
MetricService	Boolean	MetricService=true implies that the metric is retrieved by the Oracle-specific DMS Metric Service. In this case the previous columnOrder property is a list of column names and the 'metric' property indicates the actual DMS table name.	Optional
identityCol	String	<p>The Mbean object name key (or a semi-colon separated list of keys) that will be extracted from the Mbean ObjectName and surfaced as key columns in the resultant metric.</p> <p>If the value 'canonical' is specified, an additional key metric column with the complete Mbean object name is returned by the fetchlet.</p> <p>This property makes sense only if the previous metric property is an ObjectName pattern that matches more than one Mbean on the server.</p>	Optional
autoRowID	String	<p>Prefix for an automatically generated key column. The suffix is sequential numbers starting at 1.</p> <p>For example, autoRowID set to ROW_ generates a key column at position 0 with values ROW_1, ROW_2, and so on up to the number of rows returned.</p> <p>This is usually the case if none of the other columns (JMX attributes selected) are unique and multiple rows are returned as a result of multiple mbeans matches and mbean pattern.</p>	Optional
useCache	Boolean	Applicable only when MetricService=true and indicates if metric service cache needs to be used	Optional
ServerNames	String	Applicable only when MetricService=true and is a semicolon list of server names from which the DMS metrics need to be retrieved. This is relevant only when collecting these metrics from the AdminServer (that is, serviceURL points to AdminServer through farm association), which has metrics from all managed servers	Optional
valueWhenNoMBean	Number	Typically used for response metrics and has the value that the fetchlet returns as a single row and column when no mbeans are found that match the given mbean pattern (in the previous metric property).	Optional
valueWhenDown	Number	Typically used for response metrics. This has the value that the fetchlet returns as a single row and column when the connection to the server fails due to a connection exception (indicating that the server is down).	Optional

Table 21–17 (Cont.) JMX Fetchlet Major Input Parameters

Name	Type	Description	Use
admlMap	String	Applicable only when MetricService=true and is an XML snippet that indicates what adml parameters need to be passed for this adml table. (Oracle recommends using jmxcli to generate this).	Optional

Notes:

1. The JMX fetchlet is used to retrieve primarily JMX attributes from Mbeans on a target MbeanServer. It can also retrieve attributes from multiple MBeans of the same kind in the form of a table (with multiple rows where each row represents a matching MBean).

For example, if an MBean ObjectName pattern specifies servlets, (that is, *:Type=ServletRuntime, * in the metric property), and the columnOrder specifies A1;A2;A3, then the resultant metric will have one row for each servlet.

2. If the metric data must be obtained using a JMX operation (this is not typical for collecting metrics), then the QueryDescriptor property operation must specify the JMX operation name and the arguments are an XML representation of the parameters to be passed into the JMX operation.

For example, the following QueryDescriptor indicates the invocation of a JMX operation called "getNumUserSessions" with a single string argument with a value="total".

Example 21–15 Specifying a JMX Operation Name

```
<Metric NAME="GetNumUserSessions" TYPE="TABLE" USAGE_TYPE="HIDDEN">
  <Display>
    <Label NLSID="GetNumUserSession">GetNumUserSession</Label>
  </Display>
  <TableDescriptor>
    <ColumnDescriptor NAME="Get Num User Sessions" TYPE="STRING">
      <Display>
        <Label NLSID="Get Num User Sessions">Get NumUser Sessions</Label>
      </Display>
    </ColumnDescriptor>
  </TableDescriptor>
  <QueryDescriptor FETCHLET_ID="JMX">
    <Property NAME="serviceURL" SCOPE="ASSOCTGT" ASSOCIATION_
NAME="farm">serviceURL</Property>
    <Property NAME="UserName" SCOPE="ASSOCTGT" OPTIONAL="TRUE" ASSOCIATION_
NAME="farm">UserName</Property>
    <Property NAME="password" SCOPE="ASSOCTGT" OPTIONAL="TRUE" ASSOCIATION_
NAME="farm">password</Property>
    <Property NAME="instName.parameter"
SCOPE="INSTANCE">instName</Property>
    <Property NAME="metric"
SCOPE="GLOBAL">oracle.forms.FormsJ2EEApplication.%instName.parameter%:,type=Runtim
e,*</Property>
    <Property NAME="operation" SCOPE="GLOBAL">getNumUserSessions</Property>
    <Property NAME="columnOrder" SCOPE="GLOBAL">getNumUserSessions</Property>
    <Property NAME="arguments" SCOPE="GLOBAL">
      <![CDATA[<arguments>
<argument type="java.lang.String">
```

```

        <value>total</value>
    </argument>
</arguments>]]>
    </Property>
</QueryDescriptor>
</Metric>

```

3. A QueryDescriptor for the JMX fetchlet contains JMX connection information. This is usually in the form of a serviceURL. If the serviceURL property is not available in the QueryDescriptor, then the combination of MachineName, Port, protocol, and service properties must be present in the QueryDescriptor to provide connection information to the JMX fetchlet.

21.12 Web Services Fetchlet

In target metadata files generated by the Web Services Command-Line tool, the `<QueryDescriptor>` element specifies the properties that will be passed to the Web Services fetchlet when being invoked.

Input Parameters

Table 21–18 lists the supported properties:

Table 21–18 Web Services Fetchlet Properties

Name	Description	Use	Comments
ServiceName	Web service name	Required. Service Name must be prefixed with a valid namespace.	All referenced namespaces are specified by the property <i>"Namespace"</i>
PortName	Web service port name	Required. Port Name must be prefixed with a valid namespace.	All referenced namespaces are specified via the property <i>"Namespace"</i>
OperationName	Web service operation name	Required. Operation Name must be prefixed with a valid namespace.	All referenced namespaces are specified by the property <i>"Namespace"</i>
ServiceEndpoint	Web service endpoint	Required. A valid URL.	
WsdIURL	Web service WSDL URL	Optional. A valid URL.	Required only if it is a RPC/Encoded Web service
ParameterStyle	SOAP parameter mapping style	Optional. - BARE - WRAPPED	Optional only if it is a RPC/Encoded or REST-ful Web service
Payload	Web service operation request payload	Required. Must be specified using the CDATA section.	
SOAPBindingStyle	SOAP binding style	Optional. - DOCUMENT - RPC	Optional only if it is a RPC or Encoded Web service

Table 21–18 (Cont.) Web Services Fetchlet Properties

Name	Description	Use	Comments
SOAPBindingUse	SOAP binding use	Optional - <i>ENCODED</i> - <i>LITERAL</i>	Optional only if it is a RPC or Encoded Web service
SOAPVersion	SOAP version	Optional - <i>SOAP_1_1</i> - <i>SOAP_1_2</i>	Optional only if it is an RPC or Encoded Web service
MessageType	Web service message type	Optional - <i>SOAP</i> - <i>REST</i>	Optional only if it is a RPC or Encoded Web service
SecurityPolicy	Security policy	Required - <i>NONE</i> - <i>BASIC_AUTHENTICATION</i>	
Namespace	Set of all namespaces referenced	Optional. Contains all the namespaces referenced in the metric Specify using notation: [ns0="uri0"][ns1="uri1"] Example: [ns0="http://type.abc.com"] [ns1="http://app.abc.com"]	
ColType	Collection result column type	Required List of metric column type (separated by comma) Example: msgId:STRING,source:STRING,detail:STRING	
RowType	Collection result row type	Required List of XPath expression corresponding to metric columns (separated by comma) For example: //ns0:eventResponse/msgId , //ns0:eventResponse/source ,	
SSLKeyStoreCredential	SSL keystore credentialSet name	Optional A valid CredentialSet of a Store Credential Type defined in the <CredentialInfo>	Must be defined as a monitoring credential.
SSLTrustStoreCredential	SSL truststore credentialset name	Optional A valid CredentialSet of a StoreCredential Type defined in the <CredentialInfo>	Must be defined as a monitoring credential.

Table 21–18 (Cont.) Web Services Fetchlet Properties

Name	Description	Use	Comments
UserCredential	User token credentialset name	Optional A valid CredentialSet of a AliasCredential or CSFKeyCredential Type defined in the <CredentialInfo>	Must be defined as a monitoring credential.
ValueWhenDown	Default response when target is down	Required (only for response metric). Not required for regular metric. For Response metric, when a target is down, this value (if specified) will be returned.	A target is considered as down when the Fetchlet catches a ConnectionException.

Examples

[Example 21–16](#) provides an example of a metric definition for Remote Procedure Call (RPC) or encoded Web services and [Example 21–17](#) provides an example of a metric definition for doc or literal Web services.

Example 21–16 Metric Definition for RPC or Encoded Web Service

```

<Metric NAME="getVacantRooms" TYPE="TABLE">
  <Display>
    <Label NLSID="NLSID_GET_VACANT_ROOMS">getVacantRooms</Label>
  </Display>

  <TableDescriptor>
    <ColumnDescriptor IS_KEY="TRUE" NAME="roomID" TYPE="STRING">
      <Display>
        <Label NLSID="COL_ROOM_ID">roomID</Label>
      </Display>
    </ColumnDescriptor>
    <ColumnDescriptor IS_KEY="FALSE" NAME="floor" TYPE="STRING">
      <Display>
        <Label NLSID="COL_FLOOR">floor</Label>
      </Display>
    </ColumnDescriptor>
    <ColumnDescriptor IS_KEY="FALSE" NAME="number" TYPE="STRING">
      <Display>
        <Label NLSID="COL_NUMBER">number</Label>
      </Display>
    </ColumnDescriptor>
    <ColumnDescriptor IS_KEY="FALSE" NAME="rate" TYPE="STRING">
      <Display>
        <Label NLSID="COL_RATE">rate</Label>
      </Display>
    </ColumnDescriptor>
    <ColumnDescriptor IS_KEY="FALSE" NAME="roomType" TYPE="STRING">
      <Display>
        <Label NLSID="COL_ROOM_TYPE">roomType</Label>
      </Display>
    </ColumnDescriptor>
    <ColumnDescriptor IS_KEY="FALSE" NAME="smoking" TYPE="STRING">
      <Display>
        <Label NLSID="COL_SMOKING">smoking</Label>
      </Display>
    </ColumnDescriptor>
  </TableDescriptor>
</Metric>

```

```

    <ColumnDescriptor IS_KEY="FALSE" NAME="available" TYPE="STRING">
      <Display>
        <Label NLSID="COL_AVAILABLE">available</Label>
      </Display>
    </ColumnDescriptor>
  </TableDescriptor>
<QueryDescriptor FETCHLET_ID="WSF">
  <Property NAME="SecurityPolicy" SCOPE="INSTANCE">NONE</Property>
  <Property NAME="WsdURL" SCOPE="INSTANCE">wsdlURL</Property>
  <Property NAME="ServiceEndpoint" SCOPE="INSTANCE">serviceURL</Property>
  <Property NAME="ServiceName"
SCOPE="GLOBAL">ns0:SimpleHotelServiceRE</Property>
  <Property NAME="PortName" SCOPE="GLOBAL">ns0:HotelService</Property>
  <Property NAME="OperationName" SCOPE="GLOBAL">getVacantRooms</Property>
  <Property NAME="MessageType" SCOPE="GLOBAL">SOAP</Property>
  <Property NAME="SOAPBindingStyle" SCOPE="GLOBAL">RPC</Property>
  <Property NAME="SOAPBindingUse" SCOPE="GLOBAL">ENCODED</Property>
  <Property NAME="ParameterStyle" SCOPE="GLOBAL">BARE</Property>
  <Property NAME="SOAPVersion" SCOPE="GLOBAL">SOAP_1_1</Property>
  <Property NAME="Namespace"
SCOPE="GLOBAL"><![CDATA[[[ns1="http://hotel.apps.muws/"] [ns0="http://hotel.apps.muws/rpc/"]]]></Property>
  <Property NAME="RowType"
SCOPE="GLOBAL">//ns1:getVacantRoomsResponse/return/item/@roomID, //ns1:getVacantRoomsResponse/return/item/floor,

//ns1:getVacantRoomsResponse/return/item/number, //ns1:getVacantRoomsResponse/return/item/rate, //ns1:getVacantRoomsResponse/return/item/roomType,

//ns1:getVacantRoomsResponse/return/item/smoking, //ns1:getVacantRoomsResponse/return/item/available</Property>
  <Property NAME="ColType"
SCOPE="GLOBAL">roomID:STRING, floor:STRING, number:STRING, rate:STRING, roomType:STRING, smoking:STRING, available:STRING</Property>
  <Property NAME="Payload" SCOPE="GLOBAL"><![CDATA[<soap:Envelope
xmlns:soap="http://schemas.xmlsoap.org/soap/envelope/"
xmlns:ns="http://hotel.apps.muws/"
xmlns:xsd="http://www.w3.org/2001/XMLSchema"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xmlns:soapenc="http://schemas.xmlsoap.org/soap/encoding/">
  <soap:Body soap:encodingStyle="">
    <ns:getVacantRooms/>
  </soap:Body>
</soap:Envelope>]]></Property>
</QueryDescriptor>
</Metric>

```

Example 21-17 Metric Definition for Doc or Literal Web Service

```

<Metric NAME="square" TYPE="TABLE">
  <Display>
    <Label NLSID="NLSID_SQUARE">square</Label>
  </Display>
  <TableDescriptor>
    <ColumnDescriptor IS_KEY="FALSE" NAME="arg0" TYPE="STRING">
      <Display>
        <Label NLSID="COL_ARG0">arg0</Label>
      </Display>
    </ColumnDescriptor>
  </TableDescriptor>

```

```

    </TableDescriptor>
    <QueryDescriptor FETCHLET_ID="WSF">
      <Property NAME="SecurityPolicy" SCOPE="INSTANCE">NONE</Property>
      <Property NAME="ServiceEndpoint" SCOPE="INSTANCE">serviceURL</Property>
      <Property NAME="ServiceName"
SCOPE="GLOBAL">ns0:CalculatorService</Property>
      <Property NAME="PortName" SCOPE="GLOBAL">ns0:CalculatorPort</Property>
      <Property NAME="OperationName" SCOPE="GLOBAL">square</Property>
      <Property NAME="MessageType" SCOPE="GLOBAL">SOAP</Property>
      <Property NAME="SOAPBindingStyle" SCOPE="GLOBAL">DOCUMENT</Property>
      <Property NAME="SOAPBindingUse" SCOPE="GLOBAL">LITERAL</Property>
      <Property NAME="ParameterStyle" SCOPE="GLOBAL">WRAPPED</Property>
      <Property NAME="SOAPVersion" SCOPE="GLOBAL">SOAP_1_1</Property>
      <Property NAME="Namespace"
SCOPE="GLOBAL"><![CDATA[[ns0="http://tests.jaxws.oracle.com/" ]
[ns1="http://www.oracle.com/jaxws/tests"]]]></Property>
      <Property NAME="RowType"
SCOPE="GLOBAL">//ns1:squareResponse/arg0</Property>
      <Property NAME="ColType" SCOPE="GLOBAL">arg0:STRING</Property>
      <Property NAME="Payload" SCOPE="GLOBAL"><![CDATA[<soap:Envelope
xmlns:soap="http://schemas.xmlsoap.org/soap/envelope/">
      <soap:Body xmlns:ns1="http://www.oracle.com/jaxws/tests">
        <ns1:square>
          <arg0>%square.arg00001%</arg0>
        </ns1:square>
      </soap:Body>
    </soap:Envelope>]]></Property>
    </QueryDescriptor>
  </Metric>

```

21.12.1 Using Credentials for Authentication

If basic authentication is required, then you must configure or define the following in the metric definition:

1. Set the SecurityPolicy property to BASIC_AUTHENTICATION:

```

    <Property NAME="SecurityPolicy" SCOPE="INSTANCE">BASIC_
AUTHENTICATION</Property>

```

2. Add the following properties to the <QueryDescriptor> element:

```

    <Property NAME="UserCredential" SCOPE="GLOBAL"> UserCredentialSet
  </Property>
    <CredentialRef NAME="UserCredentialSet">UserCredentialSet</CredentialRef>

```

3. Define the credential type after the <Metric> tag:

```

.....
    <Property NAME="UserCredential" SCOPE="GLOBAL">UserCredentialSet
  </Property>
    <CredentialRef NAME="UserCredentialSet">UserCredentialSet
  </CredentialRef>
  </QueryDescriptor>
  </Metric>
  <CredentialInfo>
    <CredentialType NAME="AliasCredential">
      <Display>
        <Label NLSID="CRED_TYPE">Alias Credential Type</Label>
      </Display>
    </CredentialType>
  </CredentialInfo>

```

```

    <CredentialTypeColumn NAME="Alias">
      <Display>
        <Label NLSID="CRED_ALIAS">Alias (i.e. username, encryption key,
signature key, etc)</Label>
      </Display>
    </CredentialTypeColumn>
    <CredentialTypeColumn NAME="Password">
      <Display>
        <Label NLSID="CRED_PASSWORD">Password for the alias</Label>
      </Display>
    </CredentialTypeColumn>
  </CredentialType>
  <CredentialSet NAME="UserCredentialSet" USAGE="MONITORING">
    <AllowedCredType TYPE="AliasCredential"/>
  </CredentialSet>
</CredentialInfo>

```

Example 21–18 Using Keystore and Truststore for SSL

```

.....
  <Property NAME="SSLTrustStoreCredential"
SCOPE="GLOBAL">SSLTrustStoreCredentialSet</Property>
  <Property NAME="SSLKeyStoreCredential"
SCOPE="GLOBAL">SSLKeyStoreCredentialSet</Property>
NAME="SSLTrustStoreCredentialSet">SSLTrustStoreCredentialSet</CredentialRef>
  <CredentialRef
NAME="SSLKeyStoreCredentialSet">SSLKeyStoreCredentialSet</CredentialRef>
  </QueryDescriptor>
</Metric>
<CredentialInfo>
  <CredentialType NAME="StoreCredential">
    <Display>
      <Label NLSID="CRED_TYPE">Store Credential Type</Label>
    </Display>
    <CredentialTypeColumn NAME="StoreLocation">
      <Display>
        <Label NLSID="CRED_STORE_LOCATION">Store Location</Label>
      </Display>
    </CredentialTypeColumn>
    <CredentialTypeColumn NAME="StoreType">
      <Display>
        <Label NLSID="CRED_STORE_TYPE">Store Type</Label>
      </Display>
    </CredentialTypeColumn>
    <CredentialTypeColumn NAME="StorePassword">
      <Display>
        <Label NLSID="CRED_STORE_PASSWORD">Store Password</Label>
      </Display>
    </CredentialTypeColumn>
  </CredentialType>
  <CredentialSet NAME="SSLTrustStoreCredentialSet" USAGE="MONITORING">
    <AllowedCredType TYPE="StoreCredential"/>
  </CredentialSet>
  <CredentialSet NAME="SSLKeyStoreCredentialSet" USAGE="MONITORING">
    <AllowedCredType TYPE="StoreCredential"/>
  </CredentialSet>
</CredentialInfo>

```

21.13 WS-Management Fetchlet

In target metadata files generated by the *wsmancli* Command-Line Tool, the `<QueryDescriptor>` element specifies the properties that will be passed to the WSManagement fetchlet when being invoked.

Input Parameters

Table 21–19 provides a complete list of the supported properties:

Table 21–19 WS Management Fetchlet Properties

Name	Description	Use
ResourceURI	URI of a resource class representation or instance representation (wsman:ResourceURL)	Required Any valid URI according to RFC 3986
To	Transport address of a service (wsa:To).	Required Any valid network transport address.
Action	wsa:Action identifies which operation is to be carried out against the resource.	Required Current release only supports "http://schemas.xmlsoap.org/ws/2004/09/transfer/Get".
TransferOperation	Name of the WS-Transfer operation.	Required Current release only supports "GET".
Locale	Specifies the language that the client requests (and sometimes requires) and the response text to be translated into (wsman:Locale)	Optional Any valid value for the standard XML attribute xml:lang
MaxEnvelopeSize	The size to indicate that client expects a response whose total SOAP envelope does not exceed the specified number of octets (wsman:MaxEnvelopeSize)	Optional Value should not be less than 8192
OperationTimeout	The value to indicate that client expects a response or a fault within the specified time (wsman:OperationTimeout)	Optional Specify the value using format xs:duration (see http://www.w3.org/2001/XMLSchema:duration).
OptionSet	A set of switches to the service to modify or refine the nature of the request (wsman:OptionSet).	Optional Specify the values using the notation: [<OptionName1>, value:<value1>, type:<type1>, mustComply:<true false>][<OptionName2>, value:<value2>, type:<type>, mustComply:<true false>][...]

Table 21–19 (Cont.) WS Management Fetchlet Properties

Name	Description	Use
ReplyTo	The header to be present in all request messages when a reply is required (wsa:ReplyTo).	Optional It should be either a valid address for a new connection using any transport supported by the service or the URI <code>http://schemas.xmlsoap.org/ws/2004/08/addressing/role/anonymous</code> (see WS-Addressing)
SelectorSet	Set of selectors that identify the instance of resource to be accessed (wsman:SelectorSet)	Required Specify the value using the format below: <code>[S1, V1][S2, V2]...[Sn, Vn]</code> Where - S1, S2, ..., Sn are Selector names - V1, V2, ..., Vn are Selector values
SecurityPolicy	Security policy	Required - <i>NONE</i> - <i>BASIC_AUTHENTICATION</i>
Namespace	Set of all namespaces referenced	Required Specify using notation: <code>[ns0="uri0"][ns1="uri1"]..</code> Example: <code>[ns0="http://type.abc.com"][ns1="http://app.abc.com"]</code>
ColType	Collection result column type	Required List of metric column types (separated by comma) Example: <code>msgId:STRING,source:STRING,detail:STRING</code>
RowType	Collection result row type	Required List of XPath expression corresponding to metric columns (separated by comma) For example: <code>//ns0:eventResponse/msgId, //ns0:eventResponse/source</code>
SSLKeyStoreCredential	SSL keystore credentialSet name	Optional A valid CredentialSet of a StoreCredential Type defined in the <CredentialInfo>
SSLTrustStoreCredential	SSL truststore credentialSet name	Optional A valid CredentialSet of a StoreCredential Type defined in the <CredentialInfo> tag.

Table 21–19 (Cont.) WS Management Fetchlet Properties

Name	Description	Use
UserCredential	User token credentialSet name	Optional A valid CredentialSet of a AliasCredential or CSFKeyCredential Type defined in the <CredentialInfo> tag.
ValueWhenDown	Default response when target is down	Required (only for response metric). Not required for regular metric. For Response metric, when a target is down, this value (if specified) will be returned. A target is considered as down when the Fetchlet catches a ConnectionException

Examples

[Example 21–19](#) provides an example of a metric definition using the WS-Management fetchlet.

Example 21–19 Metric definition for using the WS-Management Fetchlet

```
<Metric NAME="trafficLight" TYPE="TABLE">
  <Display>
    <Label NLSID="NLSID_TRAFFIC_LIGHT">trafficLight</Label>
  </Display>
  <TableDescriptor>
    <ColumnDescriptor IS_KEY="YES" NAME="name" TYPE="STRING">
      <Display>
        <Label NLSID="COL_NAME">name</Label>
      </Display>
    </ColumnDescriptor>
    <ColumnDescriptor IS_KEY="FALSE" NAME="color" TYPE="STRING">
      <Display>
        <Label NLSID="COL_COLOR">color</Label>
      </Display>
    </ColumnDescriptor>
    <ColumnDescriptor IS_KEY="FALSE" NAME="x" TYPE="STRING">
      <Display>
        <Label NLSID="COL_X">x</Label>
      </Display>
    </ColumnDescriptor>
    <ColumnDescriptor IS_KEY="FALSE" NAME="y" TYPE="STRING">
      <Display>
        <Label NLSID="COL_Y">y</Label>
      </Display>
    </ColumnDescriptor>
  </TableDescriptor>
  <QueryDescriptor FETCHLET_ID="WSManagementFetchlet">
    <Property NAME="SecurityPolicy" SCOPE="INSTANCE">NONE</Property>
    <Property NAME="ResourceURL" SCOPE="INSTANCE">resourceURL</Property>
    <Property NAME="To" SCOPE="INSTANCE">To address</Property>
    <Property NAME="OptionSet" SCOPE="INSTANCE">optionSet</Property>
    <Property NAME="Locale" SCOPE="INSTANCE">locale</Property>
    <Property NAME="MaxEnvelopeSize" SCOPE="INSTANCE">maxEnvelopeSize</Property>
    <Property NAME="OperationTimeout"
SCOPE="INSTANCE">operationTimeout</Property>
    <Property NAME="Namespace" SCOPE="GLOBAL">
```

```

<![CDATA[ [ns1="http://schemas.wiseman.dev.java.net/traffic/1/light.xsd"]
           [ns0="http://www.w3.org/2001/XMLSchema"]
           [wsa="http://www.w3.org/2005/08/addressing"]
           [env="http://www.w3.org/2003/05/soap-envelope"]] ]></Property>
  <Property NAME="RowType"
SCOPE="GLOBAL"> //ns1:trafficlight/ns1:name, //ns1:trafficlight/ns1:color, //ns1:
  <Property NAME="ColType"
SCOPE="GLOBAL"> name:STRING, color:STRING, x:STRING, y:STRING </Property>
  <Property NAME="ReplyTo"
SCOPE="GLOBAL"> http://www.w3.org/2005/08/addressing/role/anonymous </Property>
  <Property NAME="Action"
SCOPE="GLOBAL"> http://schemas.xmlsoap.org/ws/2004/09/transfer/Get </Property>
  <Property NAME="TransferOperation" SCOPE="GLOBAL"> GET </Property>
  <Property NAME="SelectorSet" SCOPE="GLOBAL"> [name, Light1] </Property>
</QueryDescriptor>
</Metric>

```

21.13.1 Using Credentials

If basic authentication is required, then configure or define in the metric definition:

1. Set the SecurityPolicy property to BASIC_AUTHENTICATION:

```

<Property NAME="SecurityPolicy" SCOPE="INSTANCE">BASIC_
AUTHENTICATION</Property>

```

2. Add the following properties in the <QueryDescriptor> tag:

```

<Property NAME="UserCredential" SCOPE="GLOBAL"> UserCredentialSet </Property>
<CredentialRef NAME="UserCredentialSet">UserCredentialSet</CredentialRef>

```

3. Define the credential type after the <Metric> tag:

```

.....
  <Property NAME="UserCredential" SCOPE="GLOBAL">UserCredentialSet
</Property>
  <CredentialRef NAME="UserCredentialSet">UserCredentialSet
</CredentialRef>
</QueryDescriptor>
</Metric>
<CredentialInfo>
  <CredentialType NAME="AliasCredential">
    <Display>
      <Label NLSID="CRED_TYPE">Alias Credential Type</Label>
    </Display>
    <CredentialTypeColumn NAME="Alias">
      <Display>
        <Label NLSID="CRED_ALIAS">Alias (i.e. username, encryption key,
signature key, etc)</Label>
      </Display>
    </CredentialTypeColumn>
    <CredentialTypeColumn NAME="Password">
      <Display>
        <Label NLSID="CRED_PASSWORD">Password for the alias</Label>
      </Display>
    </CredentialTypeColumn>
  </CredentialType>
  <CredentialSet NAME="UserCredentialSet" USAGE="MONITORING">
    <AllowedCredType TYPE="AliasCredential"/>
  </CredentialSet>

```

```
</CredentialInfo>
```

Example 21–20 Using Keystore and Truststore for SSL

```
.....
<Property NAME="SSLTrustStoreCredential"
SCOPE="GLOBAL">SSLTrustStoreCredentialSet</Property>
  <Property NAME="SSLKeyStoreCredential"
SCOPE="GLOBAL">SSLKeyStoreCredentialSet</Property>
<CredentialRef
NAME="SSLTrustStoreCredentialSet">SSLTrustStoreCredentialSet</CredentialRef>
  <CredentialRef
NAME="SSLKeyStoreCredentialSet">SSLKeyStoreCredentialSet</CredentialRef>
</QueryDescriptor>
</Metric>
<CredentialInfo>
  <CredentialType NAME="StoreCredential">
    <Display>
      <Label NLSID="CRED_TYPE">Store Credential Type</Label>
    </Display>
    <CredentialTypeColumn NAME="StoreLocation">
      <Display>
        <Label NLSID="CRED_STORE_LOCATION">Store Location</Label>
      </Display>
    </CredentialTypeColumn>
    <CredentialTypeColumn NAME="StoreType">
      <Display>
        <Label NLSID="CRED_STORE_TYPE">Store Type</Label>
      </Display>
    </CredentialTypeColumn>
    <CredentialTypeColumn NAME="StorePassword">
      <Display>
        <Label NLSID="CRED_STORE_PASSWORD">Store Password</Label>
      </Display>
    </CredentialTypeColumn>
  </CredentialType>
  <CredentialSet NAME="SSLTrustStoreCredentialSet" USAGE="MONITORING">
    <AllowedCredType TYPE="StoreCredential"/>
  </CredentialSet>
  <CredentialSet NAME="SSLKeyStoreCredentialSet" USAGE="MONITORING">
    <AllowedCredType TYPE="StoreCredential"/>
  </CredentialSet>
</CredentialInfo>
```

21.14 REST Fetchlet

The REST fetchlet provides target monitoring for RESTful web resources. Based on input properties, this fetchlet can construct a request to communicate with the managed targets using HTTP standards. It can retrieve relevant data from the response to build and return monitoring metrics.

This release supports the following RESTful web services only:

- HTTP methods
 - GET: Define a reading access of the source without any side-effects. The resource is never changed through a GET request.
 - POST: Update an existing resource or create a new resource.

- HEAD: vCheck if a given path is serviceable.
- Media type of request or response representations
 - application/xml (both request and response)
 - application/json (both request and response)
 - text/xml (response only)
 - application/x-www-form-urlencoded (request only)
- Authentication scheme
 - Supports BASIC authentication.

21.14.1 Response Processing

The fetchlet relies on response data to construct monitoring metrics. Because the response media type can be application/xml, application/json, or text/xml, different mechanisms are adapted to process the response. [Table 21–20](#) describes the different mechanisms for each response media type.

Table 21–20 *Response Processing Mechanism*

Media Type	Mechanism
application/xml	<p>XPath Query is used for processing XML data.</p> <p>The fetchlet property, RowType, specifies a list of XPath expressions corresponding to metric columns (separated by comma) for retrieving column data.</p> <p>For example:</p> <pre><records> <ns2:Record xmlns:ns2="urn:com.office.directory"> <name>Peter</name> <phone>+1 (650) 555-0100</phone> </ns2:Record> <ns2:Record xmlns:ns2="urn:com.office.directory"> <name>John</name> <phone>+1 (650) 555-0185</phone> </ns2:Record> </records></pre> <p>Assume the monitoring metric has two columns (name and phone). The corresponding XPath expressions are:</p> <ul style="list-style-type: none">■ /records/Record/name■ /records/Record/phone <p>The following is an example of extracted data:</p> <pre>Peter, +1 (650) 555-0100 John, +1 (650) 555-0185</pre>

Table 21–20 (Cont.) Resonse Processing Mechanism

Media Type	Mechanism
application/json	<p>JSONPath is used for processing JavaScript Object Notation (JSON) data. JSONPath expressions refer to a JSON structure in the same way as XPath expressions are used in an XML document.</p> <p>For example:</p> <pre>{ "Record": [{ "name": "Peter" "phone": "+1 (650) 555-0100" }, { "name": "John" "phone": "+1 (650) 555-0185" },] }</pre> <p>Assume the monitoring metric has two columns (name and phone). The corresponding JSONPath expressions are:</p> <ul style="list-style-type: none"> ■ \$.Record.name ■ \$.Record.phone <p>The following example is an example of extracted data:</p> <pre>Peter, +1 (650) 555-0100 John, +1 (650) 555-0185</pre>
text/xml	<p>Because text is a non-structural representation, there is no way to extract any specific data from it. Instead, the entire response is returned.</p>

Input Parameters

[Table 21–21](#) provides a complete list of the supported properties.

Table 21–21 REST Fetchlet Properties

Name	Description	Optional
BaseURI	Base URI of the RESTful web service	No
RequestElementPayload	Request element payload (XML/JSON) in string format. Must be specified using the CDATA section if it is XML	Yes
RequestMetadata	Request metadata in XML format	No
SecurityPolicy	Specifies authentication scheme. Either NONE or BASIC_ AUTHENTICATION	No
Namespace	<p>Set of all namespaces referenced in the metric. Specify using notation: [ns0="uri0"][ns1="uri1"]...</p> <p>For example:</p> <pre>[ns0="http://type.abc.com"] [ns1="http://app.abc.com"]</pre>	No

Table 21–21 (Cont.) REST Fetchlet Properties

Name	Description	Optional
ColType	Collection result column type. List of metric column type (separated by comma). For example: <code>msgId:STRING,source:STRING,detail:STRING</code>	No
RowType	Collection result row type. List of path (XPath or JsonPath) expressions corresponding to metric columns (separated by comma). For example: <code>//ns0:eventResponse/msgId, //ns0:eventResponse/source, //ns0:eventResponse/detail</code>	No
SSLKeyStoreCredential	SSL keystore credential set name. It must be defined as a monitoring credential and contain these credential columns: Location, Type, Password	Yes
SSLTrustStoreCredential	SSL truststore credentialset name. It must be defined as a monitoring credential and contain these credential columns: Location, Type, Password	Yes
UserCredential	User token credentialset name. It must be defined as a monitoring credential and contain these credential columns: Alias, Password	Yes
ProxyHost	Host name of the proxy server to make the URL connection	Yes
ProxyPort	Port number of the proxy server to make the URL connection	Yes

[Example 21–21](#) shows an example of the Fetchlet Query Descriptor from a target metadata file. For more information about the target metadata files, see [Chapter 3, "Creating Target Metadata Files"](#).

Note: The fetchlet ID is RESTFetchlet.

Example 21–21 Fetchlet Query Descriptor

```
<QueryDescriptor FETCHLET_ID="RESTFetchlet">
  <Property NAME="SecurityPolicy"
SCOPE="INSTANCE">ListAll.SecurityPolicy</Property>
  <Property NAME="BaseURI" SCOPE="INSTANCE">ListAll.BaseURI</Property>
  <Property NAME="Namespace" SCOPE="GLOBAL">
    <![CDATA[ [ns0="urn:com.office.directory"] ] ]></Property>
  <Property NAME="RowType" SCOPE="GLOBAL">//ns0:Record/name, //ns0:Record/title,
//ns0:Record/phone, //ns0:Record/building, //ns0:Record/floor,
//ns0:Record/office</Property>
  <Property NAME="ColType" SCOPE="GLOBAL">name:STRING, title:STRING, phone:STRING
, building:STRING, floor:STRING, office:STRING</Property>
  <Property NAME="RequestMetadata" SCOPE="GLOBAL">
    <![CDATA[<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
      <Resource path="/">
        <resource path="lookup/list">
```

```
        <method elementDefined="false"
            accept="application/xml" name="GET"/>
    </resource>
</Resource>
]]></Property>
<Property NAME="UserCredential" SCOPE="GLOBAL">UserCredentialSet</Property>
<CredentialRef NAME="UserCredentialSet">UserCredentialSet</CredentialRef>
</QueryDescriptor>
```

21.14.2 Using HTTPS and Self-Signed Certificates

When calling an HTTPS URL with a self-signed SSL certificate from a REST fetchlet, the credential set must be specified in the target metadata file.

Example 21–22

```
<QueryDescriptor FETCHLET_ID="RESTFetchlet">
    .....
    <Property NAME="SSLTrustStoreCredential" SCOPE="GLOBAL">
        SSLTrustStoreCredentialSet</Property>
    <CredentialRef NAME="SSLTrustStoreCredentialSet">
        SSLTrustStoreCredentialSet</CredentialRef>
</QueryDescriptor>
<CredentialInfo>
    <CredentialType NAME="StoreCredential">
        <Display>
            <Label NLSID="CRED_TYPE">Store Credential Type</Label>
        </Display>
        <CredentialTypeColumn NAME="StoreLocation" IS_SENSITIVE="FALSE">
            <Display>
                <Label NLSID="CRED_STORE_LOCATION">Store Location</Label>
            </Display>
        </CredentialTypeColumn>
        <CredentialTypeColumn NAME="StoreType" IS_SENSITIVE="FALSE">
            <Display>
                <Label NLSID="CRED_STORE_TYPE">Store Type</Label>
            </Display>
        </CredentialTypeColumn>
        <CredentialTypeColumn NAME="StorePassword">
            <Display>
                <Label NLSID="CRED_STORE_PASSWORD">Store Password</Label>
            </Display>
        </CredentialTypeColumn>
    </CredentialType>
    <CredentialSet NAME="SSLTrustStoreCredentialSet" USAGE="MONITORING" CONTEXT_
TYPE="TARGET">
        <AllowedCredType TYPE="StoreCredential"/>
    </CredentialSet>
</CredentialInfo>
```

In Enterprise Manager Cloud Control Release 2 (12.1.0.3), a new fetchlet property, “SSLTrustServerCert”, has been added. If set to “TRUE”, the fetchlet will use the non-validating mode for the server certificate, and there is no need to provide or specify the SSL trust store.

Example 21–23

```
<Property NAME="SSLTrustServerCert" SCOPE="GLOBAL">TRUE</Property>
```


21.14.3 Using REST CLI to Generate Metadata

REST CLI is a client command line tool for generating target metadata and default collection files to enable the Management Agent to monitor RESTful web resources through invoking the REST fetchlet.

Use the following `emctl` command to invoke REST CLI:

```
emctl restcli
```

[Table 21–22](#) provides a list of the command-line arguments that you can use with the `emctl restcli` command.

Table 21–22 Command-line Arguments Supported by REST CLI

Argument	Description	Example
metadata	Generate target metadata	-metadata
wadl	WADL location	-wadl=http://..... -wadl=file:///.....
wsdl	WSDL location	-wsdl=http://..... -wsdl=file:///.....
username	User name to log in to the host	-username=admin
proxyhost	Host name of the proxy server	-proxyhost=proxy.example.com
proxyport	Port number of the proxy server	-proxyport=80

To use REST CLI:

1. Run the REST CLI command with the Web Application Description Language (WADL) location. For example:

```
emctl restcli
-wadl=http://host.us.example.com:17382/OfficeDirectoryBA/application.wadl
```

If the WADL location is access protected, then enter a user name and password.

[Example 21–24](#) provides an example of a user running the REST CLI tool.

2. REST CLI prompts you to enter the target type and location where the output directory will contain the generated target and collection metadata files.
3. REST CLI lists out all the available resources paths for monitoring. You must select a resource path and one of its methods to define monitoring metric for that resource.
4. REST CLI also prompts you to define the collection schedule.

When all the information is gathered from the user, the tool generates the target and default collection metadata files under the specified output directory similar to the metadata provided in [Example 21–25](#).

Example 21–24 Running REST CLI

Generate Metric Metadata for REST Web Resource Monitoring

Enter password for "weblogic" :

```
Reading WADL Document at
http://host.us.example.com:17382/OfficeDirectoryBA/application.wadl...done.
```

```
==> Enter the name for this target type : OfficeDirectory

==> Enter metadata file name [/scratch/work/metadata/OfficeDirectory.xml] :

All resource paths available for monitoring :
[1]   /add
[2]   /lookup/list
[3]   /lookup/phone
[4]   /lookup/building/people
[5]   /db/count

==> Enter the index [1-5] to select : 2
* Selected Resource Path : /lookup/list

All methods available from the selected path :
[1]   application/xml[Record] GET()
[2]   application/json[Record] GET()

==> Enter the index [1-2] to select : 1
* Selected Resource Method: application/xml[Record] GET()

Define new metric group :
==> Enter the name for this metric group [GET] : ListGet_XML

Return value(s) for the selected method :
[1]   //ns0:Record/name <string>
[2]   //ns0:Record/title <string>
[3]   //ns0:Record/phone <string>
[4]   //ns0:Record/building <string>
[5]   //ns0:Record/floor <string>
[6]   //ns0:Record/office <string>

==> Enter the index [1-6] of metric to display : 1
==> Enter the name for this metric [name] :
==> Enter the label for this metric [name] :
==> Is this a key metric <y/n>? [n] : y
==> Do you want to add another metric <y/n>? [y] :

Return value(s) for the selected method :
[1]   //ns0:Record/title <string>
[2]   //ns0:Record/phone <string>
[3]   //ns0:Record/building <string>
[4]   //ns0:Record/floor <string>
[5]   //ns0:Record/office <string>

==> Enter the index [1-5] of metric to display : 1
==> Enter the name for this metric [title] :
==> Enter the label for this metric [title] :
==> Is this a key metric <y/n>? [n] :
==> Do you want to create threshold for it <y/n>? [n] :
==> Do you want to add another metric <y/n>? [y] :

Return value(s) for the selected method :
[1]   //ns0:Record/phone <string>
[2]   //ns0:Record/building <string>
[3]   //ns0:Record/floor <string>
[4]   //ns0:Record/office <string>

==> Enter the index [1-4] of metric to display : 1
```

```

==> Enter the name for this metric [phone] :
==> Enter the label for this metric [phone] :
==> Is this a key metric <y/n>? [n] :
==> Do you want to create threshold for it <y/n>? [n] :
==> Do you want to add another metric <y/n>? [y] :

Return value(s) for the selected method :
[1] //ns0:Record/building <string>
[2] //ns0:Record/floor <string>
[3] //ns0:Record/office <string>

==> Enter the index [1-3] of metric to display : 1
==> Enter the name for this metric [building] :
==> Enter the label for this metric [building] :
==> Is this a key metric <y/n>? [n] :
==> Do you want to create threshold for it <y/n>? [n] :
==> Do you want to add another metric <y/n>? [y] :

Return value(s) for the selected method :
[1] //ns0:Record/floor <string>
[2] //ns0:Record/office <string>

==> Enter the index [1-2] of metric to display : 1
==> Enter the name for this metric [floor] :
==> Enter the label for this metric [floor] :
==> Is this a key metric <y/n>? [n] :
==> Do you want to create threshold for it <y/n>? [n] :
==> Do you want to add another metric <y/n>? [y] :

Return value(s) for the selected method :
[1] //ns0:Record/office <string>

==> Enter the index [1-1] of metric to display : 1
==> Enter the name for this metric [office] :
==> Enter the label for this metric [office] :
==> Is this a key metric <y/n>? [n] :
==> Do you want to create threshold for it <y/n>? [n] :

Setup request parameters

==> Do you want to add User/Password Credential <y/n>? [n] : y

==> Do you want to add SSL TrustStore Credential <y/n>? [n] :

==> Do you want to add SSL KeyStore Credential <y/n>? [n] :

==> Is this metric group for periodic collection <y/n>? [y] :
The following units are for collection frequency:
[1] Min
[2] Hr
[3] Day

==> Enter the index [1-3] of unit for this collection : 1
==> Enter the frequency of collection in Min : 5

==> Do you want to add another metric group <y/n>? [n] : y

All resource paths available for monitoring :
[1] /add

```

```
[2]    /lookup/list
[3]    /lookup/phone
[4]    /lookup/building/people
[5]    /db/count

==> Enter the index [1-5] to select : 3
* Selected Resource Path : /lookup/phone

All methods available from the selected path :
[1]    application/json[PhoneInfo] GET(name)
[2]    application/xml[PhoneInfo] GET(name)

==> Enter the index [1-2] to select : 1
* Selected Resource Method: application/json[PhoneInfo] GET(name)

Define new metric group :
==> Enter the name for this metric group [GET] : LookupGet_JSON

Return value(s) for the selected method :
[1]    $..name <string>
[2]    $..phone <string>

==> Enter the index [1-2] of metric to display : 1
==> Enter the name for this metric [name] :
==> Enter the label for this metric [name] :
==> Is this a key metric <y/n>? [n] : y
==> Do you want to add another metric <y/n>? [y] :

Return value(s) for the selected method :
[1]    $..phone <string>

==> Enter the index [1-1] of metric to display : 1
==> Enter the name for this metric [phone] :
==> Enter the label for this metric [phone] :
==> Is this a key metric <y/n>? [n] :
==> Do you want to create threshold for it <y/n>? [n] :

Setup request parameters
==> Enter value for query parameter "name" [%LookupGet_JSON.name0000%] : Harry
Smith

==> Do you want to add User/Password Credential <y/n>? [n] : y

==> Do you want to add SSL TrustStore Credential <y/n>? [n] :

==> Do you want to add SSL KeyStore Credential <y/n>? [n] :

==> Is this metric group for periodic collection <y/n>? [y] :
The following units are for collection frequency:
[1]    Min
[2]    Hr
[3]    Day

==> Enter the index [1-3] of unit for this collection : 1
==> Enter the frequency of collection in Min : 5

==> Do you want to add another metric group <y/n>? [n] :

Files Generated:
```

- Target Metadata file: /scratch/work/metadata/OfficeDirectory.xml
- Target Collection file: /scratch/work/metadata/OfficeDirectoryCollection.xml

Example 21–25 REST CLI-Generated Target Metadata

```
<TargetMetadata META_VER="1.0" TYPE="OfficeDirectory">
  <Display>
    <Label NLSID="NLSID_OFFICE_DIRECTORY">OfficeDirectory</Label>
    <ShortName NLSID="NLSID_OFFICE_DIRECTORY">OfficeDirectory</ShortName>
    <Description NLSID="NLSID_OFFICE_DIRECTORY">OfficeDirectory</Description>
  </Display>
  <Metric NAME="ListGet_XML" TYPE="TABLE">
    <Display>
      <Label NLSID="NLSID_LIST_GET_XML">ListGet_XML</Label>
    </Display>
    <TableDescriptor>
      <ColumnDescriptor IS_KEY="TRUE" NAME="name" TYPE="STRING">
        <Display>
          <Label NLSID="COL_NAME">name</Label>
        </Display>
      </ColumnDescriptor>
      <ColumnDescriptor IS_KEY="FALSE" NAME="title" TYPE="STRING">
        <Display>
          <Label NLSID="COL_TITLE">title</Label>
        </Display>
      </ColumnDescriptor>
      <ColumnDescriptor IS_KEY="FALSE" NAME="phone" TYPE="STRING">
        <Display>
          <Label NLSID="COL_PHONE">phone</Label>
        </Display>
      </ColumnDescriptor>
      <ColumnDescriptor IS_KEY="FALSE" NAME="building" TYPE="STRING">
        <Display>
          <Label NLSID="COL_BUILDING">building</Label>
        </Display>
      </ColumnDescriptor>
      <ColumnDescriptor IS_KEY="FALSE" NAME="floor" TYPE="STRING">
        <Display>
          <Label NLSID="COL_FLOOR">floor</Label>
        </Display>
      </ColumnDescriptor>
      <ColumnDescriptor IS_KEY="FALSE" NAME="office" TYPE="STRING">
        <Display>
          <Label NLSID="COL_OFFICE">office</Label>
        </Display>
      </ColumnDescriptor>
    </TableDescriptor>
    <QueryDescriptor FETCHLET_ID="JAXRS_Fetchlet">
      <Property NAME="ProxyHost" SCOPE="INSTANCE"
OPTIONAL="TRUE">ProxyHost</Property>
      <Property NAME="ProxyPort" SCOPE="INSTANCE"
OPTIONAL="TRUE">ProxyPort</Property>
      <Property NAME="SecurityPolicy" SCOPE="INSTANCE" OPTIONAL="FALSE">ListGet_
XML.SecurityPolicy</Property>
      <Property NAME="BaseURI" SCOPE="INSTANCE" OPTIONAL="FALSE">ListGet_
XML.BaseURI</Property>
      <Property NAME="Namespace" SCOPE="GLOBAL"
OPTIONAL="FALSE"><![CDATA[ [ns0="urn:com.office.directory"] ]]></Property>
      <Property NAME="RowType" SCOPE="GLOBAL"
```

```

OPTIONAL="FALSE">//ns0:Record/name, //ns0:Record/title, //ns0:Record/phone,
//ns0:Record/building, //ns0:Record/floor, //ns0:Record/office</Property>
    <Property NAME="ColType" SCOPE="GLOBAL"
OPTIONAL="FALSE">name:STRING, title:STRING, phone:STRING, building:STRING,
floor:STRING, office:STRING</Property>
    <Property NAME="RequestMetadata" SCOPE="GLOBAL"
OPTIONAL="FALSE"><![CDATA[<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<Resource path="/">
    <resource path="lookup/list">
        <path>lookup/list</path>
        <method elementDefined="false" accept="application/xml" name="GET"/>
    </resource>
</Resource>
]]></Property>
    <Property NAME="UserCredential" SCOPE="GLOBAL"
OPTIONAL="FALSE">UserCredentialSet</Property>
    <CredentialRef NAME="UserCredentialSet">UserCredentialSet</CredentialRef>
</QueryDescriptor>
</Metric>
<Metric NAME="LookupGet_JSON" TYPE="TABLE">
    <Display>
        <Label NLSID="NLSID_LOOKUP_GET_JSON">LookupGet_JSON</Label>
    </Display>
    <TableDescriptor>
        <ColumnDescriptor IS_KEY="TRUE" NAME="name" TYPE="STRING">
            <Display>
                <Label NLSID="COL_NAME">name</Label>
            </Display>
        </ColumnDescriptor>
        <ColumnDescriptor IS_KEY="FALSE" NAME="phone" TYPE="STRING">
            <Display>
                <Label NLSID="COL_PHONE">phone</Label>
            </Display>
        </ColumnDescriptor>
    </TableDescriptor>
    <QueryDescriptor FETCHLET_ID="JAXRS_Fetchlet">
        <Property NAME="ProxyHost" SCOPE="INSTANCE"
OPTIONAL="TRUE">ProxyHost</Property>
        <Property NAME="ProxyPort" SCOPE="INSTANCE"
OPTIONAL="TRUE">ProxyPort</Property>
        <Property NAME="SecurityPolicy" SCOPE="INSTANCE" OPTIONAL="FALSE">LookupGet_
JSON.SecurityPolicy</Property>
        <Property NAME="BaseURI" SCOPE="INSTANCE" OPTIONAL="FALSE">LookupGet_
JSON.BaseURI</Property>
        <Property NAME="Namespace" SCOPE="GLOBAL"
OPTIONAL="FALSE"><![CDATA[ [ns0="urn:com.office.directory"] ]]></Property>
        <Property NAME="RowType" SCOPE="GLOBAL"
OPTIONAL="FALSE">$.name, $.phone</Property>
        <Property NAME="ColType" SCOPE="GLOBAL"
OPTIONAL="FALSE">name:STRING, phone:STRING</Property>
        <Property NAME="RequestMetadata" SCOPE="GLOBAL"
OPTIONAL="FALSE"><![CDATA[<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
<Resource path="/">
    <resource path="lookup/phone">
        <method elementDefined="false" accept="application/xml" name="GET">
            <Parameter style="query" value="Harry Son" name="name"/>
        </method>
    </resource>
</Resource>
]]></Property>

```

```

        <Property NAME="UserCredential" SCOPE="GLOBAL"
OPTIONAL="FALSE">UserCredentialSet</Property>
        <CredentialRef NAME="UserCredentialSet">UserCredentialSet</CredentialRef>
    </QueryDescriptor>
</Metric>
<CredentialInfo>
    <CredentialType NAME="CSFKeyCredential">
        <Display>
            <Label NLSID="CRED_TYPE">CSF-Key Credential Type</Label>
        </Display>
        <CredentialTypeColumn NAME="CSFKey">
            <Display>
                <Label NLSID="CRED_CSFKEY">Alias CSF Key</Label>
            </Display>
        </CredentialTypeColumn>
    </CredentialType>
    <CredentialType NAME="AliasCredential">
        <Display>
            <Label NLSID="CRED_TYPE">Alias Credential Type</Label>
        </Display>
        <CredentialTypeColumn NAME="Alias">
            <Display>
                <Label NLSID="CRED_ALIAS">Alias (i.e. username, encryption key,
signature key, etc)</Label>
            </Display>
        </CredentialTypeColumn>
        <CredentialTypeColumn NAME="Password">
            <Display>
                <Label NLSID="CRED_PASSWORD">Password for the alias</Label>
            </Display>
        </CredentialTypeColumn>
    </CredentialType>
    <CredentialSet NAME="UserCredentialSet" USAGE="MONITORING">
        <AllowedCredType TYPE="CSFKeyCredential"/>
        <AllowedCredType TYPE="AliasCredential"/>
    </CredentialSet>
</CredentialInfo>
<InstanceProperties>
    <InstanceProperty NAME="ProxyHost" CREDENTIAL="FALSE" OPTIONAL="TRUE">
        <Display>
            <Label NLSID="PROP_PROXY_HOST">Proxy Server Name</Label>
        </Display>
    </InstanceProperty>
    <InstanceProperty NAME="ProxyPort" CREDENTIAL="FALSE" OPTIONAL="TRUE">
        <Display>
            <Label NLSID="PROP_PROXY_PORT">Proxy Server Port</Label>
        </Display>
    </InstanceProperty>
    <InstanceProperty NAME="ListGet_XML.SecurityPolicy"
        CREDENTIAL="FALSE" OPTIONAL="FALSE">
        <Display>
            <Label NLSID="PROP_LIST_GET_XML_SECURITY_POLICY">[ListGet_XML]
Authentication/Web Service Policy</Label>
        </Display>
    </InstanceProperty>
    <InstanceProperty NAME="ListGet_XML.BaseURI" CREDENTIAL="FALSE"
OPTIONAL="FALSE">
        <Display>
            <Label NLSID="PROP_LIST_GET_XML_BASE_URI">[ListGet_XML] Resource Base
URI</Label>

```

```
        </Display>
    </InstanceProperty>
    <InstanceProperty NAME="LookupGet_JSON.SecurityPolicy"
        CREDENTIAL="FALSE" OPTIONAL="FALSE">
        <Display>
            <Label NLSID="PROP_LOOKUP_GET_JSON_SECURITY_POLICY">[LookupGet_JSON]
Authentication/Web Service Policy</Label>
        </Display>
    </InstanceProperty>
    <InstanceProperty NAME="LookupGet_JSON.BaseURI" CREDENTIAL="FALSE"
OPTIONAL="FALSE">
        <Display>
            <Label NLSID="PROP_LOOKUP_GET_JSON_BASE_URI">[LookupGet_JSON] Resource
Base URI</Label>
        </Display>
    </InstanceProperty>
</InstanceProperties>
</TargetMetadata>
```

Out-of-Box Associations

Enterprise Manager provides a common set of association types that should meet the needs of most integrators. As an integrator you are encouraged to become familiar with these association types and use them if applicable. As an integrator, you should also update the Table of Integrators and Documents with links to the documents describing your association types and your usage of all association types (allowed_pairs).

The following tables provide details on the out-of-box associations:

- [Table A-1, "application_contains"](#)
- [Table A-2, "app_composite_contains"](#)
- [Table A-3, "authenticated_by"](#)
- [Table A-4, "composite_contains \(abstract\)"](#)
- [Table A-5, "cluster_contains"](#)
- [Table A-6, "connects_through"](#)
- [Table A-7, "contains \(abstract\)"](#)
- [Table A-8, "depends_on\(abstract\)"](#)
- [Table A-9, "deployed_on"](#)
- [Table A-10, "exposes"](#)
- [Table A-11, "hosted_by"](#)
- [Table A-12, "installed_at"](#)
- [Table A-13, "internal_contains \(for internal OMS use only\)"](#)
- [Table A-14, "managed_by"](#)
- [Table A-15, "monitored_by"](#)
- [Table A-16, "provided_by"](#)
- [Table A-17, "runs_on \(abstract\)"](#)
- [Table A-18, "stores_on"](#)
- [Table A-19, "stores_on_db"](#)
- [Table A-20, "uses \(abstract\)"](#)

Table A–1 *application_contains*

Basic Details	Source/Destination	Description	Usage
extends: Contains Core/Extended: Core Display Name: "contains (app component)" Inverse:member _of_application	Source: any aggregate type which is not kind of composite, for example oracle_ emrep. For the source type of composite type, user should use composite_contains or assoc types extended from it. Destination: member entity types, for example, j2ee_ applications for Enterprise Manager console and backend service are members for oracle_emrep Cardinality: 0..*	Capture the membership between application and its members. The source is an aggregation of the members. One member can be part of multiple aggregations.	Indicate a membership of an application, Can be used in the topology.

Table A–2 *app_composite_contains*

Basic Details	Source/Destination	Description	Usage
extends: application_ contains, composite_ contains Core/Extended: Core Display Name: "contains (app_ composite component)" Inverse:member _of_composite_ app	Source: any applicaion that is also a composite type, for example siebel_server. For the source type of cluster type, user should use cluster_ contains Destination: member entity types, for example, siebel_ component_group is a member of siebel_ server Cardinality: 0..*	Capture the membership between application and its members. The membership is also a kind of composition. One member can be part of only one composition	Indicate a membership of an application, Can be used in the topology.

Table A–3 *authenticated_by*

Basic Details	Source/Destination	Description	Usage
extends: depends_on Core/Extended: Core Display Name: "authenticated_ by)" Inverse:authenti cates	Source: An ME that requires authentication Destination: Me providing the Authentication, for example: oracle_ldap Cardinality: 0..1	Capture the membership between application and its members. The membership is also a kind of composition. One member can be part of only one composition	Indicate a membership of an application, Can be used in the topology.

Table A-4 *composite_contains (abstract)*

Basic Details	Source/Destination	Description	Usage
extends: contains,uses Core/Extended: Core Display Name: "CompositeContains" Inverse: member_of_ composite	Source: Source: any composite types Destination: Members of the composite. Cardinality: 0..*	<p>A form of aggregation which requires that a part instance be included in at most one composite at a time. For example: If a database D1 is part of rac cluster R1, it cannot be part of another cluster R2. This is used to place a box around the source and all its members to indicate that the members cannot be part of another source.</p> <p>The restriction applies to specific concrete association type extended from composite_contains. An ME can be a destination of no more than 1 assoc of type T if T extends composite_contains. But an ME can be destination with different source ME as long as the concrete composite_contains types are different.</p>	Current usage is to link cluster targets (rac_database,cluster,weblogic_cluster) to its members. Framework functionality, such as topology viewer, can use it.

Table A-5 *cluster_contains*

Basic Details	Source/Destination	Description	Usage
extends: composite_ contains Core/Extended: Core Display Name: "contains (in cluster)" Inverse:member _of_cluster	Source: A cluster target type, such as RAC or Cluster Destination: Member type of the cluster. The cluster member types should be the same Cardinality: 1..*	<p>Cluster membership, where the members are of the same type and together provide scalability and redundancy. Also indicates composite containment; Cluster A cluster_contains B implies that B cannot be member of another cluster C.</p>	Tools like Consolidation Planner need to know cluster membership. Also, all the general tools such as Topology Viewer.

Table A–6 connects_through

Basic Details	Source/Destination	Description	Usage
extends: depends_on Core/Extended: Core Display Name: "connects_ through" Inverse: connects	Source: An ME which is connecting to another ME via a intermediate path. Example: an application connecting to database via listener Destination: Me providing the access point for another ME (oracle_listener, oracle_ apache, slb) Cardinality: 0..*	Application connects_ through Listener, which exposes database. Service connects_through oracle_apache, which exposes oracle_oc4j.	Used in Applications topology to represent the connection to the end via an intermediate path. Example: connects_through listener, which exposes database. The source will also have a direct functional dependency directly on the end point (application stores_on_ db database) In functional view, the direct dependency of stores_on_db will be shown, in physical view, the listener link will be shown.

Table A–7 contains (abstract)

Basic Details	Source/Destination	Description	Usage
extends: none Core/Extended: Core Display Name: "contains" Inverse: member_of	Any source ME and its member ME types Cardinality: 0..*	Indicates containment membership. A contains B implies that B is one of the parts that make up A. All containment relationships should be captured by integrators except that system membership is captured by the OMS and TC containment is not required to be represented via an assoc instance. Integrator needs to use the concreted types, which extend from "contains"	However, user can query instances for all the association types which extend from the "constains". Framework functionality, such as topology viewer, can query this type of associations.

Table A–8 depends_on(abstract)

Basic Details	Source/Destination	Description	Usage
extends: uses Core/Extended: Core Display Name: "Depends on" Inverse: depended_on_ by	Source: any ME type Destination: any ME type Cardinality: 0..*	For any ME A that depends on ME B for its availability. If ME B is not available, the availability of A may be impacted.	Framework functionalities, such as RCA and Topology Viewer, can use it.

Table A–9 *deployed_on*

Basic Details	Source/Destination	Description	Usage
extends: runs_on, member_of_application Core/Extended: Core Display Name: "Deployed on" Inverse: deploys	Source: any ME except target component Destination: j2ee container Cardinality: 1..*	Application A is deployed into a J2EE? container B.	Topology viewer can display the application deployed on a j2ee server .

Table A–10 *exposes*

Basic Details	Source/Destination	Description	Usage
extends: uses Core/Extended: Core Display Name: "exposes" Inverse: exposed_by	Source: ME providing the access point for another ME (oracle_listener, oracle_apache, slb) Destination: ME of target being accessed via the source. Cardinality: 0..*	Some ME's functionality is exposed through other ME, such as oracle_listener exposes oracle_database to application, oracle_apache expose oc4j.Capture more semantics for uses: What the listener can do gets impacted if the db goes down, but listener itself does not go down, it is degraded mode. Used to represent the entry points for targets/systems. Listener exposes oracle_database. In this case it is strictly not the database target which is providing the accesspoint, the oracle database can be thought of providing the services which are available via the listener. oracle_http exposes oc4j. Here again, the oc4j target may not be providing the http service, the http service could be running as a seperate service which lets an application connect to oc4j.	Framework functionality, such as topology viewer, can use it.

Table A-11 *hosted_by*

Basic Details	Source/Destination	Description	Usage
extends: runs_on Core/Extended: Core Display Name: "hosted_by" Inverse: host_for	Source: any ME type except a system, service, or target component Destination: host Cardinality: 1	For any target T that is hosted_by H, the process(es) that comprise T execute on host H. A target can be hosted by only one host.	Used to locate the targets running on the given host.

Table A-12 *installed_at*

Basic Details	Source/Destination	Description	Usage
extends: uses Core/Extended: Core Display Name: "installed_at" Inverse: install_home_for	Source: any ME except target component Destination: A ME representing an install home Cardinality: 1	A installed at B indicates that B represents the install home for A Example: oracle database --> installed_at --> oracle_home.	Used to denote the link to the install home where the software for the ME is installed. Used in patching to get to the oracle home where the target is installed.

Table A-13 *internal_contains (for internal OMS use only)*

Basic Details	Source/Destination	Description	Usage
extends: contains Core/Extended: Core Display Name: "InternalContains" Inverse: internal_member_of	Source: System/Group Destination: A system can contain any ME except group, a group can contain any ME except target component Cardinality: 0..*	A special form of an association that specifies a whole-part relationship between the system and a component part. The component part can exist independent of the system and can be part of multiple systems. System A contains B implies that B is one of the parts that make up A. B can also be included in other system C	User should normally use Group/System API to find the members. But use can query the 'internal_contains' against table/view.

Table A-14 *managed_by*

Basic Details	Source/Destination	Description	Usage
extends: monitored_by, uses Core/Extended: Core Display Name: "managed_by" Inverse: manages	Source: any ME type except target component Destination: an ME type that can provide management functionality for other ME types. For example, oracle_cs can manage oracle_ database, oracle_ listener etc. Cardinality: 1..* (A specific allowed_pair can have stricter cardinality, such as 1)	The destination ME may work as watchdog and can start source target. The entity which manages another entities can make change to the managed entities, which the monitored_by doesn't have this semantic.	Framework functionality, such as topology viewer, can use the association

Table A-15 *monitored_by*

Basic Details	Source/Destination	Description	Usage
extends: Core/Extended: Core Display Name: "monitored_by" Inverse: monitors	Source: any ME type except target component Destination: agent Cardinality: 1 (cardinality is 1 at any given moment)	For any target T that is monitored by an agent. Example: target T--> monitored_by -> agent A .	Used in agent synchrononization/avail ability calculations/framework k code

Table A-16 *provided_by*

Basic Details	Source/Destination	Description	Usage
extends: depends_on Core/Extended: Core Display Name: "provided_by" Inverse: provides	Source: an ME, typically representing some kind of service, such as the db service, fusion product, webservice Destination: A system or a target which is providing the service Cardinality: 1	A provided_by B indicates that the service of B is provided by ME A.	Topology viewer can display the association

Table A-17 *runs_on (abstract)*

Basic Details	Source/Destination	Description	Usage
extends: depends_on Core/Extended: Core Display Name: "runs_on" Inverse: runs	Source: any ME except target component Destination: any ME which provides infrastructure for other MEs to run, such as VM. Cardinality: 1	A run_on B indicates that B provides infrastructure for some processes of A to execute. Note: processes is used in the English sense and does not indicate OS processes.	Framework functionality, such as topology viewer, can use it.

Table A–18 stores_on

Basic Details	Source/Destination	Description	Usage
extends: depends_on Core/Extended: Core Display Name: "stores_on" Inverse: stores	Source: typically a target which stores data. For example: oracle_database or sql server. Destination: an ME representing storage. For example: netapp filer or exadata. Cardinality: 0..*	Indicates the link to the target representing the storage of the bits. A stored_on B indicates that B provides infrastructure for storage of bits of A, Example: datafile-->stored_on-->netapp_filer. The stored data can be static or updatable.	Used to denote the link to the storage infrastructure. This is to visually locate the storage details in the topology.

Table A–19 stores_on_db

Basic Details	Source/Destination	Description	Usage
extends: stores_on Core/Extended: Core Display Name: "Data Repository" Inverse: data_repository_for	Source: an ME that stores data in a database. Destination: ME providing the database repository for storing the data Example: application stores_on_db oracle_database. Cardinality: 0..*	Represents depends_on in that if the database server is down, the source can be down.	Used in Applications topology to represent the database where an applications data is stored.

Table A–20 uses (abstract)

Basic Details	Source/Destination	Description	Usage
extends: None Core/Extended: Core Display Name: "uses" Inverse: used_by	Source: any ME type Destination: any ME type Cardinality: 0..*	For any ME A that depends on ME B for its working but does not affect availability.	Used in topology pages to indicate non availability dependency

A

- accessibility guidelines, 8-107
- accessing Enterprise Manager data, 8-37
- ActionScript, 8-8
- activity content parameter, 8-20
- ActivityController base class, 8-22
- adding a target instance, 13-11
- Adding Reports, 4-1
- adding targets manually, 11-8
- Adobe Flash Builder, 8-97
- Adobe Flex SDK, 8-96
- advanced metric collection, defining, 3-17
- advanced metrics, defining, 3-11
- advanced plug-in, creating, 1-5
- Anonymous PL/SQL Block, specifying, 4-10
- ANT_HOME environment variable, 8-96
- ANT_OPTS environment variable, 8-96
- Apache Ant, 8-95
- application activities, defining, 8-21
- application deployment views, 19-3
- appModel property, 8-18
- area charts, 8-72
- asynchronous service request handling, 8-6
- automatic discovery, configuring, 11-12
- automatic target discovery, 11-1
- automation services
 - about, 8-54
 - running jobs, 8-55
 - submitting jobs, 8-55
- availability region, 8-69
- Availability Status Icon, 4-13

B

- bar charts, 8-73
- basic metric collection, defining, 3-17
- basic plug-in, creating, 1-4
- basic response metric group, defining, 3-9
- blackout views, 19-9
- BulkSqlQuery interface, 8-47

C

- Chart Element, 4-21
- Chart Title, 4-24

- Chart Type, 4-22
- charts
 - area charts, 8-72
 - bar charts, 8-73
 - column charts, 8-74
 - defining, 8-70
 - horizontal charts, 8-73
 - line charts, 8-70
 - pie charts, 8-75
 - vertical bar charts, 8-74
- checking job status, 8-58
- collected configuration data, 6-2
- collecting target configuration data, 6-1
- column charts, 8-74
- Column Group End Column, 4-12
- Column Group Header, 4-12
- Column Group Start Column, 4-12
- columns, transient, 3-21
- commands
 - emcli add_target, 13-11
 - emcli import_update, 13-8
 - emctl register oms metadata, 13-12, 13-13
 - empdk create_plugin, 13-6
 - empdk validate_plugin, 13-5
- compliance content, example, 12-26
- compliance examples, 12-34
- compliance framework
 - defining, 12-24
 - example, 12-26
 - syntax, 12-24
 - tags, 12-25, 12-27
- compliance real-time monitoring views, 19-27
- compliance standard rules, 12-2
- compliance standards
 - adding, 12-1
 - defining, 12-21
 - example, 12-23
 - process, 12-1
 - syntax, 12-21
 - tags, 12-22
- compliance views, 19-15
- compliance XML, packaging, 12-32
- configuration collection tables, 6-1, 6-4
- configuration data, upgrading, 6-15
- configuration management tables, 6-2
- configuration management views, 19-32

- configuration metadata, 6-8
- configuration metadata file, 6-2
- configuration metadata XML file, 6-4
 - COLUMN element, 6-10
 - elements, 6-9
 - example, 6-12
 - METADATA element, 6-9
 - packaging, 6-13
 - TABLE element, 6-10
- configuring automatic discovery, 11-12
- configuring Flash Builder, 8-97
- convert_mp command options, 16-4
- convert_mp command-line utility, 16-2
- converting a metadata plug-in archive, 16-5
- converting a metadata-only MPCUI
 - implementation, 8-28
- creating connectors, 12-9
- creating event-specific customization XML, 9-3
- creating plug-in archive, 13-6
- creating plugin_registry.xml file, 2-7
- creating plug-ins
 - adding targets, 1-1
 - advanced plug-in, 1-5
 - basic plug-in, 1-4
 - deploying, 1-1
 - designing, 1-1
 - developing, 1-1
 - importing into Enterprise Manager, 1-1
 - intermediate plug-ins, 1-5
 - packaging, 1-1
 - staging, 1-1
 - testing, 1-1
 - validating, 1-1
- creating plugin.xml file, 2-3
- credential information
 - retrieving, 8-61
- credential region, 8-70
- custom configuration specification views, 19-40
- custom data source, 8-41
 - creating, 8-41
 - updating, 8-43
- CustomDataSource.setRow method, 8-43
- CustomDataSource.setTimestampedRows
 - method, 8-43
- customization specification, 9-2
- customizing Incident Manager, 9-1

D

- data service tag, 8-17
- data services, 8-6
- data source
 - line chart, 8-71
- data source, binding, 8-42
- database configuration views, 19-44
- default collection file, 1-5, 3-2
- default collection file, creating, 3-15
- default collection metadata elements, overview, 3-18
- default collection metadata file
 - CollectionItem element, 3-19

- Condition, 3-20
 - MetricColl element, 3-19
 - Schedule element, 3-19
 - TargetCollection element, 3-18
- default collections metadata file, 3-9, 6-2
- default filter, overwrite description, 4-18
- define filter name, 4-14
- define filter prompt, 4-15
- defining
 - advanced metrics, 3-11
 - defining a management user interface, 8-2
 - defining a plug-in, 2-1
 - introduction, 2-1
 - defining compliance framework, 12-24
 - defining management user interface, process, 8-2
 - defining metrics, 3-8
 - defining navigation, 8-18
 - defining pages, 8-13, 8-22
 - defining target type metadata, 3-4
- deleting jobs, 8-58
- demo sample Flex UI, elements, 8-100
- demo sample MPCUI, 8-99
- demo sample project, setting up, 8-98
- deployed plug-in
 - modifying, 8-101
- deploying plug-ins, 13-8, 13-9
- developing plug-ins, 1-1
- development environment options, MPCUI, 8-95
- development guidelines, reports, 4-33
- dialogs
 - defining, 8-77
 - displaying, 8-78
 - registration, 8-77
- dialogs, defining, 8-25
- discovery content, 11-6
- discovery content, packaging, 11-7
- discovery examples, 11-10
- discovery framework, 11-3, 11-5
- discovery inputs, 11-5
- discovery metadata elements, 11-4
- discovery process, 11-1
- discovery script
 - creating, 11-5
 - example, 11-3
 - variables, 11-5
- discovery scripts, 11-7, 11-8
- discovery XML, creating, 11-2
- discovery XSD, 11-2
- DLF file, 12-29
- DMS Fetchlet/Agent Integration Instructions, 21-33
- dynamic instance properties, 3-7
- Dynamic Monitoring Service, 21-30
- Dynamic Monitoring Service (DMS) fetchlet, 21-30
- Dynamic Time Selector, 4-29

E

- EDK, 1-2, 1-5, 8-2, 11-2
 - downloading, 1-2
 - installing, 1-3

- EDK, components, 1-2
- EM CLI utility, 13-11
 - setting up, 13-8
- emagent_perl.trc, 11-4
- emcli add_target command, 13-11
- emcli import_update command, 13-8
- emctl register oms metadata command, 13-12, 13-13
- emd_common.pl file, 11-4
- empdk create_plugin command, 13-6
- empdk tool, 1-2
- empdk validate_plugin command, 13-5
- empty tabel, display, 4-17
- Empty Table Text, 4-18
- empty table, header type, 4-17
- empty table, headers, 4-17
- Enterprise Manager data, accessing, 8-37
- entity types, 12-9
- entity types, filtering, 12-10
- events views, 19-53, 19-60
- event-specific customization metadata elements, 9-4
- event-specific customization XML, 9-3
- event-specific customization XSD, 9-3
- Extensibility Development Kit, see EDK

F

- facet, definition, 12-10
- facets, 12-9
- fetchlet, definition, 3-8
- fetchlets
 - DMS, 21-30
 - HTTP data, 21-19
 - JDBC, 21-35
 - JMX, 21-41
 - OS command, 21-2
 - overview, 21-1
 - REST, 21-54
 - SNMP, 21-15
 - SQL, 21-9
 - URL timing, 21-25
 - URLXML, 21-23
 - WBEM, 21-37
 - web services, 21-44
 - WS-Management, 21-50
- file locations
 - compliance DLF files, 12-32
 - compliance XML, 12-31
 - compliance_rule.xml, 13-4
 - compliance.dlf, 13-5
 - configuration metadata XML file, 6-13
 - default_collections.xml, 13-4
 - derivedAssoc_rule.xml, 13-4
 - discovery JARs, 11-8
 - discovery metadata, 11-6
 - discovery.xml, 13-5
 - job_type.xml, 13-4
 - MPCUI metadata XML file, 8-27
 - MPCUI SWF binary file, 8-27
 - mpcui.xml, 13-5
 - plugin_registry.xml, 13-3

- plugin.xml, 13-3
- report.xml, 13-4
- target_type.xml, 13-4
- target-type_ecmdef.xml, 13-4
- files
 - configuration metadata, 6-2
 - default collections, 6-2
 - DLF, 12-29
 - metric definition, 3-9
 - MPCUI metadata file, 8-8
 - plugin_discovery.xml, 11-4
 - target type metadata, 6-2
- Fill, 4-22
- filter name, default, 4-16
- filter name, null default, 4-16
- filter names, translate, 4-16
- filter tip text, 4-16
- filtering entity types, 12-10
- Flash Builder, configuring, 8-97
- Flex Builder, 8-95
- Flex implementation
 - ActionScript, 8-8
 - defining home page, 8-8
 - MXML, 8-7
 - SWF binary file, 8-8
- flex implementation, 8-7
- Flex implementation, process, 8-3
- Flex SDK, 8-95
- Flex UI metadata file
 - example, 8-8
- framework
 - discovery, 11-3, 11-5

G

- generic discovery integration example, 11-2
- getData method, 8-38
- getIntegrationClass method, 8-21
- getIntegrationClass() method, 8-21
- getTargetInfo() method, 8-49
- grouping similar metrics, 3-16
- guided discovery, defining UI, 8-87
- Guided Resolution region
 - adding customizations, 9-14
- Guided Resolution region, customizing, 9-2

H

- hardware views, 19-63
- Height, 4-23, 4-27
- home page customizations, migrating, 8-107
- horizontal charts, 8-73
- Horizontal or Vertical, 4-23
- HTTP Data Fetchlets, 21-19
- HTTP data fetchlets, 21-19
- hyperlinks, tables, 4-19

I

- icons, defining, 8-85
- importing plug-in into Enterprise Manager, 13-8

- importing the plug-in, prerequisites, 13-8
- Incident Details region, 9-2
 - adding customizations, 9-12
- Incident Manager
 - customizing, 9-1
- incidents and problems region, 8-69
- InfoDisplay class, 8-81
- InfoItem class, 8-81
- information displays
 - defining, 8-81
- information item
 - defining, 8-81
- Information Publisher, 4-1
- init method, 8-24
- init(Train) method, 8-26
- inputParams element, 8-20
- installing the EDK, 1-3
- instance properties, defining, 3-7
- integration class, 8-5, 8-21
- integration metadata, defining, 8-12
- intermediate plug-in, 1-5
- inventory views, 19-65
- invokeActivity directive, 8-19
- invokeActivity method, 8-35
- Is PL/SQL Statement, 4-9, 4-23

J

- Java
 - Management Extensions framework, 17-18
- Java content, required by discovery, 11-8
- JDBC Fetchlet, 21-35
- JDBC fetchlet, 21-35
- JMX, 17-18
- JMX command line tool
 - syntax, 17-32
 - usage, 17-20, 17-33
- JMX fetchlet, 21-41
- job service, 8-55
- job status, checking, 8-58
- job summary region, 8-70
- job type definitions, converting, 16-8
- job views, 19-71

L

- label-value pairs, 8-81
- Legend Position, 4-23, 4-27
- legend, controlling, 8-72
- line chart data source, 8-71
- line charts, 8-70
- Link Destination, 4-28
- links, defining, 8-83
- Linux patching views, 19-80
- list filter names, 4-15
- logging
 - adding to your code, 8-93
 - options for output, 8-94

M

- management repository view examples, 19-159
- management template views, 19-81
- management user interface
 - defining, 8-1
- Maximum Number of Rows, 4-11
- MenuMetadata element, 8-19
- Message Style, 4-28
- Message Text, 4-28
- metadata
 - basic plug-in, 2-2
 - configuration, 6-2
 - default collection elements, 3-18
 - default collection file, 3-2
 - defining target types, 3-4
 - definitions, 3-2
 - discovery, 11-4
 - plugin.xml file, 2-3
 - target type definition file, 1-4
 - target type file, 3-2
 - updating deployed files, 13-12
 - versioning, 3-4
- metadata plug-in archive, converting, 16-5
- metadata plug-in framework, 16-1
- metadata registration service (MRS), 13-12
- metadata-based UI MPCUI metadata file
 - example, 8-8
- metadata-only implementation, 8-7
 - limitations, 8-12
- metadata-only implementation, process, 8-4
- metadata-only MPCUI implementation,
 - converting, 8-28
- Metric Column Name, 4-26
- metric definition files, 3-9
- Metric Details Element, 4-26
- Metric Name, 4-26
- metric services, 8-37
- metric views, 19-86
- metric, definition, 3-8
- metrics, defining, 3-8
- MetricValuesDataService tag, 8-17, 8-38
- MGMT\$AGENTS_MONITORING_
 - TARGETS, 19-147
- MGMT\$ALERT_CURRENT, 19-89
- MGMT\$ALERT_HISTORY, 19-95
- MGMT\$APPL_PATCH_AND_PATCHSET, 19-108
- MGMT\$APPLIED_PATCHES, 19-108
- MGMT\$APPLIED_PATCHSETS, 19-109
- MGMT\$AVAILABILITY_CURRENT, 19-94
- MGMT\$AVAILABILITY_HISTORY, 19-95
- MGMT\$BLACKOUT_HISTORY, 19-9
- MGMT\$BLACKOUTS, 19-10
- MGMT\$CA_EXECUTIONS, 19-71
- MGMT\$CA_TARGETS, 19-71
- MGMT\$CCC_ALL_OBS_BUNDLES, 19-27
- MGMT\$CCC_ALL_OBSERVATIONS, 19-28
- MGMT\$CCC_ALL_VIOLATIONS, 19-29
- MGMT\$CCS_DATA, 19-43
- MGMT\$CCS_DATA_SOURCE_VISIBLE, 19-41
- MGMT\$CCS_DATA_VISIBLE, 19-41, 19-42

MGMT\$COMPLIANCE_STANDARD, 19-17
 MGMT\$COMPLIANCE_STANDARD_
 GROUP, 19-18
 MGMT\$COMPLIANCE_STANDARD_RULE, 19-15
 MGMT\$COMPLIANCE_SUMMARY, 19-31
 MGMT\$COMPLIANCE_TREND, 19-31
 MGMT\$COMPLIANT_TARGETS, 19-30
 MGMT\$COMPOSITE_CS_EVAL_SUMMARY, 19-21
 MGMT\$CS_EVAL_SUMMARY, 19-19
 MGMT\$CS_GROUP_EVAL_SUMMARY, 19-24
 MGMT\$CS_RULE_EVAL_SUMMARY, 19-23
 MGMT\$CS_TARGET_ASSOC, 19-24
 MGMT\$CSA_COLLECTIONS, 19-32
 MGMT\$CSA_FAILED, 19-35
 MGMT\$CSA_HOST_COOKIES, 19-37
 MGMT\$CSA_HOST_CPUS, 19-38
 MGMT\$CSA_HOST_CUSTOM, 19-37
 MGMT\$CSA_HOST_IOCARDS, 19-39
 MGMT\$CSA_HOST_NICS, 19-39
 MGMT\$CSA_HOST_OS_COMPONENTS, 19-36
 MGMT\$CSA_HOST_OS_FILESYSTEMS, 19-40
 MGMT\$CSA_HOST_OS_PROPERTIES, 19-40
 MGMT\$CSA_HOST_RULES, 19-38
 MGMT\$CSA_HOST_SW, 19-37
 MGMT\$CSR_CURRENT_VIOLATION, 19-25
 MGMT\$CSR_VIOLATION_CONTEXT, 19-26
 MGMT\$DB_CONTROLFILES, 19-46
 MGMT\$DB_DATAFILES, 19-45
 MGMT\$DB_DBNINSTANCEINFO, 19-46
 MGMT\$DB_FEATUREUSAGE, 19-47
 MGMT\$DB_INIT_PARAMS, 19-48
 MGMT\$DB_LICENSE, 19-49
 MGMT\$DB_OPTIONS, 19-52
 MGMT\$DB_REDOLOGS, 19-49
 MGMT\$DB_ROLLBACK_SEGS, 19-50
 MGMT\$DB_SGA, 19-51
 MGMT\$DB_TABLESPACES, 19-44
 MGMT\$DB_TABLESPACES_ALL, 19-52
 MGMT\$EM_ECM_MOS_PROPERTIES, 19-148
 MGMT\$EM_ECM_TARGET_FRESHNESS, 19-148
 MGMT\$EM_HOMES_PLATFORM, 19-107
 MGMT\$EM_RULE_VIOL_CTXT_DEF, 19-27
 MGMT\$ESA_ALL_PRIVS_REPORT, 19-135
 MGMT\$ESA_ANY_DICT_REPORT, 19-135
 MGMT\$ESA_ANY_PRIV_REPORT, 19-135
 MGMT\$ESA_AUDIT_SYSTEM_REPORT, 19-136
 MGMT\$ESA_BECOME_USER_REPORT, 19-136
 MGMT\$ESA_CATALOG_REPORT, 19-136
 MGMT\$ESA_CONN_PRIV_REPORT, 19-137
 MGMT\$ESA_CREATE_PRIV_REPORT, 19-137
 MGMT\$ESA_DBA_GROUP_REPORT, 19-137
 MGMT\$ESA_DBA_ROLE_REPORT, 19-138
 MGMT\$ESA_DIRECT_PRIV_REPORT, 19-138
 MGMT\$ESA_EXMPT_ACCESS_REPORT, 19-138
 MGMT\$ESA_KEY_OBJECTS_REPORT, 19-138
 MGMT\$ESA_OH_OWNERSHIP_REPORT, 19-139
 MGMT\$ESA_OH_PERMISSION_REPORT, 19-139
 MGMT\$ESA_POWER_PRIV_REPORT, 19-139
 MGMT\$ESA_PUB_PRIV_REPORT, 19-140
 MGMT\$ESA_SYS_PUB_PKG_REPORT, 19-140
 MGMT\$ESA_TABSP_OWNERS_REPORT, 19-140
 MGMT\$ESA_TRC_AUD_PERM_REPORT, 19-140
 MGMT\$ESA_WITH_ADMIN_REPORT, 19-141
 MGMT\$ESA_WITH_GRANT_REPORT, 19-141
 MGMT\$ESM_COLLECTION_LATEST, 19-141
 MGMT\$ESM_FILE_SYSTEM_LATEST, 19-142
 MGMT\$ESM_PORTS_LATEST, 19-142
 MGMT\$ESM_SERVICE_LATEST, 19-142
 MGMT\$ESM_STACK_LATEST, 19-142
 MGMT\$EVENT_ANNOTATION, 19-57, 19-63
 MGMT\$EVENTS, 19-56, 19-62
 MGMT\$EVENTS_LATEST, 19-55, 19-61
 MGMT\$HOMES_AFFECTED, 19-108
 MGMT\$HOSTPATCH_GROUPS, 19-80
 MGMT\$HOSTPATCH_GRP_COMPL_HIST, 19-81
 MGMT\$HOSTPATCH_HOST_COMPL, 19-81
 MGMT\$HOSTPATCH_HOSTS, 19-80
 MGMT\$HW_CPU_DETAILS, 19-63
 MGMT\$HW_IO_DEVICES, 19-65
 MGMT\$HW_NIC, 19-64
 MGMT\$HW_NIC_BONDS, 19-64
 MGMT\$INCIDENT_ANNOTATION, 19-55, 19-61
 MGMT\$INCIDENT_CATEGORY, 19-54, 19-60
 MGMT\$INCIDENT_TARGET, 19-55, 19-61
 MGMT\$INCIDENTS, 19-53, 19-60
 MGMT\$J2EE_APPLICATION, 19-3
 MGMT\$J2EEAPP_EJBCOMPONENT, 19-4
 MGMT\$J2EEAPP_JRFWS, 19-4
 MGMT\$J2EEAPP_JRFWSOPER, 19-5
 MGMT\$J2EEAPP_JRFWSPOLICY, 19-5
 MGMT\$J2EEAPP_JRFWSUPPORT, 19-6
 MGMT\$J2EEAPP_WEBAPPCOMPONENT, 19-7
 MGMT\$J2EEAPP_WSCONFIG, 19-8
 MGMT\$J2EEAPP_WSPORTCONFIG, 19-8
 MGMT\$JOB_ANNOTATIONS, 19-79
 MGMT\$JOB_EXECUTION_HISTORY, 19-76
 MGMT\$JOB_NOTIFICATION_LOG, 19-79
 MGMT\$JOB_STEP_HISTORY, 19-78
 MGMT\$JOB_TARGETS, 19-75
 MGMT\$JOBS, 19-74
 MGMT\$MANAGEABLE_ENTITIES, 19-148
 MGMT\$METRIC_CATEGORIES, 19-86
 MGMT\$METRIC_COLLECTION, 19-87
 MGMT\$METRIC_CURRENT, 19-98
 MGMT\$METRIC_DAILY, 19-101
 MGMT\$METRIC_DETAILS, 19-97
 MGMT\$METRIC_ERROR_CURRENT, 19-88
 MGMT\$METRIC_ERROR_HISTORY, 19-88
 MGMT\$METRIC_HOURLY, 19-99
 MGMT\$OH_CLONE_PROPERTIES, 19-111
 MGMT\$OH_COMP_DEP_RULE, 19-112
 MGMT\$OH_COMP_INST_TYPE, 19-112
 MGMT\$OH_COMPONENT, 19-111
 MGMT\$OH_CRS_NODES, 19-110
 MGMT\$OH_DEP_HOMES, 19-110
 MGMT\$OH_FILE, 19-116
 MGMT\$OH_HOME_INFO, 19-109
 MGMT\$OH_INSTALLED_TARGETS, 19-118
 MGMT\$OH_INV_SUMMARY, 19-117
 MGMT\$OH_PATCH, 19-114

- MGMT\$OH_PATCH_FIXED_BUG, 19-115
- MGMT\$OH_PATCHED_COMPONENT, 19-115
- MGMT\$OH_PATCHED_FILE, 19-116
- MGMT\$OH_PATCHSET, 19-113
- MGMT\$OH_VERSIONED_PATCH, 19-113
- MGMT\$OS_COMPONENTS, 19-103
- MGMT\$OS_FS_MOUNT, 19-105
- MGMT\$OS_HW_SUMMARY, 19-103
- MGMT\$OS_INIT_SERVICES, 19-107
- MGMT\$OS_KERNEL_PARAMS, 19-105
- MGMT\$OS_LIMITS, 19-107
- MGMT\$OS_MODULES, 19-106
- MGMT\$OS_PATCH_SUMMARY, 19-104
- MGMT\$OS_PATCHES, 19-106
- MGMT\$OS_PROPERTIES, 19-106
- MGMT\$OS_SUMMARY, 19-102
- MGMT\$PA_RECOM_METRIC_SOURCE, 19-117
- MGMT\$PROBLEM_ANNOTATION, 19-59
- MGMT\$PROBLEMS, 19-58
- MGMT\$SERVICETAG_INSTANCES, 19-142
- MGMT\$SERVICETAG_REGISTRY, 19-143
- MGMT\$STORAGE_REPORT_DATA, 19-143
- MGMT\$STORAGE_REPORT_DISK, 19-145
- MGMT\$STORAGE_REPORT_ISSUES, 19-145
- MGMT\$STORAGE_REPORT_KEYS, 19-144
- MGMT\$STORAGE_REPORT_LOCALFS, 19-146
- MGMT\$STORAGE_REPORT_NFS, 19-146
- MGMT\$STORAGE_REPORT_PATHS, 19-144
- MGMT\$STORAGE_REPORT_VOLUME, 19-145
- MGMT\$TARGET, 19-65
- MGMT\$TARGET_ASSOCIATIONS, 19-68
- MGMT\$TARGET_FLAT_MEMBERS, 19-69
- MGMT\$TARGET_MEMBERS, 19-69
- MGMT\$TARGET_METRIC_COLLECTIONS, 19-91
- MGMT\$TARGET_METRIC_SETTINGS, 19-92
- MGMT\$TARGET_PROPERTIES, 19-70
- MGMT\$TARGET_TYPE, 19-66
- MGMT\$TARGET_TYPE_DEF, 19-68
- MGMT\$TARGET_TYPE_PROPERTIES, 19-70
- MGMT\$TEMPLATE_METRIC_SETTINGS, 19-84
- MGMT\$TEMPLATE_METRICCOLLECTION, 19-83
- MGMT\$TEMPLATE_POLICY_SETTINGS, 19-82
- MGMT\$TEMPLATES, 19-81
- MGMT\$VT_EXA_CTRL_VSERVER_TAGS, 19-153
- MGMT\$VT_VM_CONFIG, 19-150
- MGMT\$VT_VM_SW_CFG, 19-151
- MGMT\$VT_VM_VDISKS, 19-152
- MGMT\$VT_VM_VNIC, 19-151
- MGMT\$VT_VSP_CONFIG, 19-153
- MGMT\$WEBLOGIC_APPLICATIONS, 19-118
- MGMT\$WEBLOGIC_CLUSTER, 19-134
- MGMT\$WEBLOGIC_DOMAIN, 19-133
- MGMT\$WEBLOGIC_EJBCOMPONENT, 19-119
- MGMT\$WEBLOGIC_FILESTORE, 19-119
- MGMT\$WEBLOGIC_JDBCdatasource, 19-119
- MGMT\$WEBLOGIC_JDBCMULTIDS, 19-121
- MGMT\$WEBLOGIC_JMSCONNFACTORY, 19-121
- MGMT\$WEBLOGIC_JMSQUEUE, 19-122
- MGMT\$WEBLOGIC_JMSSERVER, 19-122
- MGMT\$WEBLOGIC_JMSTOPIC, 19-123

- MGMT\$WEBLOGIC_JOLTCONNPOOL, 19-123
- MGMT\$WEBLOGIC_JVMSYSPROPS, 19-124
- MGMT\$WEBLOGIC_MACHINE, 19-124
- MGMT\$WEBLOGIC_NETWORK_CHANNELS, 19-125
- MGMT\$WEBLOGIC_NODEMANAGER, 19-125
- MGMT\$WEBLOGIC_OAMCONFIG, 19-134
- MGMT\$WEBLOGIC_OPSSYSROP, 19-133
- MGMT\$WEBLOGIC_RACONFIG, 19-126
- MGMT\$WEBLOGIC_RAOUTBOUNDCONFIG, 19-126
- MGMT\$WEBLOGIC_RESOURCECONFIG, 19-128
- MGMT\$WEBLOGIC_SERVER, 19-128
- MGMT\$WEBLOGIC_STARTSHUTCLASSES, 19-130
- MGMT\$WEBLOGIC_VIRTUALHOST, 19-130
- MGMT\$WEBLOGIC_WEBAPPCOMPONENT, 19-131
- MGMT\$WEBLOGIC_WORKMANAGER, 19-131
- MGMT\$WEBLOGIC_WSCONFIG, 19-132
- migrating home page customizations, 8-107
- monitor target instances, 13-11
- monitoring entity types, 12-9
- monitoring scripts, 11-7
- monitoring views, 19-89
- MpApplication class, 8-21
- MPCUI, 8-1
 - providing online help, 8-110
- MPCUI application
 - application binary file, 8-21
 - defining, 8-20
- MPCUI concepts
 - activity, 8-5
 - integration class, 8-5
 - page, 8-5
 - services, 8-6
 - asynchronous service request handling, 8-6
 - data services, 8-6
 - operation services, 8-6
 - URL, 8-6
- MPCUI development environment options, 8-95
- MPCUI framework, 8-20, 8-37, 8-107
- MPCUI framework services, 8-6
- MPCUI implementation
 - packaging, 8-27
- MPCUI metadata elements, 8-10
- MPCUI metadata file
 - ActivityDefinition element, 8-11
 - creating, 8-8
 - Integration element, 8-11
 - MenuMetadata element, 8-11
 - SqlStatements element, 8-10
 - SwfFiles element, 8-11
 - UIMetadata element, 8-11
- MXML language, 8-5, 8-7

N

- Name Value Pair Display, 4-8
- named credentials sets, 8-61
- Named SQL statements, 16-7

- name-value pairs, adding, 9-2
- navigation, defining, 8-18, 8-34
- Not Yet Managed targets, 11-4
- Null Data String Substitute, 4-11
- Number of Rows to Show, 4-9

O

- online help, defining in MPCUI, 8-110
- OPAR, 13-6, 16-2
- operating system views, 19-102
- operation services, 8-6
- Oracle home directory patching views, 19-107
- Oracle home directory views, 19-109
- Oracle plug-in archive file
 - see* OPAR
- Oracle WebLogic cluster views, 19-134
- Oracle WebLogic domain views, 19-133
- Oracle WebLogic Server views, 19-118
- oracle_home target, 6-12
- OS Command Fetchlets, 21-2
- OS Command fetchlets, 21-2
- OSFetchlet, 21-2
- OSLinesFetchlet, 21-4
- OSLineTokenFetchlet, 21-6
- out-of-box compliance frameworks, 12-9
- out-of-box policy groups, 12-3
- Overwrite Default Button Text, 4-18
- Overwrite Default Filter Tip Text, 4-18

P

- packaged regions
 - availability region, 8-69
 - credentials region, 8-70
 - incidents and problems region, 8-69
 - including, 8-69
 - job summary region, 8-70
- packaged SQL, 8-45
- packaged SQL, writing, 8-47
- packaging compliance XML, 12-32
- packaging discovery content, 11-6, 11-7
- packaging discovery XML, 11-6
- packaging tool, 1-2
- page class, 8-22, 8-23
- page controller, 8-22, 8-23
- page layout components
 - defining, 8-67
 - regions, 8-68
- page model, 8-23
- page.invokeActivity method, 8-35
- pages, defining, 8-13, 8-22
- pie charts, 8-75
- PL/SQL report definition, 16-7
- plug-in
 - basic metadata, 2-2
 - converting existing, 16-1
 - creating archive, 13-6
 - defining, 2-1
 - deploying, 13-8, 13-9

- designing, 1-4
- importing, 13-8
- packaging, 1-2, 16-2, 16-9
- packaging SQL, 8-48
- registering, 16-2
- staging, 13-2
- UI options, 8-6
- validating, 1-2, 13-5
- plug-in archive, 13-6
- plug-in archive, importing, 13-8
- plug-in conversion process, 16-1
- plug-in creation process, 1-1
- plug-in definition files
 - creating, 2-3
 - plugin-registry.xml, 2-3
 - plugin.xml, 2-3
- plug-in definition process, 2-1
- plug-in deployment, 13-10
- plug-in development, getting started, 1-1
- plug-in ID
 - plug-in tag, 2-2
 - product ID, 2-2
 - vendor ID, 2-2
- plug-in identifier, 2-2
 - see* plug-in ID, 2-2
- plug-in metadata format, 16-2
- plug-in stage area, 8-59
- plug-in staging directory, 6-13, 11-6, 12-32
- plug-in upgrade process, 16-2
- plug-in version
 - about, 1-1
 - defining, 2-2
- plugin_discovery.xml
 - AutomaticDiscovery element, 11-4
 - BasicDiscoveryInfo element, 11-5
 - DiscoveryInfo element, 11-4
 - DiscoveryInput element, 11-5
 - DiscoveryModule element, 11-4
 - EmTargetDiscovery element, 11-4
 - SupportedAgentOSList element, 11-4
 - TypesDiscovered element, 11-5
- plugin_discovery.xml file, 11-4
- plugin_registry.xml file
 - creating, 2-7
 - elements, 2-7
 - example, 2-7
 - Plugin, 2-8
 - PlugInLibrary element, 2-8
 - TargetCollections element, 2-8
 - TargetTypes element, 2-8
 - Version attribute, 1-2
- plugin-registry.xml, 16-3
- plugin.xml, 16-3
- plugin.xml file
 - AgentSideCompatibility element, 2-5
 - creating, 2-3
 - elements, 2-4
 - example, 2-3
 - plugin element, 2-4
 - PluginAttributes element, 2-5

- PluginDependencies element, 2-5
- PluginID element, 2-4
- PluginOMSOSAruId element, 2-4
- PluginVersion attribute, 1-2
- PluginVersion element, 2-4
- TargetTypeList element, 2-5
- post-conversion steps for plug-ins, 16-6
- preferred credentials, 8-61
- prerequisites for adding compliance standards, 12-2
- prerequisites, collection configuration data, 6-2
- process
 - compliance standards, 12-1
 - discovery, 11-1
 - plug-in conversion, 16-1
 - plug-in definition, 2-1
 - target configuration data collection, 6-1
 - target metadata files creation, 3-1
 - validation, packaging, and deployment, 13-1
- processing cursor, displaying, 8-83
- processing window, defining, 8-84
- promoting Not Yet Managed targets, 11-10
- pull metrics, 3-8
- push metrics, 3-8

R

- RAW metrics, 6-17
- real-time monitoring facets, 12-10
 - creating, 12-12
 - example, 12-13
 - tags, 12-12
 - time windows, 12-14
- real-time monitoring rules, 12-8
 - creating, 12-16
 - example, 12-19
 - tags, 12-16
- Receivelet, 20-1
- receivelet, definition, 3-9
- registering event-specific customizations, 9-18
- remote operations, 8-59
- RemoteOp service, 8-59
- RemoteOp.performOperation method, 8-60
- Render Image in Column, 4-14
- report definition files, creating, 4-3
- Report Definitions Page, 4-2
- report definitions, converting, 16-7
- report definitions, updating, 4-5
- report testing, interactive, 4-4
- ReportDefinition tag, 4-6
- Report-Wide Parameters, 4-29
- repository check-based rules, 12-2
- repository rule definition
 - example, 12-4
- repository rule syntax, 12-3
- repository rule syntax, description, 12-6
- Response metric, JMX, 17-20, 17-28, 17-33
- REST CLI, 21-59
- REST fetchlet, 21-54
- RESTful web resources, 21-54
- reusable credentials UI components, 8-63

S

- scripts for remote operation, packaging, 8-59
- security
 - Web Services, 17-4
- security views, 19-135
- Separate Rows as Delimiters, 4-13
- Separate Rows for Values in Cell, 4-13
- service request performance, monitoring, 8-50
- service requests
 - batching, 8-52
- service requests, automated polling, 8-51
- service tag views, 19-142
- setRows method, 8-43
- Severity Icon, 4-13
- Show Values in Legend, 4-25
- similar metrics for collection, 3-16
- Slices as Percentage, 4-25
- SNMP Fetchlet, 21-15
- SNMP fetchlet, 21-15
- SNMP Receivelets, 20-1
- SOAP, 17-2
- software library, setting up, 13-8
- Sort Column, 4-8
- Sort Order, 4-8
- Split Table into Multiple Tables by Column, 4-11
- SQL Fetchlet, 21-9
- SQL fetchlet, 21-9
- SQL filter, 4-15
- SQL or PL/SQL queries, reports, 4-3
- SQL or PL/SQL Statement, 4-9, 4-10, 4-24
- SQL query service, 8-45
- SQL statements, packaged, 8-45
- SQLDataService tag, 8-18, 8-45
- SqlQuery interface, 8-46
- Stacked Bar Chart, 4-24
- staging directory structure, 13-2
- staging the plug-in, 13-2
- standard collection metrics, 6-17
- static instance properties, 3-7
- stopping jobs, 8-58
- storage reporting views, 19-143
- submit method, 8-55
- supported customizations, 9-2
- system home pages, defining, 8-28
- systemUiIntegration metadata XML file, 8-31

T

- Table Element Parameters, 4-7
- table header text, overwrite, 4-17
- tables
 - custom data provider, 8-76
 - data service, 8-75
 - defining, 8-75
 - getting selected rows, 8-77
- target configuration data collections
 - defining, 3-18
 - process, 6-1
- target configuration data, collecting, 6-1
- target credentials, defining, 3-5

- target definition files
 - overview, 3-2
- target descriptors
 - TargetMetadata and Display, 3-4
- target discovery, defining, 11-1
- target instance properties, 3-7
- target instance, adding, 13-11
- target metadata, 3-9
- target metadata files
 - creating, 3-1
- target metadata files creation process, 3-1
- target navigator, 8-86
- target services
 - associated targets service, 8-49
 - availability service, 8-50
 - metric metadata service, 8-49
 - target properties service, 8-49
 - working with, 8-48
- Target Type, 4-14, 4-26
- target type facets, 12-10
- target type metadata file, 3-2, 3-9, 6-2
 - creating, 3-3
 - example, 3-3
 - naming, 3-4
- target views, 19-147
- target.getAssociatedTargets() method, 8-49
- target.getAvailability() method, 8-50
- Target.getMetric() method, 8-49
- target.getMetricMetadata () method, 8-49
- targets, adding manually, 11-8
- task automation, 8-54
- test metric, 3-12
- testing discovery, 11-8
- testing Incident Manager, 9-18
- Text Element Parameters, 4-28
- Time Period, 4-8, 4-22, 4-27
- time window facet
 - example, 12-15
 - tags, 12-14
- tracing service, 8-50
- train controller, 8-80
- train events, 8-80
- train pages, defining, 8-26
- train state, 8-80
- trainDone method, 8-26
- trains
 - defining, 8-79
 - definition example, 8-79
 - train controller, 8-80
 - train events, 8-80
 - train state, 8-80
- trains, defining, 8-26
- transient columns, 3-21
- translation support, 12-29
- type properties, defining, 3-5

U

- UI options, 8-6
- updating deployed metadata files, 13-12

- URL Fetchlet (raw), 21-20
- URL Line Token Fetchlet, 21-22
- URL Lines Fetchlet, 21-21
- URL Timing Fetchlet, 21-25
- URL timing fetchlet, 21-25
- UrlEm.homepageUrl method, 8-36
- URLXML Fetchlet, 21-23
- URLXML fetchlet, 21-23

V

- validating the plug-in, 13-5
- verification tool, 1-2
- version, plug-in, 1-1
- versioning metadata, 3-4
- vertical bar charts, 8-74
- views
 - application deployment, 19-3
 - blackout, 19-9
 - compliance, 19-15
 - compliance real-time monitoring, 19-27
 - configuration management, 19-32
 - custom configuration specification, 19-40
 - database configuration, 19-44
 - events, 19-53, 19-60
 - examples, 19-159
 - hardware, 19-63
 - inventory, 19-65
 - job, 19-71
 - Linux patching, 19-80
 - management template, 19-81
 - metric, 19-86
 - monitoring, 19-89
 - operating system, 19-102
 - Oracle home directory, 19-109
 - Oracle home directory patching, 19-107
 - Oracle WebLogic cluster, 19-134
 - Oracle WebLogic domain, 19-133
 - Oracle WebLogic Server, 19-118
 - security, 19-135
 - service tag, 19-142
 - storage reporting, 19-143
 - target, 19-147
- VT Target Views, 19-150

W

- WBEM fetchlet, 21-37
- Web Services, 17-2
 - command-line tool, 17-2
 - monitoring, 17-2
- web services fetchlet, 21-44
- Width, 4-25, 4-27
- WSDL, 17-2
- WS-Management fetchlet, 21-50

Y

- Y-Axis Label, 4-25

