

Oracle® Enterprise Manager

Cloud Administration Guide

12c Release 3 (12.1.0.3)

E28814-08

June 2013

E28814-08

Copyright © 2011, 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 is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.

This software 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 | xix |
| Audience | xix |
| Supported Plug-ins | xix |
| Documentation Accessibility | xix |
| Related Documents | xx |
| Conventions | xx |

Part I Cloud Management Overview

1 Cloud Management - An Overview

| | | |
|-------|---|------|
| 1.1 | Introduction to Cloud Management | 1-1 |
| 1.2 | Managing the Cloud Management Lifecycle With Enterprise Manager | 1-2 |
| 1.2.1 | Planning | 1-2 |
| 1.2.2 | Setting Up the Cloud | 1-3 |
| 1.2.3 | Building the Cloud | 1-4 |
| 1.2.4 | Testing and Deploying a Service | 1-4 |
| 1.2.5 | Monitoring and Managing the Cloud | 1-4 |
| 1.2.6 | Metering, Charging, and Optimization | 1-5 |
| 1.3 | Understanding the Oracle Cloud Service Models | 1-5 |
| 1.3.1 | Oracle Cloud Service Models | 1-5 |
| 1.3.2 | Oracle Cloud Anatomy | 1-6 |
| 1.3.3 | IaaS Components | 1-8 |
| 1.3.4 | DBaaS and MWaaS Components | 1-9 |
| 1.3.5 | TaaS Components | 1-9 |
| 1.4 | Accessing Oracle Cloud Features | 1-9 |
| 1.4.1 | Enterprise Manager Cloud Control Console | 1-10 |
| 1.4.2 | Self Service Portal | 1-10 |

Part II Setting Up and Enabling Cloud

2 How to Enable Cloud

| | | |
|-------|---|-----|
| 2.1 | Enterprise Manager Cloud Deployment Models | 2-1 |
| 2.2 | Enabling Infrastructure as a Service (IaaS) | 2-1 |
| 2.2.1 | Infrastructure as a Service (IaaS): Super Administrator Tasks | 2-2 |
| 2.2.2 | Infrastructure as a Service (IaaS): Cloud Administrator Tasks | 2-2 |

| | | |
|---------|--|-----|
| 2.2.3 | Infrastructure as a Service (IaaS): Self Service Administrator Tasks | 2-3 |
| 2.2.4 | Infrastructure as a Service (IaaS): Self Service User Tasks..... | 2-3 |
| 2.3 | Enabling Platform as a Service (PaaS)..... | 2-3 |
| 2.3.1 | Platform as a Service (PaaS): Super Administrator Tasks | 2-3 |
| 2.3.2 | Platform as a Service (PaaS): Cloud Administrator Tasks..... | 2-3 |
| 2.3.3 | Platform as a Service (PaaS): Self Service Administrator Tasks | 2-4 |
| 2.4 | Enabling Database as a Service (DBaaS)..... | 2-4 |
| 2.4.1 | Provisioning Databases Using DBCA Templates | 2-4 |
| 2.4.2 | Cloning Databases Using RMAN Backup | 2-5 |
| 2.4.3 | Cloning Databases Using Snap Clone | 2-6 |
| 2.4.4 | Schema as a Service | 2-6 |
| 2.4.4.1 | Creating a Schema Using Export Schema Objects Provisioning Profile | 2-6 |
| 2.4.4.2 | Creating an Empty Schema | 2-7 |
| 2.5 | Enabling Middleware as a Service (MWaaS)..... | 2-8 |
| 2.5.1 | Middleware as a Service (MWaaS): Cloud Administrator Tasks | 2-8 |
| 2.5.2 | Middleware as a Service (MWaaS): Self Service Administrator Tasks | 2-8 |
| 2.5.3 | Middleware as a Service (MWaaS): Self Service User Tasks | 2-8 |
| 2.6 | Enabling Testing as a Service (TaaS)..... | 2-9 |
| 2.6.1 | Testing as a Service (TaaS): Cloud Administrator Tasks | 2-9 |
| 2.6.2 | Testing as a Service (TaaS): Self Service Administrator Tasks | 2-9 |
| 2.6.3 | Testing as a Service (TaaS): Self Service User Tasks | 2-9 |

3 Setting Up the Cloud Management Infrastructure

| | | |
|---------|--|------|
| 3.1 | Setting Up the Software Library | 3-1 |
| 3.2 | Setting Up Self Update | 3-3 |
| 3.2.1 | Setting Up Enterprise Manager Self Update Mode | 3-4 |
| 3.2.2 | Assigning Self Update Privileges to Users..... | 3-4 |
| 3.2.3 | Setting Up the EM CLI Utility (Optional) | 3-5 |
| 3.3 | Deploying the Required Plug-ins | 3-5 |
| 3.3.1 | Downloading the Plug-Ins to the Software Library | 3-6 |
| 3.3.1.1 | Downloading Plug-Ins in Online Mode | 3-6 |
| 3.3.1.2 | Downloading Plug-Ins in Offline Mode | 3-7 |
| 3.3.2 | Deploying Plug-Ins to Oracle Management Service..... | 3-9 |
| 3.3.2.1 | Deploying the Plug-ins in Graphical Mode..... | 3-10 |
| 3.3.2.2 | Deploying the Plug-ins in Silent Mode | 3-10 |
| 3.4 | Defining Roles and Assigning Users..... | 3-10 |
| 3.4.1 | Creating a Custom Role for Self Service Application Users..... | 3-12 |
| 3.4.2 | Creating a User and Assigning Roles | 3-12 |
| 3.5 | Configuring LDAP Authentication | 3-13 |
| 3.6 | Configuring Privilege Delegation Settings..... | 3-13 |
| 3.7 | Customizing the Self Service Login Page | 3-14 |
| 3.7.1 | Configuring the Self Service Login Page..... | 3-14 |
| 3.7.2 | Switching Back to the Enterprise Manager Login Page | 3-16 |
| 3.7.3 | Routing SSA Requests to a Specific OMS Pool..... | 3-17 |

Part III Setting Up and Using Infrastructure as a Service

4 Setting Up Cloud for IaaS

| | | |
|--------|---|------|
| 4.1 | Getting Started..... | 4-1 |
| 4.2 | Cloud Infrastructure Privileges | 4-2 |
| 4.3 | Registering the Oracle VM Manager..... | 4-6 |
| 4.3.1 | Updating the Oracle VM Manager URL | 4-8 |
| 4.3.2 | Discovering and Promoting the Oracle VM Manager (Optional) | 4-9 |
| 4.3.3 | Monitoring Configuration (OVM Manager)..... | 4-11 |
| 4.4 | Synchronizing the Oracle VM Manager Targets | 4-12 |
| 4.4.1 | Automatic Synchronization | 4-12 |
| 4.4.2 | Manual Synchronization..... | 4-13 |
| 4.5 | Discovering a Virtual Server | 4-13 |
| 4.6 | Rediscovering a Virtual Server | 4-14 |
| 4.7 | Setting Up Networks | 4-14 |
| 4.7.1 | Generating MAC Addresses | 4-15 |
| 4.7.2 | Configuring a VLAN Group | 4-16 |
| 4.7.3 | Creating Networks | 4-17 |
| 4.8 | Creating a Network Profile | 4-19 |
| 4.9 | Registering Storage Servers | 4-20 |
| 4.9.1 | Registering the File Server..... | 4-21 |
| 4.9.2 | Registering the Storage Array | 4-22 |
| 4.9.3 | File Server Details | 4-25 |
| 4.9.4 | Storage Array Details | 4-25 |
| 4.9.5 | Local File Server Details..... | 4-26 |
| 4.10 | Creating a Storage Repository | 4-26 |
| 4.10.1 | Performing Administrative Operations on a Storage Repository | 4-27 |
| 4.11 | Presenting the Storage Repository | 4-28 |
| 4.12 | Importing Assemblies, Templates, and Virtual Disks into the Storage Repository..... | 4-29 |
| 4.12.1 | Creating a Repository Export (Optional) | 4-31 |
| 4.13 | Using Self Update to Download Oracle VM Templates and Virtual Assemblies | 4-32 |
| 4.14 | Creating a Virtual Server Pool | 4-33 |
| 4.14.1 | Editing a Virtual Server Pool | 4-36 |
| 4.15 | Creating a Zone | 4-37 |
| 4.16 | Setting Up the Storage Quality of Service (QoS) | 4-37 |
| 4.16.1 | Setting Up the Network Type | 4-38 |
| 4.16.2 | Setting Up the Machine Sizes..... | 4-39 |
| 4.16.3 | Configuring the Software Library User..... | 4-39 |
| 4.16.4 | Configuring the Software Library User with External Management Agent..... | 4-39 |
| 4.16.5 | Defining the Request Purging Policy | 4-40 |
| 4.17 | Patching the Oracle VM Server | 4-40 |
| 4.17.1 | Configuring the YUM Repository | 4-40 |
| 4.17.2 | Upgrading the Virtual Server | 4-41 |
| 4.18 | Acknowledging OVM Events | 4-41 |
| 4.19 | Managing the NTP Configuration..... | 4-41 |
| 4.20 | Importing Virtual Machines | 4-43 |
| 4.21 | Acknowledging the Repository Events | 4-45 |

5 Defining and Managing Cloud Policies

| | | |
|---------|---|------|
| 5.1 | Managing Cloud Policies | 5-1 |
| 5.2 | Viewing Cloud Policies | 5-2 |
| 5.3 | Viewing Policies for a Target | 5-3 |
| 5.4 | Defining a Performance Policy | 5-4 |
| 5.5 | Defining a Schedule-Based Policy | 5-6 |
| 5.6 | Activating and Deactivating Policies | 5-7 |
| 5.7 | Viewing Policy Groups | 5-7 |
| 5.8 | Creating a Policy Group | 5-9 |
| 5.8.1 | Example 1 | 5-9 |
| 5.8.1.1 | Policy for Shutting Down Guest Virtual Machine | 5-10 |
| 5.8.1.2 | Policy for Starting Up Guest Virtual Machine | 5-10 |
| 5.8.1.3 | Creating a Policy Group (Start Up and Shut Down) | 5-11 |
| 5.8.2 | Example 2 | 5-11 |
| 5.8.2.1 | Policy to Scale Down a Tier Instance | 5-11 |
| 5.8.2.2 | Policy to Scale Up a Tier Instance | 5-12 |
| 5.8.2.3 | Creating a Policy Group (Scale Up and Scale Down) | 5-13 |
| 5.9 | DRS and DPM Policies | 5-13 |
| 5.9.1 | Creating Automated DRS and DPM Policies | 5-15 |

6 Setting Up the IaaS Self Service Portal

| | | |
|-----|--|-----|
| 6.1 | Setting Up the Self Service Portal | 6-1 |
| 6.2 | Uploading Large Files on the Self Service Portal | 6-7 |

7 Using the IaaS Self Service Portal

| | | |
|---------|---|-----|
| 7.1 | Using the Self Service Portal to Request Server Resources | 7-1 |
| 7.1.1 | Viewing My Requests | 7-2 |
| 7.1.2 | Viewing My Servers | 7-3 |
| 7.1.2.1 | Server Details Page | 7-3 |
| 7.1.2.2 | Scaling Up a Tier Instance | 7-4 |
| 7.1.2.3 | Scaling Down a Tier Instance | 7-4 |
| 7.1.3 | Viewing Storage | 7-4 |
| 7.1.4 | Viewing Chargeback Details | 7-5 |
| 7.1.5 | My Library | 7-5 |
| 7.1.6 | Viewing Policies | 7-6 |
| 7.1.7 | Viewing My Preferences | 7-6 |
| 7.2 | Requesting Servers for a Specific Duration | 7-7 |

8 Administering and Monitoring an IaaS Cloud

| | | |
|-------|--|-----|
| 8.1 | Viewing the Infrastructure Cloud Home Page | 8-1 |
| 8.2 | Viewing the OVM Manager Home Page | 8-3 |
| 8.3 | Viewing and Managing Members | 8-5 |
| 8.4 | Managing Zones | 8-5 |
| 8.4.1 | Viewing the Zone Home Page | 8-5 |
| 8.4.2 | Creating or Editing a Zone | 8-6 |
| 8.4.3 | Deleting a Zone | 8-7 |

| | | |
|-----------|--|------|
| 8.4.4 | Deleting a Zone from Enterprise Manager | 8-7 |
| 8.5 | Viewing the Virtual Server Pool Home Page..... | 8-7 |
| 8.6 | Managing Virtual Servers | 8-8 |
| 8.6.1 | Virtual Server Home Page..... | 8-9 |
| 8.6.2 | Editing a Virtual Server | 8-9 |
| 8.6.3 | Upgrading a Virtual Server | 8-10 |
| 8.6.4 | Starting and Stopping a Virtual Server..... | 8-10 |
| 8.6.5 | Performing Maintenance on the Virtual Server | 8-11 |
| 8.7 | Managing Guest Virtual Machines..... | 8-11 |
| 8.7.1 | Creating a Guest Virtual Machine..... | 8-11 |
| 8.7.2 | Guest Virtual Machine Home Page..... | 8-12 |
| 8.7.3 | Migrating a Guest Virtual Machine | 8-12 |
| 8.7.4 | Starting and Shutting Down a Guest Virtual Machine | 8-13 |
| 8.7.4.1 | Starting a Guest Virtual Machine | 8-13 |
| 8.7.4.2 | Restarting a Guest Virtual Machine..... | 8-13 |
| 8.7.4.3 | Stopping a Guest Virtual Machine..... | 8-13 |
| 8.7.4.4 | Killing a Guest Virtual Machine..... | 8-13 |
| 8.7.5 | Cloning a Guest Virtual Machine..... | 8-14 |
| 8.7.6 | Editing a Guest Virtual Machine | 8-16 |
| 8.7.7 | Saving a Guest Virtual Machine as a Template..... | 8-18 |
| 8.8 | Deploying Guest Virtual Machines | 8-19 |
| 8.8.1 | Getting Started | 8-19 |
| 8.8.2 | Deployment Options..... | 8-20 |
| 8.8.3 | Prerequisites | 8-20 |
| 8.8.3.1 | Creating and Storing Virtualization Components in the Software Library..... | 8-21 |
| 8.8.3.1.1 | Creating an Assembly Component..... | 8-21 |
| 8.8.3.1.2 | Creating a Template Component..... | 8-23 |
| 8.8.3.1.3 | Creating an ISO Component..... | 8-25 |
| 8.8.3.1.4 | Creating a Virtual Disk..... | 8-26 |
| 8.8.4 | Provisioning Guest Virtual Machines Using Oracle Virtual Assemblies (OVA) | 8-28 |
| 8.8.5 | Provisioning a Guest Virtual Machine Using Oracle VM Templates | 8-37 |
| 8.8.6 | Provisioning a Guest Virtual Machine Using an ISO Image | 8-41 |
| 8.8.7 | Creating PXE Bootable Guest Virtual Machines | 8-43 |
| 8.9 | Viewing the Infrastructure Request Dashboard..... | 8-45 |
| 8.9.1 | Viewing Request Details..... | 8-46 |
| 8.9.2 | Viewing the Assembly / Tier Instance Page | 8-46 |

Part IV Setting Up and Using Platform as a Service

9 Setting Up the PaaS Infrastructure

| | | |
|-------|---|-----|
| 9.1 | Getting Started..... | 9-1 |
| 9.2 | Adding Hosts..... | 9-2 |
| 9.3 | Creating a PaaS Infrastructure Zone | 9-2 |
| 9.3.1 | Editing a PaaS Infrastructure Zone | 9-6 |
| 9.3.2 | Deleting a PaaS Infrastructure Zone | 9-7 |
| 9.4 | Creating a PaaS Pool..... | 9-7 |

10 Monitoring the PaaS Infrastructure

| | | |
|------|---|------|
| 10.1 | Viewing the Middleware and Database Cloud Home Page..... | 10-1 |
| 10.2 | Viewing the Middleware and Database Request Dashboard Page..... | 10-3 |
| 10.3 | Viewing All Requests Page..... | 10-4 |
| 10.4 | Viewing the Service Instances Page | 10-5 |
| 10.5 | Viewing the Service Template Details Page..... | 10-6 |
| 10.6 | Viewing the PaaS Infrastructure Zones Page | 10-6 |
| 10.7 | Viewing the PaaS Infrastructure Zone Home Page | 10-7 |

11 Setting Up a DBaaS Cloud

| | | |
|------------|--|-------|
| 11.1 | Getting Started..... | 11-1 |
| 11.2 | Setting Up Credentials for Provisioning | 11-3 |
| 11.3 | Provisioning Database Software..... | 11-3 |
| 11.4 | Deploying the Database | 11-4 |
| 11.5 | Configuring the Oracle Listener | 11-4 |
| 11.6 | Registering and Managing Storage Servers | 11-5 |
| 11.6.1 | Before You Begin..... | 11-5 |
| 11.6.2 | Overview of Registering Storage Servers..... | 11-6 |
| 11.6.3 | Prerequisites for Registering Storage Servers..... | 11-6 |
| 11.6.3.1 | Configuring Storage Servers | 11-6 |
| 11.6.3.1.1 | Obtaining NetApp Hardware Privileges and Licenses | 11-7 |
| 11.6.3.1.2 | Obtaining Sun ZFS Hardware Privileges and Licenses..... | 11-8 |
| 11.6.3.2 | Customizing Storage Agent Proxy..... | 11-9 |
| 11.6.3.2.1 | Acquiring Third Party Licenses | 11-9 |
| 11.6.3.2.2 | Uploading Storage Vendor SDK | 11-10 |
| 11.6.3.2.3 | Overriding the Default SDK | 11-10 |
| 11.6.3.2.4 | Overriding Third Party Server Components..... | 11-11 |
| 11.6.3.3 | Granting Roles and Privileges | 11-12 |
| 11.6.3.3.1 | Accessing Security Class | 11-12 |
| 11.6.3.3.2 | Granting General Privileges | 11-12 |
| 11.6.3.3.3 | Granting Target Privileges..... | 11-13 |
| 11.6.3.3.4 | Granting Roles | 11-13 |
| 11.6.3.3.5 | Granting Privileges for Provisioning..... | 11-13 |
| 11.6.4 | Registering Storage Servers..... | 11-13 |
| 11.6.5 | Administering the Storage Server | 11-16 |
| 11.6.5.1 | Synchronizing Storage Servers..... | 11-16 |
| 11.6.5.2 | Enabling or Disabling Snap Clone | 11-17 |
| 11.6.5.3 | Deregistering Storage Servers..... | 11-18 |
| 11.6.6 | Managing Storage Servers..... | 11-18 |
| 11.6.6.1 | Managing Storage Allocation | 11-19 |
| 11.6.6.2 | Managing Storage Access Privileges | 11-19 |
| 11.6.6.3 | Viewing Storage Registration Overview and Hierarchy..... | 11-20 |

12 Setting Up the DBaaS Self Service Portal

| | | |
|------|---|------|
| 12.1 | Setting Up the Database Cloud Self Service Portal..... | 12-1 |
| 12.2 | DBaaS Using Snap Clone Based Database Provisioning Profile..... | 12-2 |

| | | |
|---------|--|-------|
| 12.3 | DBaaS Using RMAN Backup Based Database Provisioning Profile | 12-3 |
| 12.4 | DBaaS Using DBCA Template Based Database Provisioning Profile | 12-3 |
| 12.5 | Creating a Database Pool | 12-4 |
| 12.6 | Configuring Request Settings | 12-5 |
| 12.7 | Setting Up Quotas | 12-6 |
| 12.8 | Setting Up Profiles and Service Templates | 12-8 |
| 12.9 | Creating a Database Provisioning Profile | 12-8 |
| 12.9.1 | Creating a Database Provisioning Profile Using Snapshots | 12-9 |
| 12.9.2 | Creating a Database Provisioning Profile Using RMAN Backup | 12-11 |
| 12.9.3 | Creating a Database Provisioning Profile Using Existing RMAN Backup | 12-14 |
| 12.9.4 | Creating a Database Provisioning Profile Using DBCA Template | 12-16 |
| 12.9.5 | Creating a Database Provisioning Profile Using Export Schema Objects | 12-18 |
| 12.10 | Creating a Database Service Template | 12-21 |
| 12.10.1 | Creating Service Template Using Snap Clone Profile | 12-21 |
| 12.10.2 | Create Service Template Using RMAN Backup Profile | 12-25 |
| 12.10.3 | Create Service Template Using DBCA Template Profile | 12-27 |
| 12.11 | Configuring Chargeback | 12-29 |
| 12.12 | Using Schema as a Service to Create Schemas | 12-29 |
| 12.12.1 | Creating a Database Pool for Schema as a Service | 12-29 |
| 12.12.2 | Creating a Service Template for Schema as a Service | 12-31 |
| 12.13 | Pre and Post Request Creation / Deletion Scripts | 12-33 |
| 12.13.1 | Sample Scripts | 12-35 |

13 Using the DBaaS Self Service Portal

| | | |
|--------|--|------|
| 13.1 | Using the Database Cloud Self Service Portal | 13-1 |
| 13.2 | Requesting Databases and Schemas | 13-2 |
| 13.2.1 | Requesting a Schema | 13-3 |
| 13.2.2 | Requesting a Database | 13-3 |
| 13.3 | Viewing the Database Service Home Page | 13-4 |
| 13.4 | Viewing the Database Instance Home Page | 13-4 |
| 13.5 | Viewing the Cluster Database Home Page | 13-5 |

14 Monitoring a DBaaS Cloud

| | | |
|------|---|------|
| 14.1 | Viewing the Database Pool Home Page | 14-1 |
| 14.2 | Viewing the Schema Pool Home Page | 14-2 |

15 Setting Up an MWaaS Cloud

| | | |
|----------|---|-------|
| 15.1 | Getting Started with MWaaS for Physical Hosts | 15-1 |
| 15.1.1 | Creating Provisioning Profiles | 15-2 |
| 15.1.1.1 | Creating a WebLogic Domain Provisioning Profile | 15-2 |
| 15.1.1.2 | Using the Out-of-the-box Provisioning Profiles | 15-5 |
| 15.1.2 | Creating a Middleware Home | 15-7 |
| 15.1.2.1 | Provision from Oracle Middleware Home Gold Image Deployment Procedure | 15-7 |
| 15.1.2.2 | Creating an Oracle Middleware Home Gold Image | 15-9 |
| 15.1.3 | Configuring and Saving the Deployment Procedure | 15-11 |

| | | |
|--|--|-------|
| 15.2 | Getting Started with MWaaS for Virtual Hosts..... | 15-11 |
| 16 | Setting Up the MWaaS Self Service Portal | |
| 16.1 | Setting Up the Middleware Cloud Self Service Portal | 16-1 |
| 16.1.1 | Setting Up the Middleware Cloud Self Service Portal for Physical Hosts | 16-1 |
| 16.1.2 | Setting Up the Middleware Cloud Self Service Portal for Virtual Hosts | 16-2 |
| 16.2 | Creating a Middleware Pool | 16-3 |
| 16.3 | Creating a Load Balancer Configuration..... | 16-4 |
| 16.3.1 | Creating a Custom Script..... | 16-7 |
| 16.4 | Configuring Request Settings..... | 16-8 |
| 16.5 | Setting Up Quotas..... | 16-9 |
| 16.6 | Creating a Middleware Service Template | 16-10 |
| 16.7 | Creating a Middleware Physical Service Template | 16-10 |
| 16.8 | Creating a Middleware Virtual Service Template | 16-13 |
| 16.9 | Configuring Chargeback..... | 16-16 |
| 16.9.1 | Sample Scenario | 16-16 |
| 17 | Using the MWaaS Self Service Portal | |
| 17.1 | Using the MWaaS Self Service Portal..... | 17-1 |
| 17.2 | Requesting a Middleware Service | 17-4 |
| 17.3 | Viewing the Middleware Service Home Page | 17-6 |
| 17.3.1 | Deploying and Redeploying an Application..... | 17-7 |
| 17.3.1.1 | Redeploying an Application | 17-8 |
| 17.3.1.2 | Deleting an Application..... | 17-8 |
| 17.3.2 | Scaling Up and Scaling Down a Service..... | 17-8 |
| 17.3.3 | Creating a Data Source..... | 17-8 |
| 17.3.3.1 | Editing a Data Source..... | 17-9 |
| 17.4 | Uploading a Java EE Component to the Software Library | 17-9 |
| 17.5 | Viewing the Application Home Page | 17-10 |
| 18 | Monitoring a MWaaS Cloud | |
| 18.1 | Viewing the Middleware Pool Home Page..... | 18-1 |
| Part V Enabling Testing as a Service (TaaS) | | |
| 19 | About Testing as a Service | |
| 19.1 | TaaS Roles | 19-1 |
| 19.1.1 | Test Administrators..... | 19-2 |
| 19.1.2 | Test Designers | 19-2 |
| 19.1.3 | Testers..... | 19-3 |
| 20 | Setting Up Testing as a Service | |
| 20.1 | Getting Started..... | 20-1 |
| 20.2 | Creating Test Administrators..... | 20-2 |
| 20.3 | Creating an Assembly Component | 20-2 |

| | | |
|------|----------------------------------|------|
| 20.4 | Creating a Deployment Plan | 20-3 |
|------|----------------------------------|------|

21 Using the Testing as a Service Portal

| | | |
|--------|--|-------|
| 21.1 | Prerequisites | 21-2 |
| 21.2 | Test Administrator Tasks..... | 21-3 |
| 21.2.1 | Setting Up Applications..... | 21-3 |
| 21.2.2 | Setting Up Oracle Load Testing Test Drivers | 21-4 |
| 21.2.3 | Setting Up Custom Test Driver Types..... | 21-4 |
| 21.2.4 | Setting Up Quotas..... | 21-5 |
| 21.2.5 | Publishing Test Environments and Test Assets | 21-6 |
| 21.2.6 | Setting Up Chargeback Service..... | 21-6 |
| 21.3 | Test Designer Tasks: Creating Test Assets and Test Environments..... | 21-7 |
| 21.3.1 | Creating Test Assets | 21-7 |
| 21.3.2 | Creating Test Environments | 21-8 |
| 21.4 | Using the Testing Environment..... | 21-9 |
| 21.4.1 | Verifying Available Test Assets and Environments..... | 21-9 |
| 21.4.2 | Creating Tests..... | 21-10 |
| 21.4.3 | Creating and Running Trials..... | 21-11 |
| 21.4.4 | Monitoring On-Going Trials | 21-12 |
| 21.4.5 | Comparing Trial Results..... | 21-12 |
| 21.4.6 | Share Tests | 21-13 |
| 21.4.7 | Viewing Activity and Usage Information..... | 21-13 |
| 21.4.8 | Viewing Deployment Information..... | 21-14 |
| 21.4.9 | Browsing the Testing Home Page | 21-14 |

Part VI Using Consolidation Planner and Chargeback

22 Chargeback Administration

| | | |
|----------|---|-------|
| 22.1 | Overview of Chargeback | 22-1 |
| 22.1.1 | Why Implement Chargeback | 22-1 |
| 22.1.2 | Enterprise Manager Chargeback | 22-2 |
| 22.1.3 | Understanding Targets and Metrics | 22-2 |
| 22.1.4 | About Chargeback, Self Service Applications, and Zones | 22-4 |
| 22.2 | Setting Up Chargeback..... | 22-5 |
| 22.2.1 | Working with Charge Plans..... | 22-7 |
| 22.2.1.1 | Charge Plans and Effective Dates | 22-8 |
| 22.2.1.2 | Using Conditions in Charge Plans..... | 22-8 |
| 22.2.1.3 | Creating and Revising the Universal Charge Plan..... | 22-9 |
| 22.2.1.4 | Creating an Extended Charge Plan..... | 22-9 |
| 22.2.1.5 | Revising Extended Charge Plans | 22-11 |
| 22.2.2 | Creating Cost Centers | 22-12 |
| 22.2.2.1 | Setting the Cost Center Property on the Target Home Page..... | 22-12 |
| 22.2.2.2 | Creating the Cost Center in Chargeback..... | 22-13 |
| 22.2.2.3 | Importing a Business Hierarchy from an LDAP Server | 22-14 |
| 22.2.3 | Adding Targets for Chargeback | 22-15 |

| | | |
|----------|--|-------|
| 22.2.3.1 | Configuring Enterprise Manager Metric Collection for Shared Database Targets and Multitenant Container Databases | 22-16 |
| 22.2.3.2 | Selecting Targets for Chargeback | 22-16 |
| 22.2.3.3 | Identifying Shared Targets | 22-17 |
| 22.2.3.4 | Making Assignments | 22-17 |
| 22.2.3.5 | Reviewing Target Setup | 22-18 |
| 22.2.4 | Assigning Cost Centers to Targets | 22-18 |
| 22.2.5 | Assigning Charge Plans to Targets | 22-18 |
| 22.2.6 | Configuring Chargeback Settings | 22-19 |
| 22.3 | Accessing Chargeback Information | 22-20 |
| 22.3.1 | Following Usage and Charge Trends | 22-21 |
| 22.3.2 | Viewing a Target's Collected Metric Data | 22-22 |
| 22.3.3 | Generating and Distributing Chargeback Reports | 22-23 |
| 22.3.4 | Viewing Chargeback Information in the Self Service Portal | 22-25 |
| 22.3.5 | Sharing Chargeback Data with Other Applications | 22-25 |

23 Using Consolidation Planner

| | | |
|--------|---|-------|
| 23.1 | Overview of Consolidation Planner | 23-1 |
| 23.1.1 | Key Concepts | 23-2 |
| 23.2 | Consolidation Constraints | 23-3 |
| 23.2.1 | Source Server Constraints | 23-3 |
| 23.2.2 | Destination Server Constraints | 23-4 |
| 23.3 | Using Consolidation Planner | 23-4 |
| 23.3.1 | Creating a Consolidation Project | 23-5 |
| 23.3.2 | Using a Pre-configured Consolidation Scenario | 23-6 |
| 23.3.3 | Creating a Custom Consolidation Scenario | 23-7 |
| 23.3.4 | Other Scenario Creation Options | 23-10 |
| 23.3.5 | Evaluating Consolidation Scenarios | 23-11 |
| 23.3.6 | Managing Data Collections | 23-12 |

Part VII Using the Cloud APIs

24 Introduction to Cloud APIs

| | | |
|----------|----------------------------|------|
| 24.1 | Introduction | 24-1 |
| 24.2 | When to Use Cloud APIs | 24-2 |
| 24.3 | Change History | 24-2 |
| 24.4 | Common Behaviors | 24-4 |
| 24.4.1 | Transport Protocol | 24-4 |
| 24.4.2 | URI Space | 24-4 |
| 24.4.3 | Media Types | 24-4 |
| 24.4.4 | Request Headers | 24-5 |
| 24.4.5 | Response Headers | 24-5 |
| 24.4.6 | HTTP Status Codes | 24-6 |
| 24.4.7 | Common Resource Attributes | 24-7 |
| 24.4.7.1 | Resource State | 24-7 |
| 24.4.8 | Collection | 24-8 |
| 24.4.9 | Error Response Message | 24-8 |

25 Cloud Resource Models

| | | |
|--------|---|-------|
| 25.1 | Enterprise Manager Cloud Resource Model 10001..... | 25-1 |
| 25.2 | About Cloud Resources | 25-2 |
| 25.3 | Resource Data Models..... | 25-3 |
| 25.3.1 | Cloud [application/oracle.com.cloud.common.Cloud+json] | 25-3 |
| 25.3.2 | ServiceTemplate [application/oracle.com.cloud.common.ServiceTemplate+json] | 25-3 |
| 25.3.3 | Zone [application/oracle.com.cloud.common.Zone+json] | 25-5 |
| 25.3.4 | Service Family Type [application/oracle.com.cloud.common.ServiceFamilyType+json] 25-6 | |
| 25.3.5 | Service Instance Type [application/oracle.com.cloud.common.InstanceType+json] | 25-6 |
| 25.3.6 | Metric [application/oracle.com.cloud.common.Metric+json] | 25-7 |
| 25.3.7 | Service Instance [application/oracle.com.cloud.common.ServiceInstance+json] .. | 25-8 |
| 25.3.8 | Quota Usage [application/oracle.com.cloud.common.QuotaUsage+json] | 25-9 |
| 25.3.9 | Service Template Finds [application/oracle.com.cloud.common.ServiceTemplateFinds+json] | 25-10 |
| 25.4 | Cloud API Examples | 25-10 |
| 25.4.1 | Cloud Resource | 25-10 |
| 25.4.2 | Service Family Type Resource | 25-14 |
| 25.4.3 | Quota Resource | 25-19 |
| 25.4.4 | Service Instance Type Resource..... | 25-20 |
| 25.4.5 | Zone Resource | 25-21 |
| 25.4.6 | Service Template Resource..... | 25-24 |
| 25.4.7 | Metric Resource..... | 25-27 |

26 Infrastructure as a Service APIs

| | | |
|----------|---|-------|
| 26.1 | Resource Models for Infrastructure as a Service | 26-1 |
| 26.1.1 | Changes for the Infrastructure as a Service | 26-2 |
| 26.1.2 | IaaS Zone [application/oracle.com.cloud.iaas.Zone]..... | 26-2 |
| 26.1.3 | AssemblyInstance [application/oracle.com.cloud.common.AssemblyInstance+json] 26-3 | |
| 26.1.4 | Scalability Group [application/oracle.com.cloud.common.ScalabilityGroup+json] | 26-4 |
| 26.1.5 | VM [application/oracle.com.cloud.common.VM+json]..... | 26-5 |
| 26.1.6 | VNet [application/oracle.com.cloud.common.VNet+json] | 26-7 |
| 26.1.7 | NetworkInterface [application/oracle.com.cloud.common.NetworkInterface+json] | 26-7 |
| 26.1.8 | VMTemplate [application/oracle.com.cloud.common.VMTemplate+json] | 26-8 |
| 26.1.9 | AssemblyTemplate [application/oracle.com.cloud.common.AssemblyTemplate+json]. 26-9 | |
| 26.1.10 | IaaSServiceFamilyType [application/oracle.com.cloud.iaas.IaaSServiceFamilyType+json] | 26-10 |
| 26.1.11 | IaaS Quota Definition [application/oracle.com.cloud.iaas.Quota+json]..... | 26-12 |
| 26.2 | Supported Operations Examples..... | 26-12 |
| 26.2.1 | Creating a Service Instance..... | 26-13 |
| 26.2.1.1 | VM Creation | 26-13 |
| 26.2.1.2 | Assembly Instance Creation | 26-18 |

| | | |
|----------|----------------------------------|-------|
| 26.2.2 | Updating a Service Instance..... | 26-32 |
| 26.2.2.1 | Updating a VM Resource | 26-32 |
| 26.2.3 | Adding a VM Disk..... | 26-37 |
| 26.2.4 | Deleting a Service Instance..... | 26-38 |
| 26.2.5 | Listing Service Instances..... | 26-40 |
| 26.2.6 | Searching Service Templates..... | 26-41 |
| 26.3 | Support for Version 1000 | 26-44 |

27 Database as a Service Family APIs

| | | |
|--------|--|-------|
| 27.1 | Resource Model for Database as a Service | 27-1 |
| 27.1.1 | DB Zone | 27-2 |
| 27.1.2 | DBPlatformTemplate..... | 27-3 |
| 27.1.3 | DBPlatformInstance..... | 27-4 |
| 27.2 | Supported Operations for DBaaS Resources | 27-6 |
| 27.3 | Database as a Service API Examples..... | 27-6 |
| 27.3.1 | Cloud | 27-6 |
| 27.3.2 | Filtering Output for Specific Resource Attributes | 27-9 |
| 27.3.3 | Service Family Type Resource | 27-11 |
| 27.3.4 | Service Instance Type resource..... | 27-13 |
| 27.3.5 | Zone | 27-13 |
| 27.3.6 | DB Platform Template Resource | 27-14 |
| 27.3.7 | Creating Databases..... | 27-16 |
| 27.3.8 | Polling the Database Creation..... | 27-18 |
| 27.3.9 | Deleting a Database Instance | 27-20 |
| 27.4 | Resource Model for Schema as a Service..... | 27-22 |
| 27.4.1 | Resource Model for Schema as a Service | 27-22 |
| 27.4.2 | Schema PlatformTemplate | 27-23 |
| 27.4.3 | SchemaPlatformInstance | 27-24 |
| 27.5 | Supported Operations for Schema as a Service Resources | 27-25 |
| 27.6 | Schema as a Service API Examples | 27-25 |
| 27.6.1 | Schema Platform Template Resource | 27-25 |
| 27.6.2 | Creating Database Services | 27-28 |
| 27.7 | Using EMCLI to Create Database Profiles | 27-33 |
| 27.7.1 | About Database Provisioning Profile T.....ypes | 27-33 |
| 27.7.2 | Creating Provisioning Profiles using EMCLI | 27-34 |
| 27.7.3 | Sample EMCLI Property Files..... | 27-35 |

28 Java as a Service APIs

| | | |
|----------|---|-------|
| 28.1 | Java as a Service Resource Data Models..... | 28-1 |
| 28.1.1 | Cloud | 28-2 |
| 28.1.1.1 | GET Operation on a Cloud URI to return its representation | 28-3 |
| 28.1.2 | JaaS Service Family Type | 28-6 |
| 28.1.2.1 | GET on JaaS Service Family Type URI to Return its Representation | 28-7 |
| 28.1.2.2 | POST on the Jaas Service Family Type URI to Return ApplicationInstanceComponent 28-8 | |
| 28.1.3 | Zone | 28-9 |
| 28.1.3.1 | GET Operation on a Zone URI to return its representation..... | 28-10 |

| | | |
|----------|--|-------|
| 28.1.4 | JavaPlatformTemplate | 28-11 |
| 28.1.4.1 | GET Operation Supported on a JavaPlatformTemplate | 28-12 |
| 28.1.4.2 | POST Operation on a JavaPlatformTemplate..... | 28-13 |
| 28.1.4.3 | GET on the Java Platform Request URL..... | 28-14 |
| 28.1.5 | JavaPlatformInstance | 28-16 |
| 28.1.5.1 | GET Operation Supported on a JavaPlatformInstance | 28-17 |
| 28.1.5.2 | DELETE Operation Supported on a JavaPlatformInstance..... | 28-18 |
| 28.1.5.3 | GET of Java Platform Request URI returned by DELETE operation | 28-19 |
| 28.1.5.4 | POST Operations Supported on a JavaPlatformInstance | 28-19 |
| 28.1.5.5 | GET of Application Instance Deployment Request | 28-20 |
| 28.1.5.6 | POST on the JavaPlatformInstance URI | 28-21 |
| 28.1.5.7 | Get of Data Source Request | 28-22 |
| 28.1.5.8 | PUT Operations Supported on a JavaPlatformInstance | 28-23 |
| 28.1.6 | ApplicationInstanceDeployment | 28-24 |
| 28.1.6.1 | GET Operation supported on an ApplicationInstanceDeployment | 28-25 |
| 28.1.6.2 | DELETE Operation supported on an ApplicationInstanceDeployment..... | 28-26 |
| 28.1.6.3 | Sample GET on Application Instance Deployment Request | 28-27 |
| 28.1.6.4 | Sample PUT Operation supported on an ApplicationInstanceDeployment.. | 28-27 |
| 28.1.7 | DataSource | 28-29 |
| 28.1.7.1 | GET Operation supported on a DataSource | 28-30 |
| 28.1.7.2 | DELETE Operation supported on a DataSource | 28-31 |
| 28.1.7.3 | GET on Data source Request | 28-32 |
| 28.1.7.4 | PUT Operation Supported on a DataSource | 28-32 |
| 28.1.7.5 | GET on Data source Request | 28-34 |
| 28.1.8 | ApplicationInstanceComponent..... | 28-35 |
| 28.1.8.1 | GET Operation supported on an ApplicationInstanceComponent | 28-36 |
| 28.1.8.2 | DELETE Operation supported on an ApplicationInstanceComponent..... | 28-37 |
| 28.2 | Application Component Filtering | 28-37 |
| 28.2.1 | Service Template Finds | 28-37 |
| 28.2.2 | Zone Finds | 28-38 |

29 Chargeback and Metering EMCLI Verbs

| | | |
|----------|--|-------|
| 29.1 | Using Custom Charge Items | 29-1 |
| 29.1.1 | list_charge_item_candidates | 29-1 |
| 29.1.2 | create_charge_item | 29-2 |
| 29.1.3 | delete_charge_item | 29-3 |
| 29.2 | Examples of Custom Charge Item Verb Use | 29-4 |
| 29.2.1 | Metric Extension (PGA) | 29-4 |
| 29.2.2 | Metric Extension (Application A Orders) | 29-6 |
| 29.2.3 | Configuration Extensions | 29-9 |
| 29.3 | Retrieving Metering Data from Enterprise Manager | 29-12 |
| 29.3.1 | get_metering_data Output | 29-14 |
| 29.3.1.1 | Examples | 29-15 |

30 SSA Administration APIs

| | | |
|------|--------------------------|------|
| 30.1 | SSA Admin Resources..... | 30-1 |
|------|--------------------------|------|

| | | |
|----------|---|-------|
| 30.1.1 | GET to get API Entry Point Resources | 30-2 |
| 30.2 | PaasZones..... | 30-3 |
| 30.2.1 | PaasZones API Examples | 30-4 |
| 30.2.1.1 | Retrieving PaaS Zones Metadata | 30-4 |
| 30.2.1.2 | Listing PaaS Zones | 30-4 |
| 30.3 | PaasZone | 30-6 |
| 30.3.1 | PaaS Zone API Examples..... | 30-7 |
| 30.3.1.1 | Creating a PaaS Zone | 30-7 |
| 30.3.1.2 | Updating a PaaS Zone | 30-9 |
| 30.3.1.3 | Deleting a PaaS Zone | 30-10 |
| 30.3.1.4 | Retrieving PaasZone details..... | 30-10 |
| 30.3.1.5 | Getting a PaasZone (Expanded)..... | 30-11 |
| 30.4 | SoftwarePools | 30-13 |
| 30.4.1 | SoftwarePools API Example | 30-13 |
| 30.4.1.1 | Listing Software Pools | 30-14 |
| 30.4.1.2 | Filtering Output based on Search Query Parameters | 30-15 |
| 30.5 | SoftwarePool..... | 30-15 |
| 30.5.1 | SoftwarePool API Examples | 30-17 |
| 30.5.1.1 | Creating a Software Pool..... | 30-17 |
| 30.5.1.2 | Updating a Software Pool | 30-18 |
| 30.5.1.3 | Deleting a Software Pool | 30-19 |
| 30.5.1.4 | Retrieving Software Pool Details | 30-20 |
| 30.5.1.5 | Expands Query Parameter | 30-21 |
| 30.5.1.6 | Retrieving Software Pool Capacity | 30-21 |
| 30.6 | SoftwarePoolMetadata | 30-22 |
| 30.6.1 | Retrieving Software Pool Metadata API Example..... | 30-23 |
| 30.7 | SoftwarePoolFilteredTargets..... | 30-25 |
| 30.7.1 | Retrieving Available Filtered Targets API Example | 30-26 |
| 30.8 | SoftwarePoolCapacity | 30-27 |
| 30.8.1 | Retrieving Software Pool Capacity API Example..... | 30-27 |
| 30.9 | Composite Resource Attributes | 30-29 |
| 30.9.1 | ValueDescriptor | 30-29 |
| 30.9.2 | ValueEntity | 30-29 |
| 30.9.3 | EntityValueHolder..... | 30-29 |
| 30.9.4 | PoolEntityCapacity | 30-30 |
| 30.9.5 | ResourceTypeCapacity | 30-30 |
| 30.10 | EMCLI Verbs for Self Service Applications | 30-30 |
| 30.10.1 | Introduction..... | 30-30 |
| 30.10.2 | Using EMCLI Verbs for SSA Operations | 30-31 |

31 Introduction to Blueprints

| | | |
|------|--|------|
| 31.1 | Introduction to Blueprints | 31-1 |
| 31.2 | Uses of the Blueprint | 31-2 |
| 31.3 | Enterprise Manager Cloud Concepts..... | 31-2 |
| 31.4 | Blueprint Concepts | 31-3 |
| 31.5 | Deploying a Blueprint..... | 31-4 |
| 31.6 | Blueprint Deployment Processing..... | 31-4 |

| | | |
|----------|---|-------|
| 31.7 | Blueprint Examples..... | 31-7 |
| 31.7.1 | Blueprint Structure and Basics..... | 31-7 |
| 31.7.1.1 | Simple Blueprint..... | 31-8 |
| 31.7.1.2 | Simple Resource: Database Service Instance | 31-8 |
| 31.7.1.3 | Intrinsic Functions | 31-9 |
| 31.7.1.4 | Simple Resource with Parameter | 31-9 |
| 31.7.1.5 | Data Section (Named Literals)..... | 31-10 |
| 31.7.2 | Putting It All Together – Multiple Interdependent Resources | 31-11 |
| 31.7.2.1 | Macro Section..... | 31-13 |
| 31.7.2.2 | Attributes of Created Resources (Dependencies) | 31-13 |
| 31.7.3 | Visual Depiction of Blueprint Processing | 31-14 |
| 31.8 | Conclusion | 31-17 |

32 Cloud Blueprints and Blueprint Reference

| | | |
|----------|--|-------|
| 32.1 | Installing the Blueprint Processor | 32-1 |
| 32.1.1 | Linux (Oracle Linux) | 32-2 |
| 32.1.1.1 | Installing Python..... | 32-2 |
| 32.1.1.2 | Installing the Blueprint Processor from Zip File..... | 32-2 |
| 32.1.1.3 | Testing the Installation | 32-3 |
| 32.1.2 | Windows | 32-4 |
| 32.1.2.1 | Installing Python..... | 32-4 |
| 32.1.2.2 | Installing the Blueprint Processor from Zip File..... | 32-5 |
| 32.1.2.3 | Testing the Installation | 32-5 |
| 32.2 | Optional Components for Graphical Summary Report | 32-6 |
| 32.3 | Running the Blueprint Processor..... | 32-6 |
| 32.4 | Blueprint Processing Phases..... | 32-7 |
| 32.4.1 | Initialization..... | 32-8 |
| 32.4.2 | Input Parameter Evaluation..... | 32-8 |
| 32.4.3 | Resource Creation..... | 32-8 |
| 32.4.4 | Outputs..... | 32-9 |
| 32.5 | Language Specifics..... | 32-9 |
| 32.6 | Overview of Blueprint Content | 32-9 |
| 32.6.1 | Inputs Section..... | 32-9 |
| 32.6.1.1 | Example 1..... | 32-10 |
| 32.6.1.2 | Example 2..... | 32-10 |
| 32.6.1.3 | Example 3..... | 32-11 |
| 32.6.2 | Resources Section..... | 32-11 |
| 32.6.3 | Outputs Section..... | 32-12 |
| 32.6.4 | Data Section | 32-13 |
| 32.6.5 | Macros Section | 32-13 |
| 32.6.6 | Expressions | 32-15 |
| 32.6.6.1 | Path Expressions..... | 32-15 |
| 32.6.6.2 | Operator Summary..... | 32-15 |
| 32.6.6.3 | Example: Viewing all values of Info | 32-17 |
| 32.6.6.4 | Examples: Viewing Blueprint Values | 32-18 |
| 32.6.6.5 | Examples: Browsing Your Cloud | 32-19 |
| 32.6.6.6 | Eval or Blueprint Expressions..... | 32-20 |

| | | |
|------------|--|-------|
| 32.6.7 | Intrinsic Functions | 32-21 |
| 32.6.7.1 | Evaluation Intrinsics | 32-22 |
| 32.6.7.1.1 | f_path(pathExpr) | 32-22 |
| 32.6.7.1.2 | Example - Value from Data Section..... | 32-22 |
| 32.6.7.1.3 | Example 2 - Value from Inputs Section..... | 32-22 |
| 32.6.7.1.4 | Example 3 - Default Input Parameter Value via Cloud Lookup | 32-22 |
| 32.6.7.1.5 | f_eval(blueprintExpr) | 32-23 |
| 32.6.7.1.6 | Example (Contrived)..... | 32-23 |
| 32.6.7.2 | Resource Access Intrinsics..... | 32-23 |
| 32.6.7.2.1 | f_getResourceAttr(bpResName, pathExpr) | 32-24 |
| 32.6.7.2.2 | f_getResourceURI(bpResName) | 32-24 |
| 32.6.7.3 | Lookup Intrinsics..... | 32-24 |
| 32.6.7.3.1 | f_getTemplateURI(name, type) | 32-24 |
| 32.6.7.3.2 | f_getZoneURI(name, type) | 32-25 |
| 32.6.7.3.3 | f_getAppCompURI(name, owner, version) | 32-25 |
| 32.6.7.4 | Debugging Intrinsics..... | 32-25 |
| 32.6.7.4.1 | | 32-25 |
| 32.6.7.4.2 | f_print(expression, [printpointMessage]) | 32-26 |
| 32.6.7.5 | Other Intrinsics | 32-26 |
| 32.7 | Dealing with Errors | 32-27 |
| 32.7.1 | YAML Syntax Errors | 32-27 |
| 32.7.2 | Protocol Version Mismatch | 32-27 |
| 32.7.3 | Expression Evaluation Error | 32-28 |
| 32.7.4 | Cloud Resource Creation Error | 32-30 |
| 32.8 | Simulation Mode..... | 32-32 |
| 32.9 | Debugging with the Blueprint Processor | 32-32 |
| 32.9.1 | Printing Intermediate Results | 32-33 |
| 32.9.2 | Pause Points..... | 32-34 |
| 32.9.3 | Breakpoints | 32-35 |
| 32.9.4 | Debugger Commands | 32-35 |
| 32.9.4.1 | Path Command | 32-35 |
| 32.9.4.2 | Continue Command..... | 32-36 |
| 32.9.4.3 | Exit Command | 32-37 |
| 32.9.4.4 | Eval Command | 32-37 |
| 32.10 | Tips and Hints | 32-37 |
| 32.10.1 | Editing YAML - Notepad ++ Example..... | 32-38 |
| 32.10.2 | YAML Duplicate and Name / Value Pairs..... | 32-38 |
| 32.10.3 | Explicit Dependencies..... | 32-38 |
| 32.10.4 | Hint: Use '-T' Option..... | 32-38 |
| 32.10.5 | Help Forums..... | 32-38 |

Index

Preface

This book describes how to set up a Private Cloud, manage and deploy virtualization targets with Oracle Enterprise Manager 12c Release 2. The preface covers the following:

- [Audience](#)
- [Supported Plug-ins](#)
- [Documentation Accessibility](#)
- [Related Documents](#)
- [Conventions](#)

Audience

This document is intended for administrators who want to setup and manage the cloud infrastructure. It is also intended for Cloud Administrators, Self Service Administrators and Self Service Users.

Supported Plug-ins

The features in this book pertain to the following plug-ins:

- Enterprise Manager for Oracle Cloud (12.1.0.6)
- Enterprise Manager for Oracle Consolidation Planning and Chargeback (12.1.0.4)
- Enterprise Manager for Oracle Virtualization (12.1.0.5)
- Enterprise Manager for Oracle Database (12.1.0.4)
- Enterprise Manager for Storage Management (12.1.0.2)

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, see the following documents in the Enterprise Manager documentation set:

- *Oracle Enterprise Manager Cloud Control Introduction*
- *Oracle Enterprise Manager Cloud Control Basic Installation Guide*
- *Oracle Enterprise Manager Cloud Control Advanced Installation and Configuration Guide*
- *Oracle Enterprise Manager Cloud Control Administrator's Guide*

For the latest releases of these and other Oracle documentation, check the Oracle Technology Network at:

<http://www.oracle.com/technetwork/documentation/index.html#em>

Oracle Enterprise Manager also provides extensive Online Help. Click **Help** at the top of any Enterprise Manager page to display the online help window.

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, user names, text that appears on the screen, or text that you enter. |

Part I

Cloud Management Overview

This section provides an introduction to the Oracle Cloud offering and explains the various Cloud service models available.

It contains the following sections:

- [Chapter 1, "Cloud Management - An Overview"](#)

Cloud Management - An Overview

This chapter introduces Cloud Management and provides an overview of the various service offerings and components available with the Oracle Cloud platform. It also describes the Consolidation Planner, which is useful in planning your Cloud requirements, as well as various life cycle management capabilities of Oracle Enterprise Manager including resource management, metering and chargeback support.

This chapter contains the following sections:

- [Introduction to Cloud Management](#)
- [Managing the Cloud Management Lifecycle With Enterprise Manager](#)
- [Understanding the Oracle Cloud Service Models](#)
- [Accessing Oracle Cloud Features](#)

1.1 Introduction to Cloud Management

Enterprises and Cloud service providers can use Oracle Enterprise Manager to build and operate their Cloud services. The functionality provided by Enterprise Manager spans the entire Cloud lifecycle and allows you to setup and manage any type of Cloud service.

Enterprises must support hundreds or even thousands of applications to meet growing business demands. This growth has driven up the cost of acquiring and managing servers and storage. Clouds enable customers to consolidate servers, storage, and database workloads onto a shared hardware and software infrastructure.

By providing on-demand access to servers and storage in a self-service, elastically scalable and metered manner, Enterprise Manager offers the following benefits.

- **Increasing Quality of Service:** IT organizations are not only trying to drive down costs, they are also looking at solutions that will simultaneously improve quality of service in terms of performance, availability and security. Cloud consumers inherently benefit from the high availability characteristics built into the Cloud.

Organizations can also enforce a unified identity and security infrastructure as part of standardized provisioning. Thus, instead of bolting on security policies, these policies and compliance regulations are part of the provisioning process.

- **Enabling Faster Deployment:** Building the Cloud infrastructure using standard building block components (for example, servers, CPUs, storage, and network), configurations, and tools, enables a streamlined, automated, and simplified deployment process.

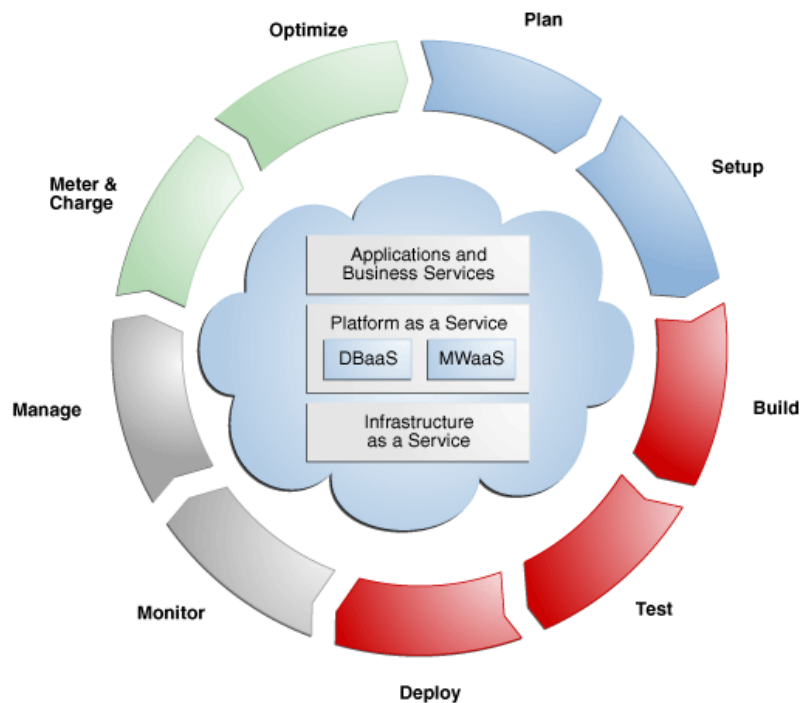
- **Providing Resource Elasticity:** The ability to grow and shrink the capacity of a given database, both in terms of storage size and compute power, allows applications the flexibility to meet the dynamic nature of business workloads.
- **Rapid Provisioning:** Databases in a Cloud can be rapidly provisioned, often by way of a self-service infrastructure, providing agility in application deployment. This reduces overall time in deploying production applications, development platforms, or creating test bed configurations.

1.2 Managing the Cloud Management Lifecycle With Enterprise Manager

Enterprise Manager allows you to manage the entire Cloud lifecycle which includes the following:

- [Planning](#)
- [Setting Up the Cloud](#)
- [Building the Cloud](#)
- [Testing and Deploying a Service](#)
- [Monitoring and Managing the Cloud](#)
- [Metering, Charging, and Optimization](#)

Figure 1–1 Cloud Lifecycle



1.2.1 Planning

Using Enterprise Manager, you can transform existing data centers into a Cloud environment. Before setting up a Cloud, you should map out your infrastructure requirements, such as the physical and virtual networks, storage arrays, applications and so on.

The Enterprise Manager Consolidation Planner is a powerful tool that helps administrators plan the Cloud architecture. It allows you to identify source and destination targets and applicable technical and functional constraints such as where the application can reside, and so on. You can generate consolidation advisories that may include plans to move from Physical to Virtual (P2V), Physical to Physical (P2P), or Physical to an Exadata solution. The Consolidation Planner can also be used to identify the database consolidation plan which is helpful when setting up Database-as-a-Service (DBaaS).

1.2.2 Setting Up the Cloud

Enterprise Manager can be used to model Infrastructure-as-a-Service (IaaS), Database-as-a-Service (DBaaS), and Middleware-as-a-Service (MWaaS) clouds. It is important to select the appropriate Cloud Service Model (as discussed in [Section 1.3, "Understanding the Oracle Cloud Service Models"](#)) that suits the needs of your users and organization. To facilitate Cloud setup, Enterprise Manager offers capabilities for both physical and virtual infrastructure.

For physical infrastructure, Enterprise Manager leverages its core automation framework of deployment procedures, job system, and the enterprise software library. It offers out of the box deployment procedures that can be used for provisioning the pre-requisite software for both databases and middleware. The same automation framework is also used to interact with third party storage systems for the purposes of data cloning and storage management.

For virtual infrastructure, it offers bare metal provisioning of hypervisor and setting up server and storage pools. Once completed, you can group all of these into zones based on functional or QoS characteristics. Enterprise Manager leverages the Virtualization Storage Connect technology, where the Cloud setup process is integrated with storage technologies like Netapp, Hitachi, Fujitsu. Administrators can define standardized service templates for databases and middleware platforms, and publish these as services. These services can represent single-tier templates or complex, multi-tier enterprise platforms.

Enterprise Manager uses components called assemblies created by the Oracle Virtual Assembly Builder (OVAB). Assemblies help package a multi-tier platform into a single metadata which can be deployed by the Enterprise Manager Cloud service. An assembly is essentially a complete multi-tier application stack - including database, application server and other middleware components - packaged as a single downloadable entity. When an assembly is deployed, the result is the creation of a set of related virtual machines representing every tier of the application stack.

Using OVAB, platform architects can model the entire platform topology graphically, define all dependencies, deployment constraints, and deliver the entire stack in the form of an assembly. This assembly can then be published to the centralized Software Library in Enterprise Manager, and be made available to developers as a Cloud service – an entire application development stack, that can be provisioned quickly.

Administrators can create different types of services depending upon the business needs. For example, administrators may offer a database service based on different versions of the Oracle database, but only the ones approved for use within the business.

Enterprise Manager supports role-driven access control. Resource limits, or quotas, are assigned to roles to control access to services. This prevents unauthorized usage of a service while also preventing a few users from using majority of the resources in the Cloud. Integration with LDAP allows Enterprise Manager to inherit enterprise roles.

1.2.3 Building the Cloud

Enterprise Manager allows entire applications or components to be packaged and published to the Cloud as a service. This expedites application development and provisioning processes within an organization.

Developers can publish utility components and applications in the form of assemblies and service templates for reuse within their groups. Similarly, allowing applications to be available as assemblies allows testing teams, business analysts or production teams to deploy pre-built applications in a few clicks.

1.2.4 Testing and Deploying a Service

After an application has been built, it needs to be tested. Enterprise Manager provides a testing portfolio that allows users to test both application changes and changes to the database. The testing solution provides the ability to capture a production load and replay in a test environment, so that the results are predictable. The testing solution also leverages the diagnostic capabilities built into the technology layers and provides prescriptions for remediation.

Enterprise Manager provides a self-service application that lets end-users deploy a service. This self service application can also be customized. End users can choose to provision application assemblies, along with databases and platforms, in an on-demand manner. For each request, users can specify the amount of underlying resources such as CPU, memory, and so on that they require for each component.

Enterprise Manager automatically provisions the requested service and the appropriate resources. The self-service application also lets users define policies to scale out or scale back resources based on schedule or performance metrics. For example, a user could set a policy to elastically scale out a Web server if the processor load on existing Web servers exceeds a certain threshold value.

1.2.5 Monitoring and Managing the Cloud

Enterprise Manager offers a number of inherent monitoring and management features that collectively comprise a full Cloud management system.

For example, Enterprise Manager provides the ability to collate targets into groups for better manageability. The Administration Group feature allows administrators to define monitoring settings, compliance standards and cloud policies through templates and also organize each target in multiple hierarchies, such as Line of Business and Lifecycle status. This allows the monitoring framework to scale to thousands of servers, databases and middleware targets in the Cloud.

Enterprise Manager's built-in Incident Management system allows you to monitor the Cloud for complex operational issues that may affect performance. You can review, suppress, escalate and remediate events that occur as needed, and even integrate incident escalation with existing support ticketing systems.

Contractual Service Level Agreements (SLAs) can be defined to govern the contract between the application owner and the provider of the Cloud. Administrators as well as users can also define management policies that automatically adjust the service resources to ensure that SLAs are met.

The configuration management capabilities of Enterprise Manager are optimized for Cloud environments. For example, Enterprise Manager can monitor vast numbers of configurations continuously to discover changes, measure drifts, pin-point configuration errors, and offer insight into system topologies - all through a single console.

Enterprise Manager Cloud management capabilities are also integrated with My Oracle Support. This integration delivers facilities such as Patch Advisories, Service Request Management, Knowledge Management right on-premise and in-context of the overall Cloud.

The IaaS, DBaaS, and MWaaS Home pages provided through the Enterprise Manager Cloud Control user interface allow Cloud administrators to get a summary view of the requests, the general state of the service such as zones, pools, servers, service instances, and databases.

1.2.6 Metering, Charging, and Optimization

The Metering and Chargeback features in Enterprise Manager enable enterprises to account for actual usage versus representative usage. Administrators can also extend the pricing models to account for fixed costs, configurations, administrative expenses, people costs, energy utilization or a combination of these.

Cloud Management also entails an ongoing optimization of resources as well as processes to make sure that the service levels are persistent. Enterprise Manager provides administrators and application users with features that help rediscover assets, re-evaluate the performance, rebalance the Cloud, and fine-tune the provisioning process. Chargeback supports basic metrics like CPU, memory, and storage usage. It also offers pricing models based on application usage, database usage, and Middleware-level metrics.

1.3 Understanding the Oracle Cloud Service Models

This section describes the available Oracle Cloud service models available.

- [Oracle Cloud Service Models](#)
- [Oracle Cloud Anatomy](#)
- [IaaS Components](#)
- [DBaaS and MWaaS Components](#)
- [TaaS Components](#)

1.3.1 Oracle Cloud Service Models

Oracle's Cloud service models can be divided into two primary categories: Infrastructure as a Service, which allows users to request the physical infrastructure required to run applications; and Platform as a Service, which provides the database and middleware components required by applications.

- **Infrastructure as a Service (IaaS) Service Model:** In the IaaS model, users request servers that are created as guest virtual machines, or Guest VMs.

The IaaS model allows users to request Guest VMs using the Self Service Portal in Enterprise Manager. It also allows users to specify an assembly or a template that is to be deployed on the requested Guest VMs. Using pre-packaged assemblies consisting of the operating system, database software and middleware software, a platform can be deployed using this service.

Users can monitor the services provided using the Self Service Portal and perform limited management operations as permitted. They can also run chargeback reports to review resource usage and chargeback amounts calculated for the resources consumed.

IaaS cloud infrastructure can be built out of Oracle hardware and software components such as Oracle VM, Oracle Solaris, Oracle Enterprise Linux, and so on, or it may have 3rd party components.

- **Platform as a Service (PaaS) Service Model:** The PaaS model allows you to create platforms onto which consumers can deploy their own applications. A platform resource is typically comprised of a host, an operating system, an Oracle WebLogic Application server - all of which can be virtualized. It can also include an Oracle database or RAC cluster.

Models available with PaaS include:

- **Database as a Service (DBaaS) Service Model:** In this model, users can request database services (single instance or RAC) through the Self Service Portal. DBaaS is implemented through several options:
 - **Virtual Machine Based:** The database is deployed as a part of a virtual assembly or template, and several virtual machines share the same physical server. This offers the maximum level of isolation (at the operating system level).
 - **Shared Cluster:** The database is deployed on existing clusterware. Typically the grid infrastructure (Oracle Clusterware, ASM) and database software is pre-installed; the Cloud service essentially consists of the deployment of databases on top of that infrastructure.
 - **Shared Installation:** The database is deployed as a single instance database on an existing database installation.
 - **Shared Database (Schema as a Service):** The database service is a schema deployment on an existing database. It is assumed for purposes of metering and chargeback that each of the consumers of the database will use a different service while accessing the database. This service model is also referred to as Schema as a Service.

As in IaaS, users are allowed to perform a few administrative tasks such as start/stop, backup, and recovery of databases. Chargeback reports are also made available for Self Service users.

- **Middleware as a Service (MWaaS):** In this model, users submit requests for middleware domains to be created. Applications can then be deployed into these domains. MWaaS is implemented through the following options:
 - * **Physical Provisioning Based:** The MWaaS platform is built using physical hosts and Fusion Middleware Provisioning.
 - * **Virtual Machine Based:** The MWaaS platform is deployed as a part of a virtual assembly or template, and several virtual machines share the same physical server. This offers the maximum level of isolation (at the operating system level).
- **Testing as a Service (TaaS):** In this model, testing can be made faster and simpler through the Cloud Testing Self Service Portal. The applications being tested can be provisioned to the private cloud using assemblies or to an existing Enterprise Manager target. The TestDrivers need to be provisioned into the private cloud.

1.3.2 Oracle Cloud Anatomy

In the Oracle implementation, Clouds are composed of logical zones. Zones are composed of logical pools and pools are composed of targets on hosts.

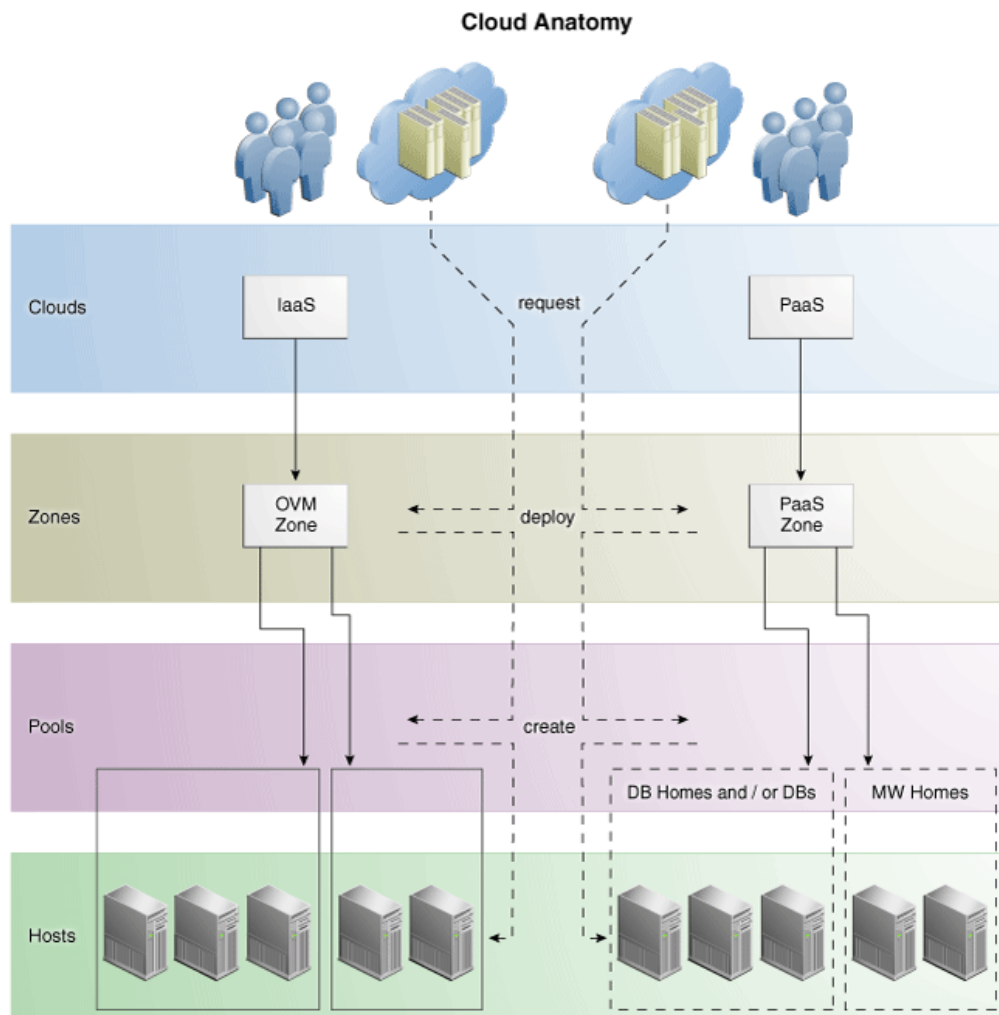
The pools in an IaaS zone are collections of one or more Oracle VM servers and their associated storage resources. The pools in DBaaS zones are collections of one or more Oracle Database homes (used for database requests), or databases (used for schema requests) of the same platform and version (for example, Oracle Database 11.2.0.3 RAC on Oracle Linux 6 x86-64), or Oracle Middleware homes of the same platform and version (for example, Oracle Database 11.2.0.2 on Linux x86-64).

In either IaaS or PaaS, Self Service users will request resources at the zone level from a catalog of templates. Enterprise Manager will then determine which pool in the chosen zone can be used to satisfy the request. The needed Enterprise Manager jobs will be initiated on one or more hosts in the selected pool to create the entities required to fulfill the request.

In an IaaS Cloud, self-service users request that servers be created. These are actually guest virtual machines, or Guest VMs. A single IaaS request may result in one or more virtual machines being created complete with database(s), middleware software and deployed applications.

In the DBaaS view of a PaaS Cloud, a self-service user can request that new databases or schemas in existing databases be created. Databases can be single instance or RAC, depending upon the zones and catalog templates to which the user has access. Similarly, in the MWaaS view of the PaaS Cloud, self-service users request that middleware domains be created.

Figure [Figure 1–2](#) shows the cloud anatomy.

Figure 1–2 Cloud Anatomy

1.3.3 IaaS Components

The IaaS Cloud model consists of the following components:

- **Cloud:** A *Cloud* is a set of storage pools, server pools and zones under the programmatic control of a Cloud Controller and the administrative control of the Cloud Administrator. The Cloud Administrator works with the Cost Center Administrator who has paid for the cloud to determine a resource allocation and charge back policy that meets their needs.
- **Zone:** A Cloud can consist of one or more *zones*. A zone is a logical grouping of resources - for example, servers and storage entities - that facilitate self-service provisioning and administration. A typical zone may consist of hundreds to thousands of servers.

A zone can be an empty zone or consist of a set of server pools. The second case may be simple to set up and will not require shared storage; however no HA and live migration is permitted within this zone.

Zones are non-overlapping, which means that a resource can only belong to one zone. However, resources within a zone may be accessible from another zone. For

example, it is possible for a virtual machine in Zone 1 to interact with a virtual machine in another zone.

- **Server Pool:** A *server pool* is a set of tightly coupled group of servers (typically up to 32 servers) that hosts a set of Guest VMs. The servers are assumed to be largely homogeneous in capabilities and in connectivity. High Availability and Live Migration is permitted within a server pool boundary. Each server pool needs to have access to a shared storage subsystem (which could just be an NFS mount point) to facilitate live migration. In addition, access to a clustered file system may be required to maintain the HA heartbeat file.
- **Storage Entity:** A *storage entity* is an individual file system or block store. Each storage entity is served by a storage pool. Some entities are free standing and will exist until they are deleted. Other storage entities that are associated with one or more Guest VMs are deleted when those VMs are retired.
- **Storage Pool:** A *storage pool* is an abstract storage system that hosts a set of storage entities. A storage pool is generally implemented using a set of storage devices such as disks, SSDs, and storage servers.

1.3.4 DBaaS and MWaaS Components

The DBaaS and MWaaS Cloud structures consist of the following:

- **PaaS Infrastructure Zone:** A *PaaS infrastructure zone* is a group of hosts or Oracle VM zones. Each resource in a zone represents a location at which a service instance is to be deployed.

Before you enable or setup DBaaS or MWaaS, you must create a PaaS Infrastructure Zone which allows you to define the placement policy constraints for a specified set of targets and the users to whom this zone will be available.

- **Software Pools:** A *software pool* is a set of homogeneous resources. A Database Pool, which is created in DBaaS, is a collection of database homes or databases based on the type of cloud service model selected. A Middleware Pool, used for MWaaS, is a collection of middleware homes.

A software pool has the following constraints:

- A target can belong to only one software pool.
- The name of the software pool and the version cannot be modified after it has been created.
- All targets in a software pool must be homogeneous.
- A service template can use multiple zones but only one software pool within each zone.

1.3.5 TaaS Components

The TaaS solution is based on the IaaS platform. Before you set up TaaS, ensure that you have set up Enterprise Manager, and the IaaS components. To use TaaS with the Oracle Load Testing TestDriver, you must download this through the Enterprise Manager self-update.

1.4 Accessing Oracle Cloud Features

Access to the Oracle Cloud features is either through the standard Enterprise Manager console, or the Self Service Portal, which is also part of Enterprise Manager.

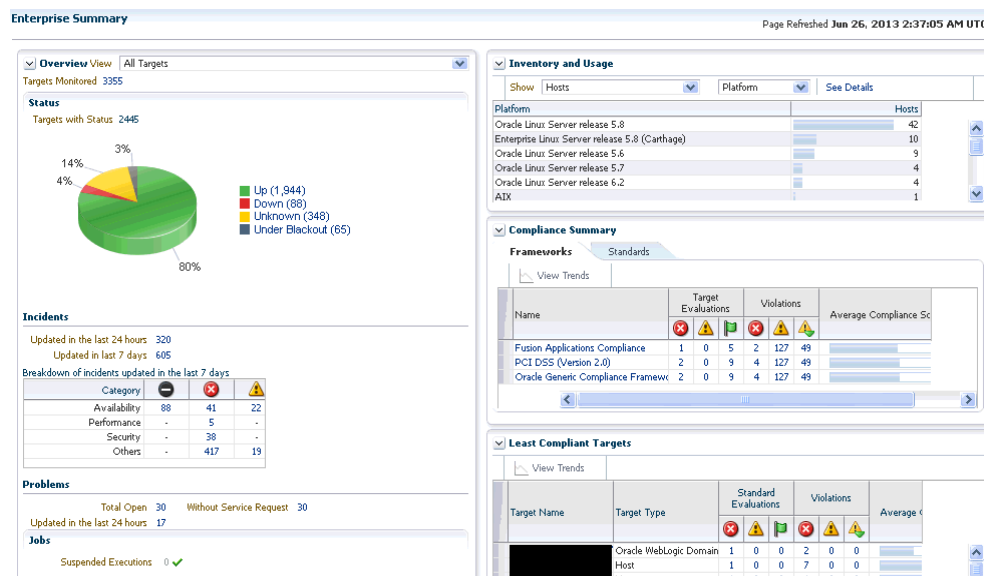
Access to the rest of the Enterprise Manager functionality is restricted. This allows enterprises to safely implement Clouds without worrying about exposing the entire infrastructure to the end users.

1.4.1 Enterprise Manager Cloud Control Console

Administrators will use the Enterprise Manager Cloud Control console to set up, monitor, and manage Cloud services. Each service is managed using a page specific to that service. For example, IaaS, DBaaS, MWaaS, and TaaS all have their own pages that can be accessed directly from the Cloud Summary page or from the Enterprise Manager menu.

The Enterprise Manager Cloud Summary page is a single pane that contains the summary of all Cloud services. Enterprise Manager enables a layer of abstraction that hides the underlying complexities of the application from the end-user. This abstraction is delivered via a self-service interface, both in Graphical User Interface (GUI) and Application Programming User Interface (API).

Figure 1–3 Enterprise Summary



1.4.2 Self Service Portal

To directly manage the Cloud infrastructure, Enterprise Manager provides an out-of-the-box Self Service Portal that allows self-service users to access Cloud services (provisioning applications) without IT intervention. It provides several pre-packaged virtual assemblies and templates for on-demand provisioning, tracks usage of services and resources, and allows data to be used for Chargeback reports and capacity planning.

The Self Service Portal is the Home Page for the Self Service User. Users who have the necessary privileges can navigate between Services pages by clicking the appropriate radio button. How you use the Portal will vary depending on the type of service you are managing.

Figure 1–4 SSA User Portal

Database Cloud Self Service Portal Page Refreshed Jun 4, 2013 4:32:19 PM PDT

Manage ☐ My Servers ☒ My Databases ☐ My Middleware ☐ My Tests My Preferences

Home Chargeback

Notifications

Databases Due to Expire in Next 7 Days: 0

Service Templates Published in Last 7 Days: 0

Usage

You have permission to use these cumulative quota allowances when making database requests.

Databases: 7

0 100

Schema Services: -1

-100 100

Memory: 6.82 GB

0 500

Storage: 38.05 GB

0 500

Databases

View

| Name | InstanceType | InstanceStatus | Start | End | ZoneName |
|---------------|-------------------|----------------|----------------------|--------------|---------------------|
| us.oracle.com | Database Instance | ↑ | Jun 03, 2013 07:52:3 | Indefinitely | yimkeyaz_11203_zon |
| us.oracle.com | Database Instance | ↓ | Jun 03, 2013 08:33:2 | Indefinitely | yimkeyaz_11203_zon |
| us.oracle.com | Database Instance | ↑ | Jun 02, 2013 09:07:2 | Indefinitely | zone slc01exf |
| us.oracle.com | Database Instance | ↑ | Jun 01, 2013 11:24:5 | Indefinitely | zone_11202_sakthi_t |
| us.oracle.com | Database Instance | ↑ | Jun 02, 2013 12:25:4 | Indefinitely | zone_11202_sakthi_t |
| us.oracle.com | Database Instance | ↑ | Jun 01, 2013 04:38:3 | Indefinitely | zone_11202_sakthi_t |
| us.oracle.com | Database Instance | ↑ | Jun 02, 2013 11:04:5 | Indefinitely | zone slc01exf |

Requests

View Name

| Request ID | Name | Status | Service Instance | ReType | Submission Date | Begin |
|------------|--------------------------------------|--------|------------------------|--------|--------------------------|--------|
| 1 | SCHEMAREQ-IMPTRC Execution Error | | | | Cre Jun 1, 2013 4:03:08 | Jun 1, |
| 4 | SCHEMAREQ-IMPTRC Success | | Service_2683886D90615B | | Cre Jun 1, 2013 4:09:57 | Jun 1, |
| 6 | RAICDBREQ_CREATE Success | | | | Cre Jun 1, 2013 4:24:24 | Jun 1, |
| 8 | | | | | Cre Jun 1, 2013 4:37:37 | Jun 1, |
| 2 | SCHEMAREQ-USRDEF Execution Error | | | | Cre Jun 1, 2013 4:05:00 | Jun 1, |
| 21 | after upg with new se Success | | | | Cre Jun 1, 2013 11:24:56 | Jun 1, |
| 22 | SSA 1 USR 1 - Sak 1n Execution Error | | | | Cre Jun 1, 2013 11:27:46 | Jun 1, |

Columns Hidden: 3

Part II

Setting Up and Enabling Cloud

This section provides an introduction to virtualization concepts and describes how to start managing virtualization targets in Enterprise Manager.

It contains the following sections:

- [Chapter 2, "How to Enable Cloud"](#)
- [Chapter 3, "Setting Up the Cloud Management Infrastructure"](#)

How to Enable Cloud

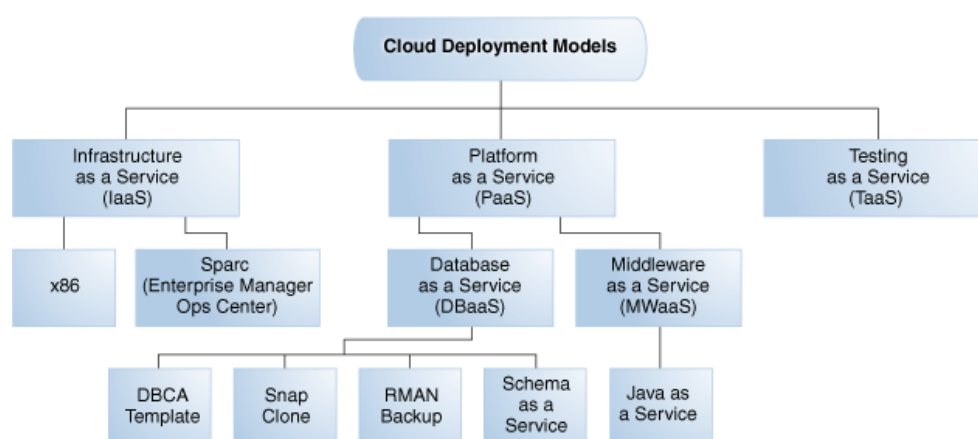
This chapter lists the tasks to be performed to enable your IaaS or PaaS Cloud. It contains the following sections:

- Enterprise Manager Cloud Deployment Models
- Enabling Infrastructure as a Service (IaaS)
- Enabling Platform as a Service (PaaS)
- Enabling Database as a Service (DBaaS)
- Enabling Middleware as a Service (MWaaS)
- Enabling Testing as a Service (TaaS)

2.1 Enterprise Manager Cloud Deployment Models

Figure 2–1 shows the cloud deployment models supported by Enterprise Manager.

Figure 2–1 Cloud Deployment Models



2.2 Enabling Infrastructure as a Service (IaaS)

This section describes the tasks that need to be performed by users with the Super Administrator, EM_CLOUD_ADMINISTRATOR, EM_SSA_ADMINISTRATOR, and EM_SSA_USER roles to setup and use IaaS.

2.2.1 Infrastructure as a Service (IaaS): Super Administrator Tasks

This section lists the tasks that can be performed by the Super Administrator.

Table 2–1 Super Administrator User Tasks (IaaS)

| Cloud Service Model | Usecase |
|------------------------------------|--|
| Infrastructure as a Service (IaaS) | <ul style="list-style-type: none"> Defining Roles and Assigning Users |

2.2.2 Infrastructure as a Service (IaaS): Cloud Administrator Tasks

This section lists the tasks that can be performed for IaaS by users with the EM_CLOUD_ADMINISTRATOR role.

Table 2–2 Cloud Administrator Tasks (IaaS)

| Cloud Service Model | Usecase |
|------------------------------------|--|
| Infrastructure as a Service (IaaS) | <ul style="list-style-type: none"> Registering the Oracle VM Manager Synchronizing the Oracle VM Manager Targets Discovering a Virtual Server Rediscovering a Virtual Server Setting Up Networks Creating a Network Profile Creating a Storage Repository Presenting the Storage Repository Importing Assemblies, Templates, and Virtual Disks into the Storage Repository Creating a Virtual Server Pool Creating a Zone Setting Up the Storage Quality of Service (QoS) Patching the Oracle VM Server Acknowledging OVM Events Managing Cloud Policies Viewing Cloud Policies Viewing Policies for a Target Defining a Performance Policy Defining a Schedule-Based Policy Activating and Deactivating Policies Viewing Policy Groups DRS and DPM Policies |
| Chargeback | <ul style="list-style-type: none"> Overview of Chargeback Setting Up Chargeback Accessing Chargeback Information |
| Consolidation Planner | <ul style="list-style-type: none"> Overview of Consolidation Planner Consolidation Constraints Using Consolidation Planner |

2.2.3 Infrastructure as a Service (IaaS): Self Service Administrator Tasks

This section lists the tasks that need to be performed for IaaS by users with the EM_SSA_ADMINISTRATOR role.

Table 2–3 SSA Administrator Tasks (IaaS)

| Cloud Service Model | Usecase |
|------------------------------------|--|
| Infrastructure as a Service (IaaS) | <ul style="list-style-type: none"> Setting Up the Self Service Portal Viewing the Infrastructure Cloud Home Page Viewing the OVM Manager Home Page Viewing and Managing Members Viewing the Virtual Server Pool Home Page Managing Guest Virtual Machines Deploying Guest Virtual Machines Viewing the Infrastructure Request Dashboard Viewing Request Details |

2.2.4 Infrastructure as a Service (IaaS): Self Service User Tasks

This section lists the tasks that can be performed for IaaS by users with the EM_SSA_USER role.

Table 2–4 SSA User Tasks (IaaS)

| Cloud Service Model | Usecase |
|------------------------------------|---|
| Infrastructure as a Service (IaaS) | <ul style="list-style-type: none"> Using the Self Service Portal to Request Server Resources Requesting Servers for a Specific Duration |

2.3 Enabling Platform as a Service (PaaS)

This section describes the tasks that need to be performed by users with the Super Administrator, EM_CLOUD_ADMINISTRATOR, EM_SSA_ADMINISTRATOR, and EM_SSA_USER roles to setup and use PaaS.

2.3.1 Platform as a Service (PaaS): Super Administrator Tasks

This section lists the tasks that need to be performed for PaaS by users with the Super Administrator role.

Table 2–5 Super Administrator User Tasks (PaaS)

| Cloud Service Model | Usecase |
|--|---|
| Platform as a Service (PaaS) | <ul style="list-style-type: none"> Defining Roles and Assigning Users |
| <ul style="list-style-type: none"> Database as a Service (DBaaS) Middleware as a Service (MWaaS) | <ul style="list-style-type: none"> Configuring Privilege Delegation Settings |

2.3.2 Platform as a Service (PaaS): Cloud Administrator Tasks

This section lists the tasks that need to be performed for PaaS by users with the EM_CLOUD_ADMINISTRATOR role.

Table 2–6 Cloud Administrator Tasks (PaaS)

| Cloud Service Model | Usecase |
|------------------------------|--|
| Platform as a Service (PaaS) | <ul style="list-style-type: none"> ▪ Adding Hosts ▪ Setting Up the Software Library ▪ Creating a PaaS Infrastructure Zone |

2.3.3 Platform as a Service (PaaS): Self Service Administrator Tasks

This section lists the tasks that need to be performed for PaaS by users with the EM_SSA_ADMINISTRATOR role.

Table 2–7 SSA Administrator Tasks (PaaS)

| Cloud Service Model | Usecase |
|------------------------------|---|
| Platform as a Service (PaaS) | <ul style="list-style-type: none"> ▪ Viewing the Middleware and Database Cloud Home Page ▪ Viewing the Middleware and Database Request Dashboard Page ▪ Viewing All Requests Page ▪ Viewing the Service Instances Page ▪ Viewing the Service Template Details Page ▪ Viewing the PaaS Infrastructure Zones Page ▪ Viewing the PaaS Infrastructure Zone Home Page |

2.4 Enabling Database as a Service (DBaaS)

Note: Before you enable Database as a Service, you must have performed the tasks listed in [Section 2.3, "Enabling Platform as a Service \(PaaS\)"](#).

This section describes the tasks that need to be performed to enable the different Database as a Service options:

- [Provisioning Databases Using DBCA Templates](#)
- [Cloning Databases Using RMAN Backup](#)
- [Cloning Databases Using Snap Clone](#)
- [Schema as a Service](#)

2.4.1 Provisioning Databases Using DBCA Templates

This section lists the tasks that need to be performed by the EM_SSA_ADMINISTRATOR and EM_SSA_USER roles to enable DBaaS while using DBCA Template based profile.

Table 2–8 SSA Administrator Tasks: DBCA Template Based Profile

| User Role | Usecase |
|----------------------|--|
| EM_SSA_ADMINISTRATOR | <ul style="list-style-type: none"> ■ Getting Started ■ Setting Up Credentials for Provisioning ■ Provisioning Database Software ■ Configuring the Oracle Listener ■ Registering and Managing Storage Servers ■ Creating a Database Pool ■ Configuring Request Settings ■ Setting Up Quotas ■ Creating a Database Provisioning Profile Using DBCA Template ■ Creating a Database Service Template ■ Configuring Chargeback |

Table 2–9 SSA User Tasks: DBCA Template Based Profile

| User Role | Usecase |
|-------------|--|
| EM_SSA_USER | <ul style="list-style-type: none"> ■ Using the Database Cloud Self Service Portal ■ Requesting a Database ■ Requesting a Schema |

2.4.2 Cloning Databases Using RMAN Backup

This section lists the tasks that need to be performed by the EM_SSA_ADMINISTRATOR and EM_SSA_USER roles to enable DBaaS while using RMAN Backup based profile.

Table 2–10 SSA Administrator Tasks: RMAN Backup Based Profile

| User Role | Usecase |
|----------------------|--|
| EM_SSA_ADMINISTRATOR | <ul style="list-style-type: none"> ■ Getting Started ■ Setting Up Credentials for Provisioning ■ Provisioning Database Software ■ Configuring the Oracle Listener ■ Creating a Database Pool ■ Configuring Request Settings ■ Setting Up Quotas ■ Creating a Database Provisioning Profile Using RMAN Backup ■ Creating a Database Service Template ■ Configuring Chargeback |

Table 2–11 SSA User Tasks: RMAN Backup Based Profile

| User Role | Usecase |
|-------------|---|
| EM_SSA_USER | <ul style="list-style-type: none"> ■ Using the Database Cloud Self Service Portal ■ Requesting a Database |

2.4.3 Cloning Databases Using Snap Clone

This section lists the tasks that need to be performed by the EM_CLOUD_ADMINISTRATOR, EM_SSA_ADMINISTRATOR, and EM_SSA_USER roles to enable DBaaS while using the Snap Clone based profile.

Table 2–12 Cloud Administrator Tasks: Snap Clone Based Profile

| User Role | Usecase |
|------------------------|--|
| EM_CLOUD_ADMINISTRATOR | <ul style="list-style-type: none"> Registering and Managing Storage Servers |

Table 2–13 SSA Administrator Tasks: Snap Clone Based Profile

| User Role | Usecase |
|----------------------|--|
| EM_SSA_ADMINISTRATOR | <ul style="list-style-type: none"> Getting Started Setting Up Credentials for Provisioning Provisioning Database Software Configuring the Oracle Listener Creating a Database Pool Configuring Request Settings Setting Up Quotas Creating a Database Provisioning Profile Using Snapshots Creating a Database Service Template Configuring Chargeback |

Table 2–14 SSA User Tasks: Snap Clone Based Profile

| User Role | Usecase |
|-------------|---|
| EM_SSA_USER | <ul style="list-style-type: none"> Using the Database Cloud Self Service Portal Requesting a Database |

2.4.4 Schema as a Service

This section lists the tasks that need to be performed to enable Schema as a Service. You can choose to:

- Create a schema using an export schema objects profile
- Create an empty schema

2.4.4.1 Creating a Schema Using Export Schema Objects Provisioning Profile

This section lists the tasks that need to be performed by users with EM_SSA_ADMINISTRATOR and EM_SSA_USER roles to create a schema using an export schema objects based provisioning profile.

Table 2–15 SSA Administrator Tasks: Export Schema Objects Based Profile

| User Role | Usecase |
|----------------------|---|
| EM_SSA_ADMINISTRATOR | <ul style="list-style-type: none"> ■ Getting Started ■ Setting Up Credentials for Provisioning ■ Provisioning Database Software ■ Configuring the Oracle Listener ■ Creating a Database Pool for Schema as a Service ■ Configuring Request Settings ■ Setting Up Quotas ■ Creating a Database Provisioning Profile Using Export Schema Objects ■ Creating a Service Template for Schema as a Service ■ Configuring Chargeback |

Table 2–16 SSA User Tasks: Export Schema Objects Based Profile

| User Role | Usecase |
|-------------|---|
| EM_SSA_USER | <ul style="list-style-type: none"> ■ Using the Database Cloud Self Service Portal ■ Requesting a Schema |

2.4.4.2 Creating an Empty Schema

This section lists the tasks that need to be performed by users with EM_SSA_ADMINISTRATOR and EM_SSA_USER roles to create an empty schema.

Table 2–17 SSA Administrator Tasks: Empty Schema

| User Role | Usecase |
|----------------------|---|
| EM_SSA_ADMINISTRATOR | <ul style="list-style-type: none"> ■ Getting Started ■ Setting Up Credentials for Provisioning ■ Provisioning Database Software ■ Configuring the Oracle Listener ■ Creating a Database Pool for Schema as a Service ■ Configuring Request Settings ■ Setting Up Quotas ■ Creating a Service Template for Schema as a Service ■ Configuring Chargeback |

Table 2–18 SSA User Tasks: Empty Schema

| User Role | Usecase |
|-------------|---|
| EM_SSA_USER | <ul style="list-style-type: none"> ■ Using the Database Cloud Self Service Portal ■ Requesting a Schema |

2.5 Enabling Middleware as a Service (MWaaS)

Note: Before you enable Middleware as a Service, you must have performed the tasks listed in [Section 2.3, "Enabling Platform as a Service \(PaaS\)"](#).

This section describes the tasks that need to be performed by users with the EM_CLOUD_ADMINISTRATOR, EM_SSA_ADMINISTRATOR, and EM_SSA_USER roles to setup and use MWaaS.

2.5.1 Middleware as a Service (MWaaS): Cloud Administrator Tasks

This section lists the tasks that need to be performed for MWaaS by users with the EM_CLOUD_ADMINISTRATOR role.

Table 2–19 Cloud Administrator Tasks (MWaaS)

| Cloud Service Model | Usecase |
|---------------------------------|--|
| Middleware as a Service (MWaaS) | <ul style="list-style-type: none"> Getting Started with MWaaS for Physical Hosts Getting Started with MWaaS for Virtual Hosts Uploading a Java EE Component to the Software Library |

2.5.2 Middleware as a Service (MWaaS): Self Service Administrator Tasks

This section lists the tasks that need to be performed for MWaaS by users with the EM_SSA_ADMINISTRATOR role.

Table 2–20 SSA Administrator Tasks (MWaaS)

| Cloud Service Model | Usecase |
|---------------------------------|--|
| Middleware as a Service (MWaaS) | <ul style="list-style-type: none"> Creating a Middleware Pool Configuring Request Settings Setting Up Quotas Creating a Middleware Physical Service Template Creating a Middleware Virtual Service Template Configuring Chargeback |

2.5.3 Middleware as a Service (MWaaS): Self Service User Tasks

This section lists the tasks that need to be performed for MWaaS by users with the EM_SSA_USER role.

Table 2–21 SSA User Tasks (MWaaS)

| Cloud Service Model | Usecase |
|---------------------------------|--|
| Middleware as a Service (MWaaS) | <ul style="list-style-type: none"> ■ Using the MWaaS Self Service Portal ■ Requesting a Middleware Service ■ Viewing the Middleware Service Home Page ■ Viewing the Application Home Page ■ Uploading a Java EE Component to the Software Library |

2.6 Enabling Testing as a Service (TaaS)

This section describes the tasks that need to be performed by users with the EM_CLOUD_ADMINISTRATOR, EM_SSA_ADMINISTRATOR, and EM_SSA_USER roles to setup and use TaaS.

2.6.1 Testing as a Service (TaaS): Cloud Administrator Tasks

This section lists the tasks that need to be performed for TaaS by users with the EM_CLOUD_ADMINISTRATOR role.

Table 2–22 Cloud Administrator Tasks (TaaS)

| Cloud Service Model | Usecase |
|-----------------------------|---|
| Testing as a Service (TaaS) | <ul style="list-style-type: none"> ■ Getting Started ■ Creating Test Administrators ■ Creating an Assembly Component ■ Creating a Deployment Plan |

2.6.2 Testing as a Service (TaaS): Self Service Administrator Tasks

This section lists the tasks that need to be performed for TaaS by users with the EM_SSA_ADMINISTRATOR.

Table 2–23 SSA Administrator Tasks (TaaS)

| Cloud Service Model | Usecase |
|-----------------------------|---|
| Testing as a Service (TaaS) | <ul style="list-style-type: none"> ■ Getting Started ■ Creating Test Administrators ■ Creating an Assembly Component ■ Creating a Deployment Plan |

2.6.3 Testing as a Service (TaaS): Self Service User Tasks

This section lists the tasks that need to be performed for TaaS by users with the EM_SSA_USER role.

Table 2–24 SSA User Tasks (TaaS)

| Cloud Service Model | Usecase |
|-----------------------------|---|
| Testing as a Service (TaaS) | Using the Testing as a Service Portal |

Setting Up the Cloud Management Infrastructure

This chapter describes the initial setup needed before you can begin using the Enterprise Manager Cloud Management solution.

The chapter includes the following sections:

- [Setting Up the Software Library](#)
- [Setting Up Self Update](#)
- [Deploying the Required Plug-ins](#)
- [Defining Roles and Assigning Users](#)
- [Configuring LDAP Authentication](#)
- [Configuring Privilege Delegation Settings](#)
- [Customizing the Self Service Login Page](#)

3.1 Setting Up the Software Library

The Software Library is a repository that stores software patches, virtual appliance images, reference gold images, application software and their associated directive scripts.

In the context of Cloud, the Software Library is the repository for PaaS database provisioning profiles and middleware deployment procedures created by PaaS self service administrators, Java EE Application components created by self service users, and virtual assemblies and templates created by IaaS administrators. The database profiles and middleware deployment procedures are then associated with an appropriate PaaS zone and made available to PaaS self service users, and similarly the virtual assemblies and templates are imported to an OVM zone and made available to IaaS self service users. IaaS self service users can also save their deployment inputs in the Software Library for subsequent use as deployment plans.

To access the Software Library page, from the **Enterprise** menu, click **Provisioning and Patching**, then select **Software Library**. The following screen appears:

Figure 3–1 Software Library Page

Software Library

Page Refreshed Aug 11, 2011 7:07:32 AM PDT

Software Library maintains entities that represent software patches, virtual appliance images, reference gold images, application software and their associated directive scripts. You can pick any of the Oracle-supplied entities, customize them or create a custom one of your own. Once defined, these reusable entities can be referenced from a Deployment Procedure to automate the patching, provisioning or deployment of the associated software.

Actions View View Edit... Delete... Find Name Search

| Name | Type | Subtype | Revision | Status | Maturity | Owner | Description |
|---|------|---------|----------|--------|----------|--------|--|
| Software Library | | | | | | ORACLE | Root Folder for Software Library entities |
| Application Server Provisioning Utilities | | | | | | ORACLE | Entities belonging to AS Provisioning |
| Bare Metal Provisioning | | | | | | ORACLE | Bare Metal Provisioning directory |
| BPELProvisioning | | | | | | ORACLE | BPEL Provisioning Entities |
| Cloud | | | | | | ORACLE | Cloud |
| Coherence Node Provisioning | | | | | | ORACLE | Coherence Node Provisioning Entities |
| Common Provisioning Utilities | | | | | | ORACLE | Directives belonging to Common Provisioning (SIDB and RACPROV) |
| Components | | | | | | SYSMAN | Components Folder |
| Directives | | | | | | SYSMAN | Directives Folder |
| Images | | | | | | SYSMAN | Images Folder |
| Networks | | | | | | SYSMAN | Networks Folder |
| Suites | | | | | | SYSMAN | Suites Folder |
| CompositeDeploy | | | | | | ORACLE | CompositeDeploy Entities |
| CVU Prerequisite-fixup components | | | | | | ORACLE | CVU Prerequisite-fixup components belonging to DB Provisioning |
| DB Provisioning | | | | | | ORACLE | Directives and Components belonging to DB Provisioning |
| Fusion Middleware Provisioning Utilities | | | | | | ORACLE | Directives belonging to FMW Provisioning |
| Java EE Provisioning | | | | | | ORACLE | Java EE Application Provisioning Entities |
| MultiOMS | | | | | | ORACLE | List of Oracle shipped Directives |
| Oracle VM Server Provisioning | | | | | | ORACLE | Oracle VM Server Provisioning directory |
| OSBProvisioning | | | | | | ORACLE | OSBProvisioning Entities |
| Patching | | | | | | ORACLE | Patching directory |
| Prerequisite-fixup components | | | | | | ORACLE | Prerequisite-fixup components Components belonging to DB Prov |
| SoaProvisioning | | | | | | ORACLE | SOA Provisioning Entities |

For the Software Library to be usable, at least one upload location must be configured. Upload File locations are locations configured for storing files uploaded by the Software Library as part of creating or updating an entity. To configure an upload file location, follow these steps:

1. Log in to Enterprise Manager as a user with EM_CLOUD_ADMINISTRATOR role.
2. From the **Setup** menu, select **Provisioning and Patching**, then select **Software Library**.
3. From the **Actions** menu, click **Administration**. The Software Library: Administration page appears where you can select the storage location.
4. Select OMS Shared File System in the Storage Type list and click **Add....**

A storage location can be of two types:

- **Upload File Locations:** Upload File Locations are locations configured for storing files uploaded by Software Library as part of creating or updating an entity. Upload File Locations support two storage options, OMS Shared File System, and OMS Agent File System. For more details on configuring the software library storage, see the *Configuring Software Library* section in the Enterprise Manager Administration Guide.
- **Referenced File Locations:** Referenced File Locations are locations that allow you, the end user to leverage your organization's existing IT infrastructure (like file servers, web servers, or storage systems). These location configurations are used by Software Library when there is a need to stage the files to host targets as part of a provisioning or patching activity. Referenced file locations can either be HTTP or Agent. For more details on configuring the software library storage, see the *Configuring Software Library* section in the Enterprise Manager Administration Guide.

5. Specify a Name and Location that is accessible to all OMSes and click **OK**.

Note: Because the storage location for the Software Library must be accessible to all OMSes as local directories, in a multi-OMS scenario, you must set up a clustered file system using OCFS2, NFS, ACFS, or DBFS. For single OMS systems, any local directory is sufficient. Ensure that sufficient storage space (more than 100 GB for production deployment of Enterprise Manager) has been allocated for the Software Library as this storage space is used to store all the cloud components.

After the Software Library storage has been configured, you can store the following:

- Assemblies, templates, and deployment plans. See [Section 8.8.3.1, "Creating and Storing Virtualization Components in the Software Library"](#).
- Database Provisioning Profiles. See [Section 12.9, "Creating a Database Provisioning Profile"](#).
- Out-of-the-box middleware provisioning profiles and middleware home gold images. See [Section 15.1.1, "Creating Provisioning Profiles"](#).

Only items that are available in the Software Library can be published for deployment by self-service users.

Note: To enable Administrators (or users) to access, and leverage an OMS Agent Filesystem Software Library Location, the owner of the Named Credential must ensure that an explicit View privilege is granted to all the Administrators accessing the OMS Agent location. To do so, you can either click **Add Grant** and add the names of the administrators while creating the Named Credential, or edit an existing Named Credential to grant privileges to other Administrators (or users) by following these steps:

- From the **Setup** menu, select **Security**, then select **Named Credentials**.
 - On the Named Credentials page, click **Manage Access**.
 - On the Manage Access page, click **Add Grant** to add a user, or **Change Privilege** to edit the privileges of an existing user. Click **Save**.
-

For more details on setting up and configuring the Software Library, see the Enterprise Manager Cloud Control Administrator's Guide.

3.2 Setting Up Self Update

The Self Update feature allows you to expand Enterprise Manager's capabilities by updating Enterprise Manager components whenever new and updated features become available between official releases. Oracle makes functional updates available between releases by publishing them to the Enterprise Manager Store, an external site that is periodically checked by Enterprise Manager to obtain information about available updates.

The updatable entities for the Oracle Cloud platform include:

- Cloud Plug-ins. See [Section 3.3, "Deploying the Required Plug-ins"](#)
- Oracle VM Templates and Assemblies. See [Section 4.13, "Using Self Update to Download Oracle VM Templates and Virtual Assemblies"](#).
- Database Provisioning Profiles. See [Section 12.9, "Creating a Database Provisioning Profile"](#).
- Out-of-the-box Middleware Provisioning Profiles. See [Section 15.1.1, "Creating Provisioning Profiles"](#).

Before you can use the Self Update feature, you must satisfy these prerequisites:

- If you are applying an update in online mode, ensure that the My Oracle Support credentials have been set up using the SYSMAN user. This is required to enable entities to be downloaded from the My Oracle Support site.
- The Software Library (also known as the local store) has been configured. Updates are downloaded to this local store before being deployed into Enterprise Manager.

Review the following sections for instructions on setting up Self Update:

- [Setting Up Enterprise Manager Self Update Mode](#)
- [Assigning Self Update Privileges to Users](#)
- [Setting Up the EM CLI Utility \(Optional\)](#)

3.2.1 Setting Up Enterprise Manager Self Update Mode

In order to set up or modify the Enterprise Manager Self Update feature, you must have Enterprise Manager Super Administrator privileges.

1. Log in to Enterprise Manager as an administrator with Super Administrator privileges.
2. From the **Setup** menu, select **Extensibility**, then select **Self Update**. The Self Update console appears with the default setup displayed.
3. From the **General** status area, click the **Connection Mode** status to set either offline or online mode. Enterprise Manager takes you to the Patching Setup page to specify online and offline settings.
4. Once the desired connection mode has been selected, return to the Self Update console.

From here you can select entity types and schedule updates from the Enterprise Manager Update Store.

3.2.2 Assigning Self Update Privileges to Users

Enterprise Manager administrators must have the requisite privileges to use the Self Update feature. The Enterprise Manager Super Administrator must assign the following Self Update roles to these administrators:

- VIEW_SELF_UPDATE: The user can view the Self Update console and can monitor the status of download and apply jobs.
- MANAGE_SELF_UPDATE: The user can schedule download and apply jobs. User can also suppress/unsuppress updates. This privilege implicitly contains VIEW_SELF_UPDATE.

- **EM_INFRASTRUCTURE_ADMIN:** The user can perform all self update operations. This privilege implicitly contains **MANAGE_SELF_UPDATE**.

By default, the Super Administrator will be granted **EM_INFRASTRUCTURE_ADMIN** privilege.

To assign Self Update privileges to regular Enterprise Manager administrators:

1. From the **Setup** menu, select **Security**, then select **Administrators**.
2. Select an administrator and click **Edit**.
3. From the Roles page, assign the appropriate Self Update roles.

3.2.3 Setting Up the EM CLI Utility (Optional)

If you plan to apply software updates in offline mode, you will need to use the Enterprise Manager Command Line Utility, or EM CLI, to import entity archives for deployment to Enterprise Manager.

A page is provided in the Enterprise Manager Cloud Control console with instructions on setting up

EMCLI. Access the page by appending `/console/emcli/download` to the URL used to access the Cloud Control console:

`https://emcc_host:emcc_port/em`

For example:

`https://emcc_host:emcc_port/em/console/emcli/download`

3.3 Deploying the Required Plug-ins

Much of the functionality available in Enterprise Manager Cloud Control is made available through *plug-ins*. As its name implies, a plug-in is a component or module that can be plugged into an existing Enterprise Manager installation to extend its management and monitoring capabilities.

The features that collectively comprise the Oracle Cloud Management solution are provided via several plug-ins which must be deployed to your Oracle Management Service (OMS). The plug-ins that must be deployed to enable each Cloud model are listed below.

You can deploy the plug-ins needed to enable Cloud in two ways:

- If you have not yet installed Enterprise Manager Cloud Control, or have not yet upgraded to the latest Enterprise Manager release, you can deploy the plug-ins as part of the installation or upgrade process. You will select the Advanced Install mode and in the Select Plug-ins screen, select the plug-ins that you wish to install.
- If you already have Enterprise Manager Cloud Control 12c installed, you must download the needed plug-ins to the Software Library. See [Section 3.3.1, "Downloading the Plug-Ins to the Software Library"](#) for instructions.

You will then deploy the plug-ins to your Oracle Management Service (OMS). See [Section 3.3.2, "Deploying Plug-Ins to Oracle Management Service"](#) for instructions.

Plug-ins Required to Enable Infrastructure as a Service (IaaS)

- Enterprise Manager for Oracle Cloud
- Enterprise Manager for Oracle Virtualization
- Enterprise Manager for Oracle Consolidation Planning and Chargeback

Plug-ins Required to Enable Database as a Service (DBaaS)

- Enterprise Manager for Oracle Cloud
- Enterprise Manager for Oracle Virtualization
- Enterprise Manager for Oracle Consolidation Planning and Chargeback
- Enterprise Manager for Oracle Database
- Enterprise Manager for Storage Management

Plug-ins Required to Enable Middleware as a Service (MWaaS)

- Enterprise Manager for Oracle Cloud
- Enterprise Manager for Oracle Virtualization
- Enterprise Manager for Oracle Consolidation Planning and Chargeback
- Enterprise Manager for Oracle Fusion Middleware

Plug-ins Required to Enable Testing as a Service (TaaS)

- Enterprise Manager for Oracle Cloud
- Enterprise Manager for Oracle Virtualization
- Enterprise Manager for Oracle Consolidation Planning and Chargeback

3.3.1 Downloading the Plug-Ins to the Software Library

You can download the plug-ins in online or offline mode. Online refers to an environment where you have Internet connectivity to connect to Enterprise Manager Store. Offline refers to an environment where you do not have Internet connectivity.

This section contains the following sections:

- [Downloading Plug-Ins in Online Mode](#)
- [Downloading Plug-Ins in Offline Mode](#)

3.3.1.1 Downloading Plug-Ins in Online Mode

To download the plug-ins in online mode, follow these steps:

1. From the **Setup** menu, select **Extensibility**, then select **Self Update**.
2. On the Self Update page, in the table, click on **Plug-in**.
3. On the Plug-in Updates page, select the plug-in available for download, and click **Download**.

Multiple selection of plug-ins is not supported.

4. In the Schedule Download dialog, select an appropriate option to schedule the download. You can also select **Immediately** which schedules the job for immediate action. Select **Notify Once downloaded** if you want to be informed once the download is complete.
5. Click **Select**.

Enterprise Manager Cloud Control submits a job to download the selected plug-in from the Enterprise Manager Store to the Software Library.

A confirmation dialog appears to confirm that the job has been submitted successfully. In this confirmation dialog, you can click **Job Details** to track the status of the job.

3.3.1.2 Downloading Plug-Ins in Offline Mode

To download the plug-ins in offline mode, follow these steps:

1. From the **Setup** menu, select **Provisioning and Patching**, then select **Offline Patching**.
2. In the Online and Offline Settings tab, select **Offline**.
3. Click **Apply**.
4. From the **Setup** menu, select **Extensibility**, then select **Self Update**.
5. On the Self Update page, click **Check for Updates**.
A message appears with a URL to an Oracle site from where the updates catalog file can be downloaded.
6. From an Internet-enabled computer, download the catalog file using the aforementioned URL.
7. Copy the downloaded catalog file to the OMS host or the Management Agent host where you plan to deploy the plug-ins.
8. Import the catalog file to Enterprise Manager. For instructions, refer to Importing Catalog Archives.
9. On the Self Update page, in the table, click **Plug-in**.
10. On the Plug-in Updates page, select the imported update that is available for download. Click **Download**.
A message appears with a URL to an Oracle site from where the update can be downloaded.
11. From a computer that is connected to the internet, download the update using the aforementioned URL.
12. Copy the downloaded file to the OMS host or the Management Agent host where you plan to deploy the plug-ins.
13. Import the downloaded plug-in archive to Enterprise Manager. For instructions, refer to Importing Plug-in Archives.

Importing Catalog Archives

To import a catalog archive, follow these steps:

1. Download the catalog archive.
2. Depending on where the catalog file has been download (to a local host or a remote host), run either of the following `emcli` commands to import the downloaded catalog archive.
 - Use this command if the catalog file has been downloaded locally to the machine on which the Oracle Management Service has been installed.

```
emcli import_update_catalog
-file="file"
-omslocal
```
 - Use this command if the catalog file has been downloaded to a remote host on which the Management Agent is running.

```
emcli import_update_catalog
-file="file"
```

```
-host="hostname"
```

In this case, you must specify the necessary credentials to access the host as follows:

```
[-credential_set_name="setname"] | -credential_name="name"  
-credential_owner="owner"
```

For more details on these commands, see the *Enterprise Manager Command Line Reference Guide*.

Importing Plug-In Archives

Import plug-in archives to Oracle Software Library in the following cases:

- When you want to deploy any non-Oracle plug-ins, that is, plug-ins that have been created by a company other than Oracle.
- When you want to import other types of entity archives when Self Update is used in offline mode.

To import a plug-in archive, follow these steps:

1. Download the external archive as described in the previous section.
2. Set up the Enterprise Manager Command Line (EM CLI) utility. To do so, from the **Setup** menu, click **Command Line Interface**. Follow the instructions outlined on the Enterprise Manager Command Line Interface Download page.
3. Import the external archive in one of the following ways, depending on where EM CLI is installed.
 - If Enterprise Manager server is on the system on which you downloaded the plug-in archive (*.opar file), run the following command:

```
emcli import_update  
-file="<path to *.opar file>"  
-omslocal
```

The `-omslocal` flag indicates that the plug-in archive path mentioned in the `-file` option is directly accessible to the EM server.

- If Enterprise Manager server is on a different system than the plug-in archive, run the following command:

```
emcli import_update  
-file="<path to *.opar file you created>"  
-host="host1.example.com"  
-credential_name="host1_creds"  
-credential_owner="admin1"
```

The command syntax is as follows:

`-file`: The absolute path to the *.opar file on the system where you created the archive.

`-host`: The target name for a 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.

Note: 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.

3.3.2 Deploying Plug-Ins to Oracle Management Service

You can deploy plug-ins to an OMS instance in graphical or silent mode. While the graphical mode enables you to deploy one plug-in at a time, the silent mode enables you to deploy multiple plug-ins at a time, thus saving plug-in deployment time and downtime, if applicable.

This section contains the following sections:

- [Deploying the Plug-ins in Graphical Mode](#)
- [Deploying the Plug-ins in Silent Mode](#)

Note:

- To view a visual demonstration on how you can deploy a plug-in to the OMS and discover targets using it, access the following URL:

https://apex.oracle.com/pls/apex/f?p=44785:24:491956260237501::NO::P24_CONTENT_ID,P24_PREV_PAGE:6000,1
 - In a multi-OMS environment, Plug-in Manager automates plug-in deployment on all the management servers.
 - A plug-in upgrade failure could put the Management Repository in an inconsistent state. Therefore it is strongly suggested that you back up the Management Repository, the Oracle Management Service, and the Software Library before upgrading the plug-in. See the *Enterprise Manager Administrator's Guide* for more details.
 - The deployment time varies from one plug-in to another, depending on the volume of data populated in the Management Repository. A page is displayed that allows you to monitor the deployment status.
 - The deployment of some plug-ins requires the OMS to be stopped, and then restarted. This process occurs automatically as part of the plug-in deployment process.
 - While deploying plug-ins to the OMS, OMS plug-in components, discovery plug-in components, and monitoring plug-in components are deployed to the OMS.
-

3.3.2.1 Deploying the Plug-ins in Graphical Mode

To deploy plug-ins to the OMS in graphical mode, follow these steps:

1. From the **Setup** menu, select **Extensibility**, then select **Plug-ins**.
2. On the Plug-ins page, select the plug-in you want to deploy.
3. From the **Deploy On** menu, select **Management Servers**.
4. In the Deploy Plug-in on Management Servers dialog, enter the Management Repository SYS password, and click **Continue**. Proceed through the steps in the dialog box.
5. Click **Deploy**.

3.3.2.2 Deploying the Plug-ins in Silent Mode

To deploy plug-ins to the OMS in silent mode, follow these steps:

1. Log in to EM CLI as follows:

```
$ORACLE_HOME/bin/emcli login -username=sysman
```
2. Run the following command:

```
$ORACLE_HOME/bin/emcli sync
```
3. To deploy the plug-ins on the OMS, run the following command:

```
emcli deploy_plugin_on_server  
-plugin="plug-in_id[:version]  
[-sys_password=sys_password]  
[-prereq_check]"
```

For example,

```
emcli deploy_plugin_on_server  
-plugin="oracle.sysman.db:12.1.0.3.0;oracle.sysman.emas:12.1.0.4.0"
```

Note: The procedure for plug-in deployment remains the same even in a multi-OMS environment. Enterprise Manager automatically detects whether it is a single-OMS or a multi-OMS environment and in case of a multi-OMS environment, Enterprise Manager automatically deploys the selected plug-in on all OMS instances.

If the plug-in deployment on any Oracle Management Service fails, perform the same steps again.

3.4 Defining Roles and Assigning Users

Roles are named groups of related system and object privileges. You can create roles and then assign them to users and to other roles. You can assign any of the existing roles to a new role and the associated privileges. Enterprise Manager contains three out-of-the-box roles for the Cloud Self Service Portal, namely:

- **EM_CLOUD_ADMINISTRATOR:** Users with this role can set up and manage the cloud infrastructure. This role is responsible for deploying the cloud infrastructure (servers, zones, storage, and networks) and infrastructure cloud operations for performance and configuration management.

- **EM_SSA_ADMINISTRATOR:** Users with this role can define quotas and constraints for the self service users and grant them access privileges. Users with this role also have provisioning and patching designer privileges that allow them to create and save deployment procedures, create and view patch plans, and support the plug-in lifecycle on the Management Agent. These privileges are required for initial setup and on going maintenance of the infrastructure.
- **EM_SSA_USER:** Users with this role, by default, can only access the Self Service Portal. An administrator with the **EM_SSA_ADMINISTRATOR** role can provide additional privileges that allow users with the **EM_SSA_USER** role to access other features in Enterprise Manager.

The table below lists the roles associated with each user.

| User Profile | EM_CLOUD_ADMINISTRATOR | EM_SSA_ADMINISTRATOR | EM_SSA_USER |
|---|---|---|---|
| Minimum roles required to create a user | <ul style="list-style-type: none"> ■ EM_CLOUD_ADMINISTRATOR ■ PUBLIC ■ EM_USER | <ul style="list-style-type: none"> ■ EM_SSA_ADMINISTRATOR ■ PUBLIC ■ EM_USER | EM_SSA_USER |
| Roles to be removed when creating a user | NONE | NONE | <ul style="list-style-type: none"> ■ PUBLIC ■ EM_USER |
| Additional roles may be added as required | | | |

The Oracle Cloud Management Self Service Portal is intended for end-users to be able to provision and manage their own cloud services. Since the functions performed by users with the **EM_CLOUD_ADMINISTRATOR** and **EM_SSA_ADMINISTRATOR** roles are consistent across Enterprise Manager, these out-of-box roles can be used as they are. All you need to create users with the **EM_CLOUD_ADMINISTRATOR** and **EM_SSA_ADMINISTRATOR** roles.

But the **EM_SSA_USER** role is used for quota assignment, and to limit access to PaaS Infrastructure zones, and service templates. In this case, the pre-defined role cannot be used as it is defined. You must create custom SSA User roles based on the standard **EM_SSA_ROLE** role as described in [Creating a Custom Role for Self Service Application Users](#). After creating a custom role, you must assign users to this role.

For example, in a DBaaS Cloud setup, you may want to create the following users:

- **CLOUD_ADMIN:** This user will have the **EM_CLOUD_ADMINISTRATOR** role and is responsible for network, system, storage, and administration activities.
- **SSA_ADMIN:** This user will have the **EM_SSA_ADMINISTRATOR** role and is responsible for database administration activities.
- **SSA_USER:** In this case, the default **EM_SSA_USER** role must be customized and a custom role must be created. A user in this role is typically a junior database administrator, developer, or tester.

For more details on the Users and Roles, see the Enterprise Manager Cloud Control Administrator's Guide.

3.4.1 Creating a Custom Role for Self Service Application Users

Typically, you need to create new SSA User roles either for different functional groups like developers, testers, production DBAs, or for different customer teams like the Siebel DBA team, BRM DBA team, and operations team for hosting custom Java applications, and so on. To create a custom SSA user role, follow these steps:

1. Log in to Enterprise Manager as a Super Administrator user.
2. From the **Setup** menu, select **Security**, then select **Roles**.
3. Click **Create** in the Roles page to launch the Create Role wizard.
4. Provide a name and description (SSA_DEV_ROLES) for the role and click **Next**.
5. From the list of Available Roles, select the EM_SSA_USER role and move it to the Selected Roles table. Click **Next**.
6. Accept the default target privileges and click **Next**.
7. Accept the default resource privileges and click **Next**.
8. Skip the Create Role: Administrators step and click **Next**.
9. Review the changes and click **Finish** to create the custom SSA user (SSA_DEV_USERS) role.

3.4.2 Creating a User and Assigning Roles

To create a user called SSA_USER1 and grant the custom role created earlier (SSA_DEV_USERS), follow these steps:

1. Log in to Enterprise Manager as a Super Administrator user.
2. From the **Setup** menu, select **Security**, then select **Administrators**.
3. Click **Create** in the Administrators page to launch the Create Administrator wizard.
4. Enter the name and password for the user (SSA_USER1) and create **Next**.
5. From the list of Available Roles, select the SSA_DEV_USERS role and move it to the Selected Roles table. Remove the EM_USER and PUBLIC roles from the Selected Roles table. Click **Next**.
6. Accept the default target privileges and click **Next**.
7. Accept the default resource privileges and click **Next**.
8. Review all the changes and click **Finish** to create the SSA_USER1 user.

Tip: To create multiple users with the same role, select the newly created user and click **Create Like**. This will create a new user that will have the same properties as the source. You can then update the name, description, and email address for the new user.

Note: Repeat these steps to create other users. For users with the EM_CLOUD_ADMINISTRATOR and EM_SSA_ADMINISTRATOR roles, the EM_USER and PUBLIC roles must be retained as these users need access to additional features.

3.5 Configuring LDAP Authentication

Oracle Enterprise Manager provides tools and procedures to help you ensure that you are managing your Oracle environment in a secure manner. Enterprise Manager's authentication framework consists of pluggable authentication schemes that let you use the type of authentication protocol best suited to your environment. The following authentication schemes are available:

- Oracle Access Manager (OAM) SSO
- Repository-Based Authentication
- SSO-Based Authentication
- Enterprise User Security Based Authentication
- Oracle Internet Directory (OID) Based Authentication
- Microsoft Active Directory Based Authentication

Enterprise User Security (EUS) provides automatic authentication to users and roles from the LDAP compliant directory server.

For more details on Enterprise User Security, see the Enterprise Manager Cloud Control Administrator's Guide.

3.6 Configuring Privilege Delegation Settings

Privilege delegation allows a logged-in user to perform an activity with the privileges of another user. Sudo and PowerBroker are privilege delegation tools that allow a logged-in user to be assigned these privileges. These privilege delegation settings will be used for all provisioning and patching activities on these hosts.

To configure privilege delegation settings on cloud hosts, follow these steps:

1. Create a Privilege Setting Template.
 1. Log in to Enterprise Manager as a Super Administrator user.
 2. From the **Setup** menu, select **Security**, then select **Privilege Delegation**.
 3. Under the Related Links section, click the **Manage Privilege Delegation Setting Templates** link.
 4. Select **Sudo** or **PowerBroker** from the Create list and click **Go**.
 5. Enter a template name, and the Sudo or PowerBroker command to be used on the target hosts. Sample values are provided in the description for the command fields. For example, the command for sudo is `/usr/bin/sudo -u %RUNAS% %COMMAND%`
 6. Click **Save**.

Note:

- If you select the PowerBroker option, you can specify an optional value in the PowerBroker Password Prompt field.
 - Check the path to the **sudo** or **pbrun** executable. For example, if you are using **sudo**, you can check this by opening a terminal to one of the hosts and run the command `which sudo`. The command returns the path to the executable.
-

2. Deploy the template to the hosts.
 1. From the **Setup** menu, select **Security**, then select **Privilege Delegation**.
 2. Click **Manage Privilege Delegation Settings Templates** in the Related Links section.
 3. Select the template that you have created and click **Apply**.
 4. Click **Add Targets** and choose the hosts for which the template is to be applied.
 5. Click **Select** to select the hosts and click **Apply**.
 6. On the Past Apply Operations page, check the Status column for all hosts. A job has been submitted to all hosts to apply this privilege delegation setting.
 7. Refresh the page using the browser refresh button, or click **Go** on this page to refresh the status for all hosts.
 8. From the **Setup** menu, select **Security**, then select **Privilege Delegation** to navigate to the Privilege Delegation page. Click the **Show** link in the Status column to confirm that the privilege delegation settings have been applied on all hosts. You can install on a maximum of 16 servers at a time. The total time required will increase if the Management Agent is installed on a large number of servers.

For more details on the configuring privilege delegation settings, see the Enterprise Manager Cloud Control Administrator's Guide.

3.7 Customizing the Self Service Login Page

You can configure Enterprise Manager and provide specific access to SSA users. To configure Enterprise Manager for SSA users, you must set some properties on the OMS and copy the required images to a specified directory. This section describes the following:

- [Configuring the Self Service Login Page](#)
- [Switching Back to the Enterprise Manager Login Page](#)
- [Routing SSA Requests to a Specific OMS Pool](#)

3.7.1 Configuring the Self Service Login Page

To launch a separate SSA login page for all SSA users, you must do the following:

- Set the following mandatory property on all OMSes:

```
$ORACLE_HOME/bin/emctl set property -name oracle.sysman.ssa.logon.ssa_oms -value true
```

If this property is not set to *true*, the standard Enterprise Manager login page is displayed.
- Set the following optional OMS properties.
 - ```
$ORACLE_HOME/bin/emctl set property -name oracle.sysman.ssa.logon.show_cloud_provider_brand -value true
```

If this property is not set to *true*, the default Oracle Enterprise Manager 12c logo is displayed.



- `$ORACLE_HOME/bin/emctl set property -name oracle.sysman.ssa.logon.show_cloud_tenant_brand -value true`

If this property is not set to *true*, the tenant logo is not displayed.

- `$ORACLE_HOME/bin/emctl set property -name oracle.sysman.ssa.logon.cloud_provider_alt_text -value "Cloud Provider"`
- `$ORACLE_HOME/bin/emctl set property -name oracle.sysman.ssa.logon.cloud_tenant_alt_text -value "Cloud Tenant"`

These properties are optional and if not set, the default values for "Cloud Provider", and "Cloud Tenant" are displayed.

- `$ORACLE_HOME/bin/emctl set property -name oracle.sysman.ssa.logon.show_disclaimer_text -value true`

If this property is not set to *true*, the default Oracle copyright message is displayed.

- `$ORACLE_HOME/bin/emctl set property -name oracle.sysman.ssa.logon.disclaimer_text -value "Customer specified Disclaimer text"`

If this property is set to *true*, the specified disclaimer text is displayed instead of the default Oracle copyright message.

- `$ORACLE_HOME/bin/emctl set property -name oracle.sysman.ssa.logon.show_em_branding_text -value true`

If this property is not set to *false*, the "Powered by Oracle Enterprise Manager" text will appear on the Self Service Login page.

- Copy the following images to the `$ORACLE_HOME/sysman/config/` directory.

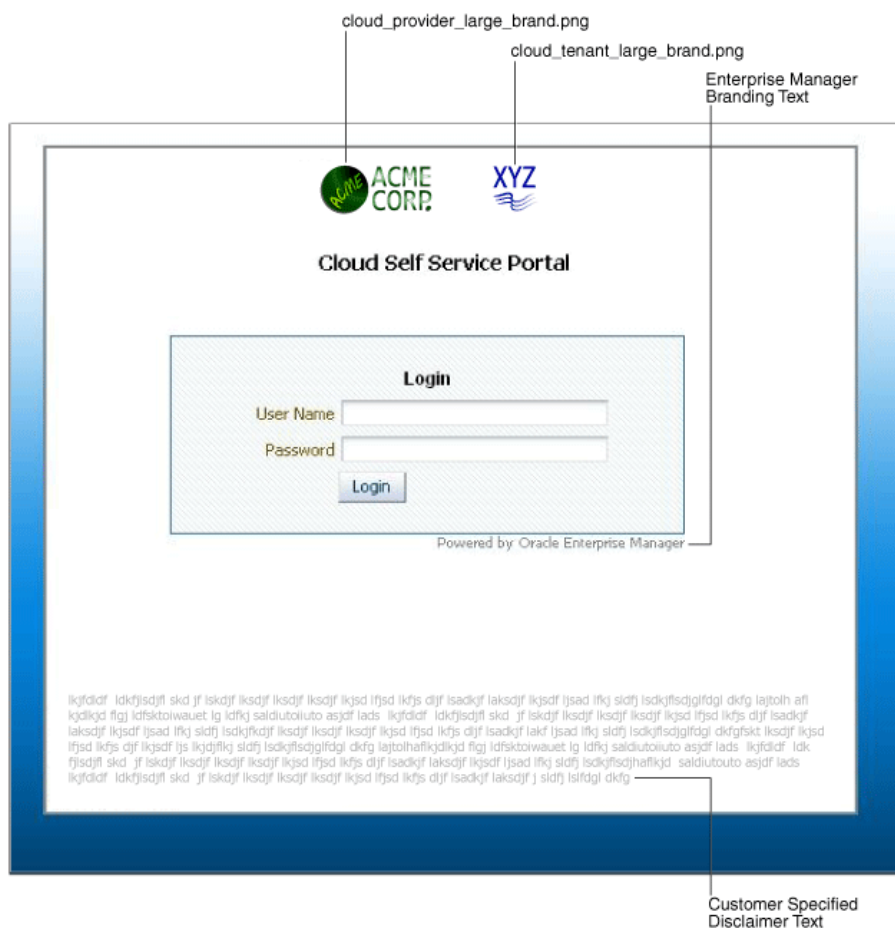
- `cloud_provider_small_brand.png`
- `cloud_tenant_small_brand.png`

If a single image is used, the maximum recommended size is 500 \* 20 px. If 2 images are used, the maximum recommended size is 200 \* 20 px per image. After login, these images are displayed instead of the Oracle logo, if the OMS properties `oracle.sysman.ssa.logon.show_cloud_provider_brand` and `oracle.sysman.ssa.logon.show_cloud_tenant_brand` are set. If the OMS property `oracle.sysman.ssa.logon.show_cloud_provider_brand` is not set to *true*, along with the tenant logo, the default Oracle logo appears.

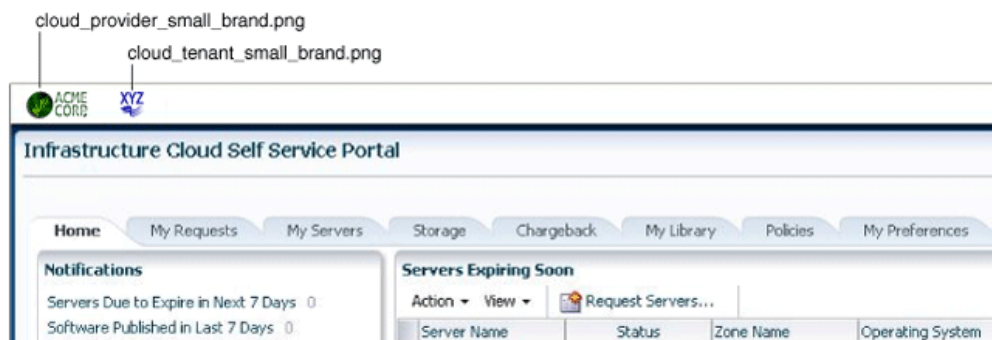
- `cloud_provider_large_brand.png`
- `cloud_tenant_large_brand.png`

If a single image is used, then the maximum recommended size is 525 \* 60 px. If 2 images are used, the maximum recommended size is 250 \* 50 px per image. These images are displayed on the login page, if the OMS properties `oracle.sysman.ssa.logon.show_cloud_provider_brand` and `oracle.sysman.ssa.logon.show_cloud_tenant_brand` are set.

For example, if ACME Corp is the Cloud Service Provider and XYZ is the Cloud Tenant, the customized login page appears as follows:

**Figure 3–2 Customized SSA Login Page**

After the SSA user has logged in, the customized Infrastructure Self Service Portal is displayed as shown below:

**Figure 3–3 Customized Post-Login Page**

### 3.7.2 Switching Back to the Enterprise Manager Login Page

To revert to the default Enterprise Manager login page, set the following property:

```
$ORACLE_HOME/bin/emctl set property -name oracle.sysman.ssa.logon.ssa_oms
-value false
```

### 3.7.3 Routing SSA Requests to a Specific OMS Pool

Oracle Management Service (OMS) is one of the core components of Enterprise Manager Cloud Control that works with the Oracle Management Agents (Management Agents) and plug-ins to discover targets, monitor and manage them, and store the collected information in a repository for future reference and analysis.

When you install Enterprise Manager for the very first time, by default, one OMS is installed along with one Management Agent. This default configuration is suitable for small environments. In larger production environments with several SSA users, you may need to install additional OMS instances to reduce the load on a single OMS and improve the efficiency of the data flow. You can then configure the Server Load Balancer (SLB) to redirect all SSA requests to a specific OMS pool. The other OMS pools will then be available for administration usage. To learn more about setting up multiple OMS instances and the SLB, see *Adding Additional Oracle Management Service* section in the *Enterprise Manager Cloud Control Basic Installation Guide*.

To redirect SSA requests, you must specify the following SLB configuration:

```
https://<slb_host_name>:<slb_em_port>/em redirecting to oms for em
```

```
https://<slb_host_name>:<slb_ssa_port>/em redirecting to oms for ssa
```

The SSA and non-SSA OMS pools are differentiated based on the port number. All requests with a particular port number will be redirected to a specific OMS pool (SSA OMS pool) and all the other requests will be redirected to the other pool.



# Part III

---

## Setting Up and Using Infrastructure as a Service

This section describes how to set up infrastructure as a service, manage virtualization targets, deploy guest virtual machines, set up, and use the self service application.

It contains the following chapters:

- [Chapter 4, "Setting Up Cloud for IaaS"](#)
- [Chapter 6, "Setting Up the IaaS Self Service Portal"](#)
- [Chapter 8, "Administering and Monitoring an IaaS Cloud"](#)
- [Chapter 5, "Defining and Managing Cloud Policies"](#)



---

## Setting Up Cloud for IaaS

This chapter describes the server, storage and networking infrastructure that needs to be defined for providing Infrastructure as a Service (IaaS) cloud using the Oracle VM technology. It contains the following sections:

- [Getting Started](#)
- [Cloud Infrastructure Privileges](#)
- [Registering the Oracle VM Manager](#)
- [Synchronizing the Oracle VM Manager Targets](#)
- [Discovering a Virtual Server](#)
- [Rediscovering a Virtual Server](#)
- [Setting Up Networks](#)
- [Creating a Network Profile](#)
- [Registering Storage Servers](#)
- [Creating a Storage Repository](#)
- [Presenting the Storage Repository](#)
- [Using Self Update to Download Oracle VM Templates and Virtual Assemblies](#)
- [Importing Assemblies, Templates, and Virtual Disks into the Storage Repository](#)
- [Creating a Virtual Server Pool](#)
- [Creating a Zone](#)
- [Setting Up the Storage Quality of Service \(QoS\)](#)
- [Patching the Oracle VM Server](#)
- [Acknowledging OVM Events](#)
- [Managing the NTP Configuration](#)
- [Importing Virtual Machines](#)
- [Acknowledging the Repository Events](#)

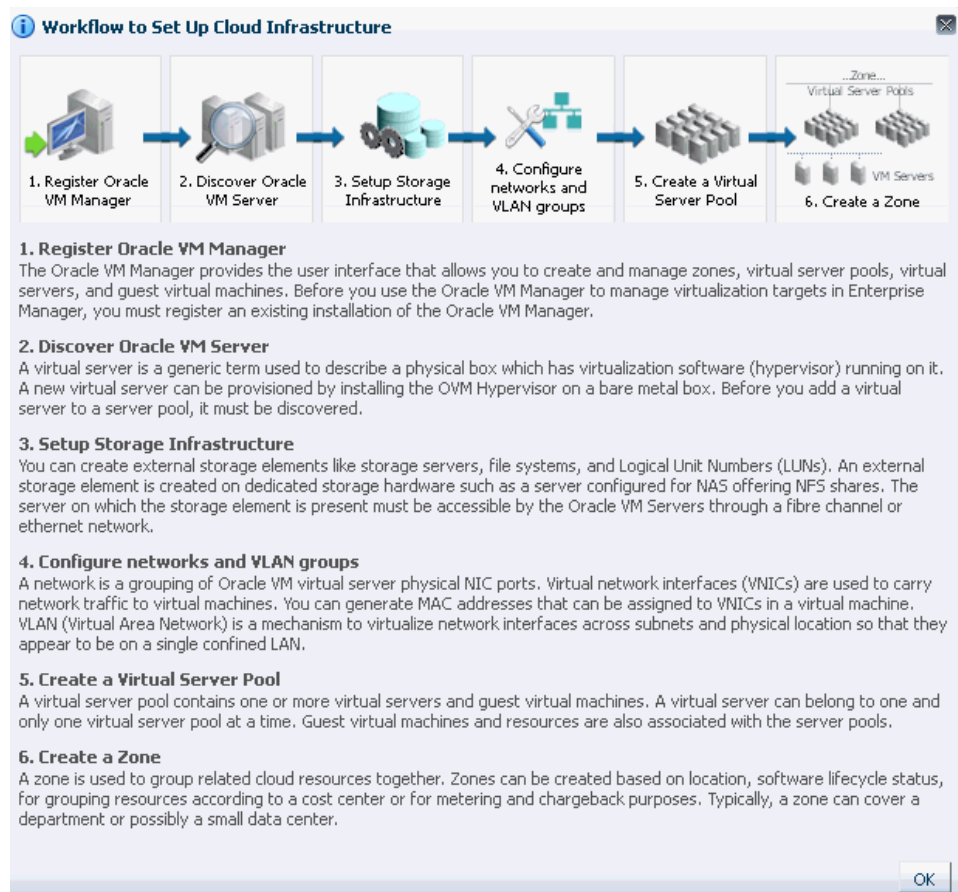
### 4.1 Getting Started

This section helps you get started with this chapter by providing an overview of the steps involved in setting up the virtualization infrastructure. Consider this section to be a documentation map to understand the sequence of actions you must perform to successfully set up virtualized systems.

1. Before you start monitoring and managing virtualization targets, you must register the Oracle VM Manager with Oracle Enterprise Manager. See [Section 4.3, "Registering the Oracle VM Manager"](#) for details.
2. You can then discover one or more Oracle VM Server targets. See [Section 4.5, "Discovering a Virtual Server"](#) for details.
3. Set up the storage infrastructure for the cloud. See [Section 4.9, "Registering Storage Servers"](#) for details.
4. Configure networks and VLAN Groups. See [Section 4.7, "Setting Up Networks"](#) for details.
5. Create a virtual server pool. See [Section 4.14, "Creating a Virtual Server Pool"](#) for details.
6. Create a zone by grouping one or more server pools. See [Section 4.15, "Creating a Zone"](#) for details.

Figure 4–1 shows the sequence of steps involved in setting up the virtualization infrastructure.

**Figure 4–1 Setting Up the Cloud Infrastructure**



## 4.2 Cloud Infrastructure Privileges

This section lists the actions that can be performed on various targets in the Infrastructure Cloud Home page when you use:



- Right click on a target in the left panel.
- Select a target on the Members region on the parent target home page and then select the **Action** menu.
- Select the **Target** menu of the corresponding Target Home page.

**Table 4–1 Infrastructure Cloud**

| Action Label         | Multi-select / Single Select | Required Target State | Privileges Required              |
|----------------------|------------------------------|-----------------------|----------------------------------|
| Register OVM Manager | Single                       |                       | Add any Target                   |
| Setup                | Single                       |                       | View on Infrastructure Cloud     |
| Members              | Single                       |                       | Operator on Infrastructure Cloud |

**Table 4–2 OVM Manager**

| Action Label                   | Multi-select / Single Select | Required Target State | Privileges Required                         |
|--------------------------------|------------------------------|-----------------------|---------------------------------------------|
| Members                        | Single                       |                       | View On OVM Manager                         |
| Edit                           | Single                       |                       | Operator on OVM Manager                     |
| Synchronize                    | Single                       |                       | Full on OVM Manager                         |
| Deregister                     | Single                       |                       | Full on OVM Manager and all its sub targets |
| Create Zone                    | Single                       |                       | Operator on OVM Manager and Add any Target  |
| Create Virtual Server Pool     | Single                       |                       | Operator on OVM Manager and Add any Target  |
| Discover Virtual Server        | Single                       |                       | Operator on OVM Manager and Add any Target  |
| Manage Unowned Virtual Servers | Single                       |                       | Operator on OVM Manager                     |
| Manage Network                 | Single                       |                       | Operator on OVM Manager                     |
| Manage Storage                 | Single                       |                       | Operator on OVM Manager                     |
| Manage Storage Repository      | Single                       |                       | Operator on OVM Manager                     |
| Manage Yum Repository          | Single                       |                       | Full on OVM Manager                         |
| Manage NTP Configuration       | Single                       |                       | Full on OVM Manager                         |

**Table 4–3 Zone**

| Action Label | Multi-select / Single Select | Required Target State | Privileges Required |
|--------------|------------------------------|-----------------------|---------------------|
| Members      | Single                       |                       | View On Zone        |
| Edit         | Single                       |                       | Operator on Zone    |
| Delete       | Single                       |                       | Full on Zone        |

**Table 4–3 (Cont.) Zone**

| Action Label                        | Multi-select / Single Select | Required Target State | Privileges Required                 |
|-------------------------------------|------------------------------|-----------------------|-------------------------------------|
| Delete only from Enterprise Manager | Single                       |                       | Full on Zone                        |
| Deploy Assembly                     | Single                       |                       | Operator on Zone and Add any Target |
| Deploy Template                     | Single                       |                       | Operator on Zone and Add any Target |
| Deploy ISO                          | Single                       |                       | Operator on Zone and Add any Target |
| Deploy PXE                          | Single                       |                       | Operator on Zone and Add any Target |
| Import Virtual Machine              | Single                       |                       | Operator on Zone and Add any Target |

**Table 4–4 Virtual Server Pool**

| Action Label                        | Multi-select / Single Select | Required Target State | Privileges Required                                |
|-------------------------------------|------------------------------|-----------------------|----------------------------------------------------|
| Members                             | Single                       |                       | View On Virtual Server Pool                        |
| Edit                                | Single                       |                       | Operator on Virtual Server Pool                    |
| Delete                              | Single                       |                       | Full on Virtual Server Pool                        |
| Delete only from Enterprise Manager | Single                       |                       | Full on Virtual Server Pool                        |
| Manage DRS / DPM Policies           | Single                       |                       | Operator on Virtual Server Pool                    |
| Deploy Assembly                     | Single                       |                       | Operator on Virtual Server Pool and Add any Target |
| Deploy Template                     | Single                       |                       | Operator on Virtual Server Pool and Add any Target |
| Deploy ISO                          | Single                       |                       | Operator on Virtual Server Pool and Add any Target |
| Deploy PXE                          | Single                       |                       | Operator on Virtual Server Pool and Add any Target |
| Import Virtual Machine              | Single                       |                       | Operator on Zone and Add any Target                |

**Table 4–5 Virtual Server**

| Action Label                        | Multi-select / Single Select | Required Target State | Privileges Required        |
|-------------------------------------|------------------------------|-----------------------|----------------------------|
| Edit                                | Single                       | Up                    | Operator on Virtual Server |
| Deregister                          | Single                       |                       | Full on Virtual Server     |
| Delete only from Enterprise Manager | Single                       |                       | Full on Virtual Server     |
| Start Maintenance                   | Multiple                     | Up                    | Operator on Virtual Server |
| Stop Maintenance                    | Multiple                     | Blackout              | Operator on Virtual Server |

**Table 4–5 (Cont.) Virtual Server**

| Action Label              | Multi-select / Single Select        | Required Target State | Privileges Required                           |
|---------------------------|-------------------------------------|-----------------------|-----------------------------------------------|
| Upgrade                   | Multiple                            |                       | Full on Virtual Server                        |
| Start                     | Multiple                            | Down                  | Operator on Virtual Server                    |
| Stop                      | Multiple                            | Up                    | Operator on Virtual Server                    |
| Restart                   | Multiple                            | Up                    | Operator on Virtual Server                    |
| Rediscover Virtual Server | Single                              |                       | Full on Virtual Server                        |
| Deploy Assembly           | Single                              |                       | Operator on Virtual Server and Add any Target |
| Deploy Template           | Single                              |                       | Operator on Virtual Server and Add any Target |
| Deploy ISO                | Single                              |                       | Operator on Virtual Server and Add any Target |
| Deploy PXE                | Single                              |                       | Operator on Virtual Server and Add any Target |
| Import Virtual Machine    | Single                              |                       | Operator on Zone and Add any Target           |
| Members                   | Multi-select / Single Select:Single |                       | View on Virtual Server                        |

**Table 4–6 Guest VM**

| Action Label                        | Multi-select / Single Select | Required Target State | Privileges Required  |
|-------------------------------------|------------------------------|-----------------------|----------------------|
| Edit                                | Single                       |                       | Operator on Guest VM |
| Delete                              | Multiple                     | Down, Metric Error    | Full on Guest VM     |
| Delete only from Enterprise Manager | Multiple                     |                       | Full on Guest VM     |
| Migrate                             | Multiple                     | Up, Down              | Operator on Guest VM |
| Move                                | Single                       | Down                  | Operator on Guest VM |
| Clone                               | Single                       | Up, Down              | Operator on Guest VM |
| Save as Template                    | Single                       | Down                  | Operator on Guest VM |
| Start                               | Multiple                     | Down                  | Operator on Guest VM |
| Stop                                | Multiple                     | Up                    | Operator on Guest VM |
| Restart                             | Multiple                     | Up                    | Operator on Guest VM |
| Kill                                | Multiple                     | Up                    | Operator on Guest VM |
| Stop and Start                      | Multiple                     | Up                    | Operator on Guest VM |
| Suspend                             | Multiple                     | Up                    | Operator on Guest VM |
| Resume                              | Multiple                     |                       | Operator on Guest VM |
| Launch VNC Console                  | Single                       |                       | Operator on Guest VM |

**Table 4–7 Assembly**

| Action Label | Multi-select / Single Select | Required Target State | Privileges Required  |
|--------------|------------------------------|-----------------------|----------------------|
| Start        | Single                       |                       | Operator on Assembly |
| Stop         | Single                       |                       | Operator on Assembly |
| Delete       | Single                       |                       | Operator on Assembly |
| Scale Up     | Single                       |                       | Operator on Assembly |
| Scale Down   | Single                       |                       | Operator on Assembly |
| Add Nodes    | Single                       |                       | Operator on Assembly |

**Table 4–8 Tier**

| Action Label | Multi-select / Single Select | Required Target State | Privileges Required |
|--------------|------------------------------|-----------------------|---------------------|
| Start        | Single                       |                       | Operator on Tier    |
| Stop         | Single                       |                       | Operator on Tier    |
| Scale Up     | Single                       |                       | Operator on Tier    |
| Scale Down   | Single                       |                       | Operator on Tier    |

## 4.3 Registering the Oracle VM Manager

The Oracle VM Manager provides the user interface that allows you to create and manage zones, virtual server pools, virtual servers, and guest virtual machines. Before you use the Oracle VM Manager to manage virtualization targets in Enterprise Manager, you must register an existing installation of the Oracle VM Manager. To register the Oracle VM Manager, you must first add it as an Enterprise Manager target. To do so, follow these steps:

---

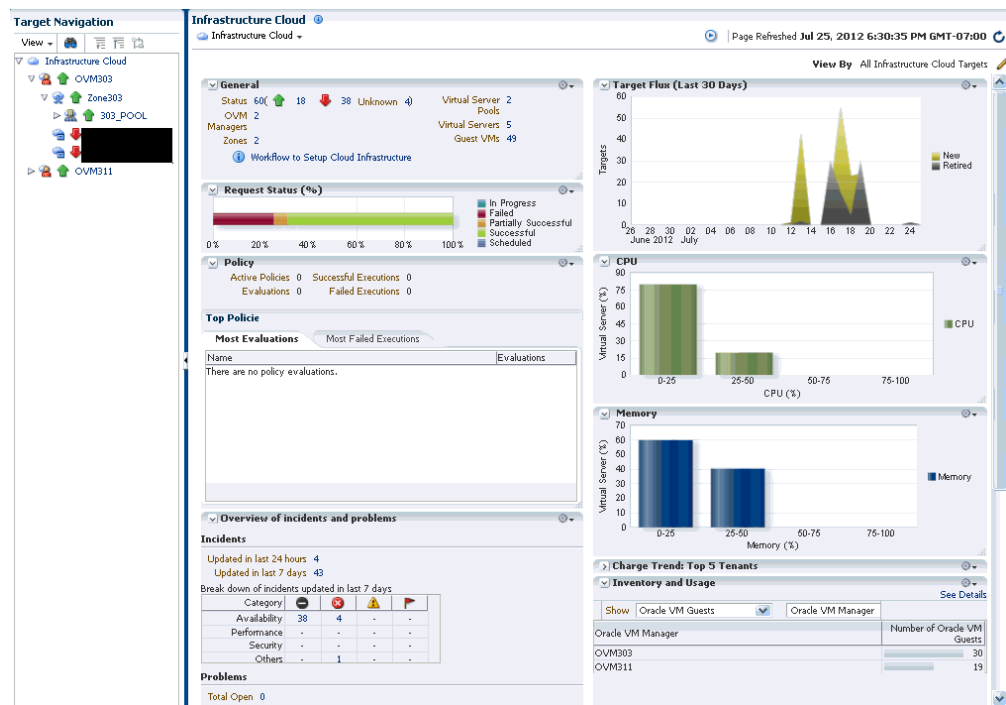
---

**Note:**

- The Oracle VM Manager that you register with Enterprise Manager must be at least version 3.0.2 or higher.
  - The Oracle VM Manager must be associated with only one Enterprise Manager installation.
  - As a best practice, it is recommended that the Management Agent and the Oracle VM Manager are present on the same physical machine.
- 
- 

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.

Figure 4–2 Infrastructure Cloud Home



- From the Infrastructure Cloud menu, select **Register OVM Manager**. The Add Oracle VM Manager page appears.

Figure 4–3 Add Oracle VM Manager

Infrastructure Cloud Page Refreshed Jun 2, 2013 4:17:53 PM GMT-07:00

[Submit](#) [Cancel](#)

### Register Oracle VM Manager

Register an existing installation of Oracle VM Manager with Enterprise Manager. Make sure that a connection can be established using the specified URL. Ensure that the Oracle VM Manager certificate has been imported into the Agent Keystore.

\* Name

\* Monitoring Agent

---

#### Oracle VM Manager

##### Connection URLs

\* Oracle VM Manager URL

☒ **TIP** URL for an existing Oracle VM Manager installation. e.g. `tcp://localhost:54321`, `tcps://<server.domain>:54322`

\* Oracle VM Manager Console URL

☒ **TIP** URL for the Oracle VM Manager Console e.g. `http://<server.domain>:<port>/`, `https://<server.domain>:<port>/`

##### Monitoring Credentials

Specify the credentials to be used for monitoring Oracle VM Manager.

\* Username

\* Password

☐ Automatic Synchronization with Oracle VM Manager

##### Administration Credentials

Specify the credentials to be used for administration of Oracle VM Manager. If not specified, it defaults to the monitoring credentials.

☐ Use Administration Credentials

\* Username

\* Password

You can register an existing installation of Oracle VM Manager by following these steps:

### Prerequisites

- If you are registering OVM Manager 3.1.1 or later, before you update the Oracle VM Manager URL, you must have configured tcps as described in Note 1490283.1.
- You must have imported the Oracle VM Manager certificate to the Agent keystore by following these steps:
  - Export the OVM Manager certificate `<JAVA_HOME>/bin/keytool -keystore <OVM_MANAGER_HOME>/ovmmCoreTcps.ks -exportcert -alias ovmm -file <file_loc_for_certificate>`
  - Import the OVM Manager certificate `<AGENT_INSTANCE_HOME>/bin/emctl secure add_trust_cert_to_jks -trust_certs_loc <file_loc_for_certificate> -alias <alias_name>`
- 1. Enter the target name for the Oracle VM Manager instance to be registered.
- 2. Select the Management Agent to monitor and manage the Oracle VM Manager instance to be registered. It is recommended that the Management Agent (with *Enterprise Manager for Oracle Virtualization* plug-in) is present on the same machine where the Oracle VM Manager instance is running.
- 3. Enter the Oracle VM Manager URL. Refer to Note 1490283.1 for details on configuring Oracle VM Manager for tcps access.

The URL must be in one of the following formats:

– `tcp://localhost:<port>`. Example: `tcp://localhost:54321` (If the Oracle VM Manager version is 3.1.1 or greater, you can specify only localhost in the URL)

– `tcps://<hostname>:<port>`. Example: `tcps://server.domain:54322`

- 4. Select the **Automatic Synchronization** checkbox to automatically synchronize the data from Oracle VM Manager into Enterprise Manager at fixed intervals of time. If you do not wish to enable this option, you can manually synchronize Oracle VM Manager with Enterprise Manager. For more details on the automatic and manual synchronization options, see [Section 4.4, "Synchronizing the Oracle VM Manager Targets"](#) for details.
- 5. Enter the administrator user name and password to connect to the Oracle VM Manager URL provided in **Step C**. If you want to use separate credentials for managing Oracle VM Manager, check the **Use Administration Credentials** check box. If these credentials are not specified, the monitoring credentials are used by default.
- 6. Enter the Oracle VM Manager's Console URL. This is the URL to enable the VNC Console for the Guest VMs. For example, `https://ovmmgr.example.com:7002/ovm/console`  
The URL must be in the following format:  
`https://<hostname>:<port>`
- 7. Click **Submit** to register the Oracle VM Manager. You can now use the Oracle VM Manager to monitor and manage the virtualization targets in Enterprise Manager.

### 4.3.1 Updating the Oracle VM Manager URL

If you had Oracle VM Manager 3.0.x registered with Enterprise Manager previously and had upgraded it to versions 3.1.x or 3.2.1, you must update the Oracle VM Manager URL.

### Prerequisites

- Before you update the Oracle VM Manager URL, you must have configured tcps as described in Note 1490283.1.
- You must have imported the Oracle VM Manager certificate to the Agent keystore by following these steps:
  - Export the OVM Manager certificate `<JAVA_HOME>/bin/keytool -keystore <OVM_MANAGER_HOME>/ovmmCoreTcps.ks -exportcert -alias ovmm -file <file_loc_for_certificate>`
  - Import the OVM Manager certificate `<AGENT_INSTANCE_HOME>/bin/emctl secure add_trust_cert_to_jks -trust_certs_loc <file_loc_for_certificate> -alias <alias_name>`

To update the Oracle VM Manager URL, follow these steps:

1. Right click on the OVM Manager target that is to be upgraded and select **Edit**.
2. Enter a tcps based (tcps://hostname:port) URL in the Oracle VM Manager field.

---

**Note:** It is recommended that you use the tcps based URL. But if the Management Agent is running on the same host on which Oracle VM Manager has been installed, you can use the tcp based URL (tcp://localhost:54321).

---

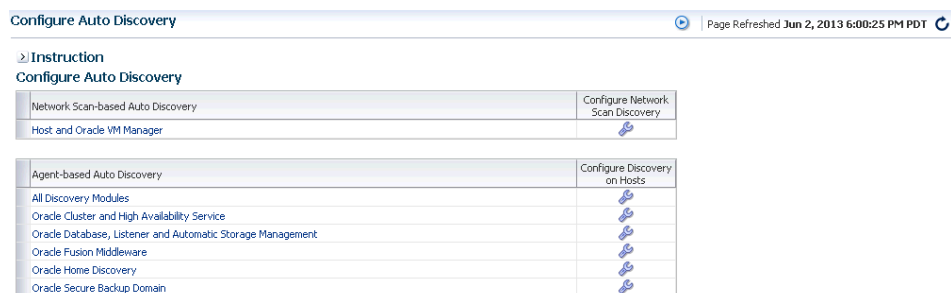
### 4.3.2 Discovering and Promoting the Oracle VM Manager (Optional)

Automatic discovery is a process that detects hosts, Oracle VM Managers, and other targets so that they can be monitored and managed by Enterprise Manager. If you know the Oracle VM Manager instances in your environment, see [Section 4.3, "Registering the Oracle VM Manager"](#) for details on registering them with Enterprise Manager. Discovery and promotion of Oracle VM Manager should be attempted if you want to scan your networks to discover unknown Oracle VM Manager instances.

To discover an Oracle VM Manager, follow these steps:

1. Log in to Enterprise Manager as a user with the Super Administrator role.
2. From the **Setup** menu, select **Add Target**, then select **Configure Auto Discovery**.

**Figure 4–4 Configure Auto Discovery**



3. In the Network Scan-based Auto Discovery table, click **Configure Network Scan Discovery** in the Host and Oracle VM Manager row. The Network Discovery page appears. Use this page to search, browse, and create OVM Manager discovery specifications.

**Figure 4–5 Network Scan Discovery**

**Network Scan Discovery** Page Refreshed Jun 2, 2013 6:03:51 PM PDT

Configure Auto Discovery > Network Scan Discovery

Use this page to search, browse, and create specifications for the discovery of host operating systems and virtual servers using Network Scan. Schedule the discovery on a configurable interval. Once the specification is created it can be reused later to scan again.

**Search**

View

| Name                | Owner | Discovered Targets | Scans | Scanning Agents | Status | Start Time | End Time | Description |
|---------------------|-------|--------------------|-------|-----------------|--------|------------|----------|-------------|
| No data to display. |       |                    |       |                 |        |            |          |             |

**Scan Status:**

Latest Scan

**Search**

View

| Scanning Agent      | Status | Discovered Targets | Start Time | End Time | IP Ranges Scanned |
|---------------------|--------|--------------------|------------|----------|-------------------|
| No data to display. |        |                    |            |          |                   |

- Click **Create**. The Network Scan Discovery: Create page appears.

**Figure 4–6 Network Scan Discovery: Create**

**Network Scan Discovery: Create**

Specify Network Scan Discovery to discover host operating systems and virtual servers. Schedule the discovery on a configurable interval.

☒ To perform Network Scan discovery, configure the scanning agent hosts with Sudo Privilege Delegation, and use credentials that have Run As 'root' attribute set.

\* Name: Host Discovery Jun 2, 2013 6:07:01 PM PDT

Description:

Owner: SYSMAN

**Network Scans**

View

| IP Ranges to Scan   | Scanning Agent |
|---------------------|----------------|
| No data to display. |                |

**Exclude IP Ranges**

**Schedule**

Start: ☒ Immediately ☐ Later  (GMT-08:00) Los Angeles - Pacific Time (PT)

Repeat:

**Credentials**

The discovery Network Scan is run as root. It is required that you set privileged host credentials or named credentials that use Sudo.

Credential: ☒ Preferred ☐ Named ☐ New

Preferred Credential Name:  Privileged Host Credentials

Credential Details: Credentials will be determined at runtime.

- Enter the Name and Description for the Discovery Specification. In the Scan Details tab, select the Agent that is to be used perform the IP scan. For each Agent, specify the IP ranges for the scan. The following formats are supported: host name, IP address, and IP range. Use space to separate values.
- Use the Job Details tab on the Host Discovery Specification Create page to schedule the discovery IP scan, allowing you to start it immediately or at a later specified time. Use the Credentials section to enter preferred credentials for the host. You can choose from Preferred Credentials, Named Credentials, or New Credentials. The discovery IP Scan is run as root. It is required that you set privileged host credentials or named credentials that use Sudo.
- You can use the Host Discovery page to check the status of discovery and view newly discovered OVM Managers.
- After the target has been discovered, from the **Setup** menu, select **Add Target**, then select **Auto Discovery Results**. Enterprise Manager displays the Auto Discovery Results page where you can review discovered unmanaged targets and promote targets to be managed by Enterprise Manager for monitoring.



- On the Network-scanned Targets tab, the target discovery results are displayed in the table at the bottom. Select the OVM Manager target and click **Promote**. The OVM Manager Registration page is displayed. See [Section 4.3, "Registering the Oracle VM Manager"](#) for details.

### 4.3.3 Monitoring Configuration (OVM Manager)

To customize or view the monitoring configuration details for the OVM Manager target, follow these steps:

- From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
- Click on an OVM Manager target in the left panel.
- From the **VM Manager** menu, select **Target Setup**, then select **Monitoring Configuration**. The Monitoring Configuration page for the OVM Manager target appears.

**Figure 4–7 Monitoring Configuration (OVM Manager)**

VM Manager ▾ Page Refreshed Jun 2, 2013 9:50:00 PM EDT ↻

**Monitoring Configuration** OK Cancel

Current Operation  
LiveSyncEnabled: false

Timeout Settings For Network Configuration Of Deployment (minutes): 60

OVM Job Retry Count: 5

OVM Job Retry Sleep Interval: 5

Timeout value for short running OVM Job (minutes): 5

OVM Manager Console URL: https://[REDACTED]

OVM Manager URL: tcp://[REDACTED]

Oracle VM Manager Password: \*\*\*\*\*

Oracle VM UUID: 0004fb0000010000f46fcbef9dcf200

Oracle VM Manager Username: \*\*\*\*\*

Timeout Settings For Product Configuration Of Deployment (minutes): 60

Timeout Value for Checking if SSH service is Running on a deployed Virtual Machine (minutes): 10

SyncDataFetchTime: 2013-05-21T08:54:18.129PDT

SyncStatus: SYNC\_SUCCEEDED

Deployed on Exalogic: false

**Monitoring**  
Oracle has automatically enabled monitoring for this target's availability and performance, so no further monitoring configuration is necessary. You can edit the metric thresholds from the target's homepage.

- The following details are displayed:
  - LiveSyncEnabled:** Indicates if auto synchronization is enabled between Enterprise Manager and Oracle VM Manager. By default, this is false. If enabled, the data is synchronized every 5 minutes.
  - Timeout Settings For Network Configuration of Deployment:** The total wait time allowed for network configuration during assembly deployment. The default is 60 minutes.
  - OVM Job Retry Count:** The number of attempts that can be made after which the OVM Job should be retired. The default number of attempts is 5.
  - OVM Job Retry Sleep Interval:** The interval between two subsequent OVM job retry attempts. The default interval is 5 seconds.
  - Timeout value for short running OVM Job (minutes):** The time interval allowed for an OVM job to be completed.
  - OVM Manager Console URL:** The URL used to launch VNC consoles for VM.
  - OVM Manager URL:** The OVM Manager registration URL used for communication between Enterprise Manager and Oracle VM Manager.

- **Oracle VM Manager Password:** The password required to access Oracle VM Manager.
  - **Oracle VM UUID:** The internal Oracle VM identifier.
  - **Oracle VM Manager Username:** The username required to access Oracle VM Manager
  - **Timeout Settings for Product Configuration of Deployment:** The total wait time allowed for product configuration during assembly deployment. The default is 60 minutes.
  - **Timeout Value for Checking if SSH Service is Running on a Deployed Virtual Machine:** The total wait time allowed to check if the SSH service is running on a Guest VM. This parameter is used as prerequisite for deploying the Management Agent on the Guest VM.
  - **SyncDataFetchTime:** The time at which the Oracle VM Manager was last synchronized.
  - **SyncStatus:** The status of the Oracle VM Manager synchronization.
  - **Deployed on Exalogic:** Indicates whether the Oracle VM Manager is inside an Exalogic rack.
5. After editing the entries, click **OK** to return to the previous page.

## 4.4 Synchronizing the Oracle VM Manager Targets

Enterprise Manager can monitor the OVM Manager and the various virtualization targets registered with the OVM Manager. The targets registered with the OVM Manager must also be monitored targets in Enterprise Manager. If some targets have been created directly in OVM Manager, you can synchronize them in two ways: options:

- Automatic Synchronization
- Manual Synchronization

### 4.4.1 Automatic Synchronization

You can enable this option when Oracle VM Manager is registered or when you Edit Oracle VM Manager. If the **Automatic Synchronization** option is enabled, all changes made in Oracle VM Manager are automatically reflected in Enterprise Manager at fixed intervals. The default interval is 1 minute. To change the default interval, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Select an Oracle VM Manager target from the right panel to navigate to the Oracle VM Manager Home page.
3. From the **Target** menu, select **Monitoring**, then select **Metrics and Collection Settings**.
4. Click **Other Collected Items** and then click the **Collection Schedule** link for the OVM Manager Events Data metric.
5. Modify the frequency as required, click **Continue**, then click **OK**.

---

---

**Note:** If you do not enable automatic synchronization at this time, you can enable it later by following these steps:

- Navigate to the OVM Manager homepage.
  - From the **Target** menu, select **Target Setup**, then select **Edit OVM Manager**.
  - Select the **Automatic Synchronization** check box and click **Submit**.
- 
- 

## 4.4.2 Manual Synchronization

When the Oracle VM Manager is registered, if the **Automatic Synchronization** option is enabled, for the targets created in Oracle VM Manager, corresponding targets are created in Enterprise Manager. If, this option is not enabled, you can manually synchronize the targets by following these steps:

1. From the **Enterprise** menu, select **Infrastructure Cloud**, then select **Home**.
2. Right click on an OVM Manager target from the left panel and select **Synchronize**.
3. A confirmation message appears. Click **Submit** to submit a job to perform synchronization. The targets in Enterprise Manager and OVM Manager are synchronized.
  - Targets that are not present in OVM Manager will be removed.
  - For targets present only in OVM Manager, corresponding targets are created in Enterprise Manager.
  - Configuration collections are refreshed for all other targets.

## 4.5 Discovering a Virtual Server

A virtual server is a generic term used to describe a physical box which has virtualization software (hypervisor) running on it. A new virtual server can be provisioned by installing the OVM Hypervisor on a bare metal box. Before you add a virtual server to a server pool, it must be discovered.

To discover an Oracle VM Server, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Right click on an OVM Manager target from the left panel and select **Discover Virtual Server**.

**Figure 4–8 Discover Virtual Servers**

ovm\_manager

Oracle VM Manager

Page Refreshed Jun 2, 2011 5:09:26 AM GMT

### Discover Virtual Servers

Submit Cancel

#### Virtual Servers

Specify long host name(FQDN) or IP address. To enter multiple entries, enter each new host name/IP address/IP range in a new line

Oracle VM Agent Credentials

Specify the Oracle VM Agent Username and Password for the virtual servers to be discovered.

\* User Name oracle

\* Password

3. Specify the host name or the IP address of the virtual server. You can enter multiple host names or IP addresses on separate lines.
4. Enter the user name and password of the Oracle VM Agent running on the virtual server to be discovered.
5. Click **Submit** to discover the virtual server.

The newly discovered Oracle VM Server contains some basic information about itself, and about any immediate connectivity to a network storage, but it is considered to be in an unconfigured state. After the storage and networking has been configured, the Oracle VM virtual servers are ready to be used as the infrastructure for cloud.

## 4.6 Rediscovering a Virtual Server

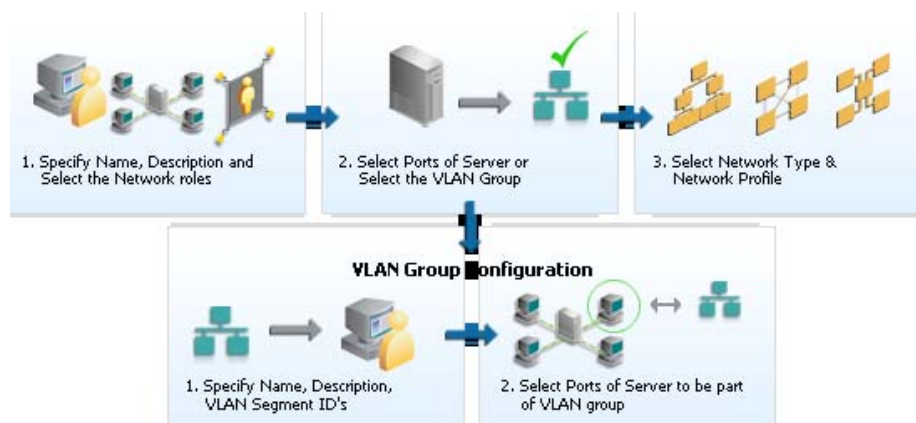
To rediscover a virtual server, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Right click on a Virtual Server target in the left panel and select **Rediscover Virtual Server**.
3. A confirmation message appears. Click **OK** to rediscover the virtual server and discover any new NICs and LUNs that have been added.

## 4.7 Setting Up Networks

This section discusses creation and use of Oracle VM Networks. You can define a name or an alias for each logical network. When you have created your networks, you connect the physical network ports on Oracle VM Servers to the logical networks. Before you define the logical networks in Enterprise Manager, you have to review your physical network configuration that you intend to use, such as VLAN and subnet usage. Setting up an Oracle VM network involves:

- Generating MAC Addresses
- Creating and configuring VLAN Groups
- Creating Networks

**Figure 4–9 Setting Up Networks**

### 4.7.1 Generating MAC Addresses

Virtual network interfaces (VNICs) are used to carry network traffic to virtual machines. You can generate MAC addresses that can be assigned to VNICs in a virtual machine. Specify a range of MAC addresses that need to be generated.

To create MAC Addresses, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Right click on an OVM Manager target from the left panel and select **Manage Network**.

**Figure 4–10 Virtual Network Interface Card Manager**

**Network**

Networks VLAN Groups **Virtual Network Interface Card Manager**

Virtual Machine networks carry network traffic to virtual network interface cards (VNICs) in the Virtual Machines. Each VNIC can be connected to one Virtual Machine network.

View   Search

| VNICs(MAC Address) | Ethernet Network | Virtual Mach (VNIC Owne |
|--------------------|------------------|-------------------------|
| 00:21:f6:12:00:00  |                  |                         |
| 00:21:f6:12:00:01  |                  |                         |
| 00:21:f6:12:00:02  |                  |                         |
| 00:21:f6:12:00:03  |                  |                         |
| 00:21:f6:12:00:04  |                  |                         |
| 00:21:f6:12:00:05  |                  |                         |
| 00:21:f6:12:00:06  |                  |                         |
| 00:21:f6:12:00:07  |                  |                         |
| 00:21:f6:12:00:08  |                  |                         |
| 00:21:f6:12:00:09  |                  |                         |
| 00:21:f6:12:00:0a  |                  |                         |
| 00:21:f6:12:00:0b  |                  |                         |
| 00:21:f6:12:00:0c  |                  |                         |
| 00:21:f6:12:00:0d  |                  |                         |
| 00:21:f6:12:00:0e  |                  |                         |
| 00:21:f6:12:00:0f  |                  |                         |
| 00:21:f6:12:00:10  |                  |                         |
| 00:21:f6:12:00:11  |                  |                         |
| 00:21:f6:12:00:12  |                  |                         |

Total Number of VNICs: 25 Used: 0 Available: 25

3. Click the **Virtual Network Interface Card Manager** tab.
4. Click **Generate**. In the Generate MAC Addresses pop-up window, you are prompted to specify an initial MAC address in 3 sets of 3 characters. This is a

two-digit number for the Initial MAC Address. Click **OK** to return to the previous page. You can now use these VNICs while creating virtual machines.

## 4.7.2 Configuring a VLAN Group

---

---

**Note:** Configuring a VLAN Group is an optional feature. If VLANs exist in your physical network, you must follow the steps in this section to use them in the cloud.

---

---

VLAN (Virtual Area Network) is a mechanism to virtualize network interfaces across subnets and physical location so that they appear to be on a single confined LAN. This concept is used to isolate (and group) the network traffic across these distributed interface, which enables secure network traffic within such nodes.

You can create multiple virtual LAN (VLANs) on the same NIC port. Each VLAN is an independent logical network operating with other VLANs over the same physical connection. Configuring VLANs involves creating one or more VLAN Groups, each of which contain multiple VLANs segments.

Each VLAN is assigned a distinct VLAN identification. The VLAN ID is used by an attached VLAN switch to segregate traffic among the different VLANs operating on the same link. Once a VLAN is configured, it functions exactly like a separate physical connection. VLANs need to be configured in the physical switches before you use them. To create a VLAN group, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Right click on an OVM Manager target from the left panel and select **Manage Network**.
3. Click the **VLAN Groups** tab and click **Create**.
4. Enter a name and description for the VLAN Group.
5. Enter one or more VLAN Segments. This is the VLAN ID that is enabled on the switch ports connecting the virtual server to the switch. The valid range is from 1 to 4094.
6. Click **Add** to select the ports to be added to the network. Select one or more virtual servers to be added to the network. If you add more than two ports of a virtual server, the first four ports will be bonded. Network bonding refers to the combination of network interfaces on one host for redundancy or increased throughput.

If your network supports jumbo frames, you must set the Maximum Transfer Unit size in the MTU field. Setting the MTU field, sets the maximum transmission rate, so that larger packets can be sent and received. You can set the MTU to a maximum of 1500 for Ethernet, 9000 for 1GbE, and 64,000 for 10GbE.

---

---

**Note:** The ports that you have selected must have the VLAN Segment IDs enabled.

---

---

7. Click **OK** to create the VLAN Group.

### 4.7.3 Creating Networks

A network is a grouping of Oracle VM virtual server physical NIC ports. This section describes the procedure to create a logical network.

#### Prerequisites

- If you choose to associate a network profile with the network, then the network profile must be created before the network is created. A network profile stores the network interface configuration (IP address, subnet mask, host name, DNS servers) that is allocated to network interfaces of newly provisioned virtual machines.

Follow these steps to create a logical network:

1. Log into Enterprise Manager Cloud Control.
2. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
3. Right click on an OVM Manager target from the left panel and select **Manage Network**.
4. Click the **Create** option in the Networks tab. You are prompted to configure the ports of the VLAN Groups. Click **Create VLAN Group** to configure VLAN Groups. See [Section 4.7.2, "Configuring a VLAN Group"](#) for details. If you want to create the networks without configuring the VLAN groups, click **Continue Network Creation**. In the Create Network: General page, enter the name and description of the network.

**Figure 4–11 Create Network: General**

**OVM:Create Network**

General Configure Ports and Vlan Interfaces Network Profile and QoS Review

Create Network : General

Back Step 1 of 4 Next Finish Cancel

Name: test\_network

Description:

Network Type: ☒ Inter-server ☐ Intra-server

Virtual Server:

**Network Roles**

| Name              | Description                                                        | Select                              |
|-------------------|--------------------------------------------------------------------|-------------------------------------|
| Server Management | Used to communicate between the Manager and Server Pool Master.    | <input type="checkbox"/>            |
| Cluster HeartBeat | This network role is used for cluster heartbeat traffic.           | <input type="checkbox"/>            |
| Storage           | This network role is used to carry storage traffic.                | <input type="checkbox"/>            |
| Live Migrate      | This network role is used for virtual machine live migration data. | <input checked="" type="checkbox"/> |
| Virtual Machine   | Used to carry network traffic to Virtual Machines.                 | <input checked="" type="checkbox"/> |

**Concepts**

- Inter-server: Accessible outside of the Server, i.e., routable through standard switches.
- Intra-server: Only available within a single Server without a route to an external physical network.

5. Specify if you are creating an inter-server or intra-server network. An inter-server network is routable through standard switches. An intra-server network is available within a single server without a route to an external physical network.
6. Specify the network role by checking one or more check boxes in the Network Roles section.

---

**Note:** To deploy a Guest VM, you must have at least one logical network with the **Virtual Machine** role.

---

- **Server Management:** Manages the physical virtual servers in a server pool, for example, to update the Oracle VM Agent on the different virtual servers.

- **Live Migrate:** Carries live migrate traffic. You can migrate virtual machines from one virtual server to another in a server pool, without changing the status of the virtual machine.
  - **Cluster Heartbeat:** Carries cluster heartbeat traffic.
  - **Virtual Machine:** Used for the network traffic between the different virtual machines in a server pool. The virtual machine role can either be an Inter-Server (routable through standard switches), or an Intra-Server (without a route to an external physical network). The Virtual Machine role must be unique on the network, which means that this network cannot have other network roles. It is however possible, and very likely, to have multiple networks with the Virtual Machine role in one Oracle VM Manager.
  - **Storage:** Used for all storage transport in a server pool. It is used by the virtual servers to connect to ethernet-based storage repositories.
7. Click **Next**. In the Configure Ports and VLAN Interfaces page, click **Add** and select a VLAN Group from the list. For more details, see [Section 4.7.2, "Configuring a VLAN Group"](#).

**Figure 4–12 Create Network: Configure Ports and VLAN Interfaces**

OVM: Create Network

General **Configure Ports and VLAN Interfaces** Network Profile and QoS Review

Create Network : Configure Ports and VLAN Interfaces

Back Step 2 of 4 Next Finish Cancel

**VLAN Interfaces**

View Add... Remove

| Port              | MAC Address                                               | MTU  | Address Type | IP Address | Netmask | Bonding Mode |
|-------------------|-----------------------------------------------------------|------|--------------|------------|---------|--------------|
| ▼ vlan_120102_100 | network.EthernetPort (5) in zowiu.us.or 00:16:3e:3a:65:8a | 1500 | None         |            |         |              |

**Ports**

View Add... Remove

| Port                        | MAC Address       | MTU  | Address Type | IP Address | Netmask | Bonding Mode |
|-----------------------------|-------------------|------|--------------|------------|---------|--------------|
| network.EthernetPort (2) in | 00:16:3e:3a:65:87 | 9000 | None         |            |         |              |

8. After adding the VLAN Segment, click **Add** the VLAN Interfaces and the ports. Select the ports to be added to network. If you select more than two ports, they will be bonded or aggregated into a single interface. By default, the Bonding Mode is Active Passive but you can change it by selecting one of the following:
- **Active Passive:** There is one active NIC. If this one goes down, another NIC becomes active.
  - **Link Aggregation:** All NICs act as one NIC which results in a higher throughput.
  - **Load Balanced:** The network traffic is equally balanced over the NICs of the machine.

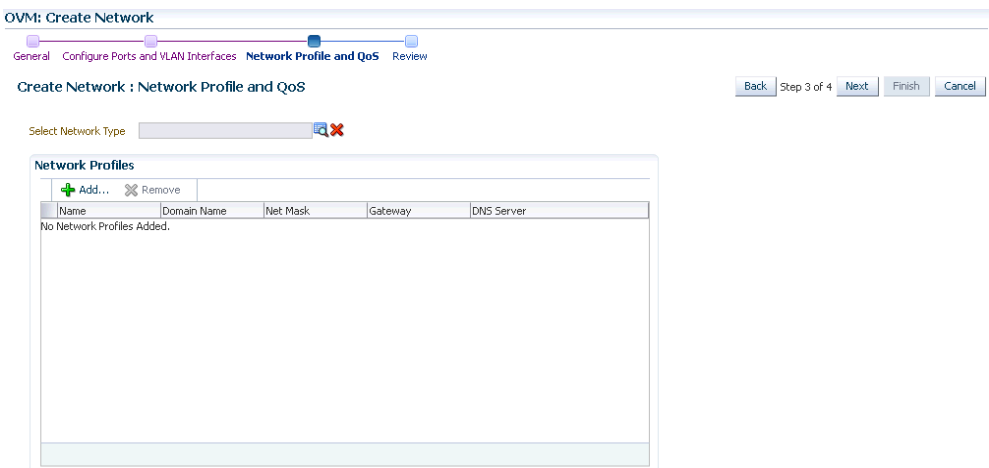
If your network supports jumbo frames, you must set the Maximum Transfer Unit size in the MTU field. Setting the MTU field, sets the maximum transmission rate, so that larger packets can be sent and received. You can set the MTU to a maximum of 1500 for Ethernet, 9000 for 1GbE, and 64,000 for 10GbE

9. Click **Next**. The Network Profile and QoS page appears.



**Note:** Defining the Network Profile and QoS is optional. A network profile stores the network interface configuration (IP address, subnet mask, host name, DNS servers) to be allocated to network interfaces of newly provisioned virtual machines.

Figure 4–13 Create Network: Network Profile and QoS



Specify the following:


- **Network Type:** You can select Internet Routable, Non Internet Routable, or a predefined set of network types. For more details, see [Section 4.16.1, "Setting Up the Network Type"](#).
- **Network Profile:** A network profile defines a set of IP Addresses, their associated host names (optionally), and common networking attributes for them. See [Section 4.8, "Creating a Network Profile"](#) for details.

10. Review the details entered so far and click **Finish** to create the network.

## 4.8 Creating a Network Profile

A network profile is used to automate assignment of IP addresses to guest virtual machines. A network profile is a list of IP address along with host names. It defines a set of IP addresses, their associated host-names, and common networking attributes for them.

Figure 4–14 Network Profile



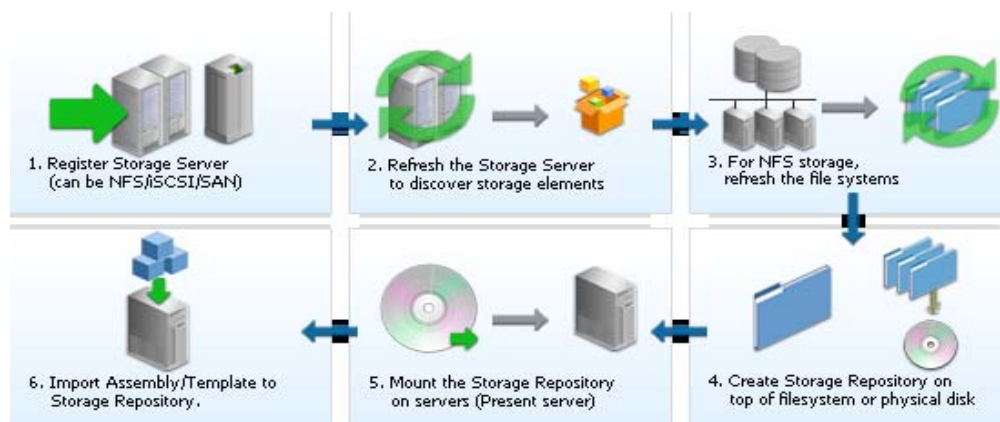
To create a network profile, follow these steps:

1. From the **Setup** menu, select **Provisioning and Patching**, then select **Network Profile**.
2. Click **Create** on the Network Profile page.
3. Enter the following details in the Create Network Profile pop-up window:
  - **Name:** Unique name that identifies the network profile name.
  - **Domain Name:** Domain for which the profile is being defined. The domain name must not begin with a dot ('.'). For example, you must specify the domain name as `us.oracle.com` and not with a dot ('.') in the front like `.us.oracle.com`.
  - **Netmask:** Network mask for the subnet IP. For example: `255.255.240.0`
  - **Gateway:** Gateway for the IP address list.
  - **DNS Servers:** Name resolution servers.
4. You can add a list of IP addresses or specify a range. If you select:
  - **List:** Specify IP addresses with host name and MAC address (optional).
  - **Range:** Specify a Hostname Pattern. The Start Value is appended to the host names generated. The First IP Address and the Last IP Address defines the range. For example: If the pattern is say `hostprod` and Start Value is 1, First IP Address is `10.1.1.1` and Last IP Address is `10.1.1.3`, the host names will be `hostprod1`, `hostprod2` and `hostprod3`.
5. Click **OK** to save the network profile.

## 4.9 Registering Storage Servers

You can create external storage elements like storage servers, file systems, and LUNs. An external storage element is created on dedicated storage hardware such as a server configured for NAS offering NFS shares. The server on which the storage element is present must be accessible by the Oracle VM Servers through a fibre channel or ethernet network. You can configure a virtual server pool with multiple such external storage devices.

**Figure 4–15 Creating Storage Servers and File Systems**

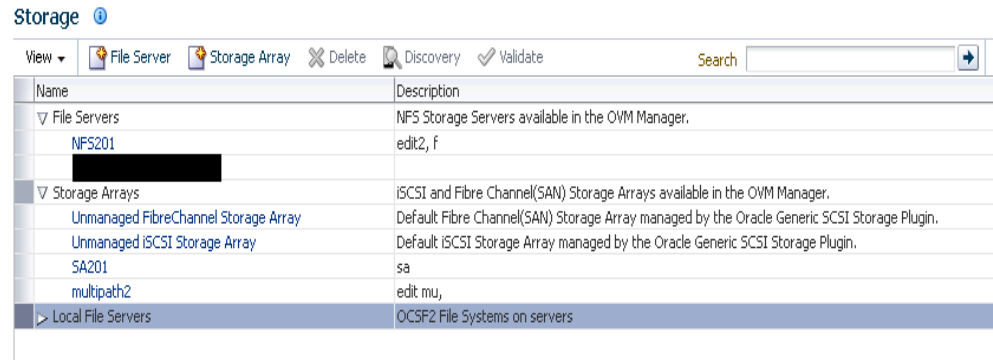


To set up a storage server, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.

2. Right click on an OVM Manager target in the left panel and select **Manage Storage**. Figure 4-16 appears.

**Figure 4-16 Storage Server**



3. A list of file servers, storage arrays, and local file servers that have been defined appears.
  - **Type:** Select the type of storage array that can be defined. A storage array is additional network provided storage that can be specifically used to store virtual disks for virtual machines. A server pool can be configured with multiple external storage devices. You can define 3 types of storage servers:
    - **iSCSI:** An iSCSI target configured or discovered on the Oracle VM Server.
    - **SAN:** A Fiber Channel SAN target configured/discovered on the Oracle VM Server.
    - **NFS:** An NFS or NAS file system mounted locally.
 Select File Server (NFS) or Storage Array (iSCSI or SAN).
4. You can register a File Server or a Storage Array by clicking the appropriate link.
5. Select a storage element and click **Discover** to start the discovery operation of the file systems available on the server. A confirmation message prompting you to refresh the file system appears. Click **OK** to refresh the file system and return to the Storage page.
6. Click on a **File Server**, **Storage Array**, or **Local File Server** link to drill down to the Details page.
7. Select a storage element and click **Validate**. A confirmation message prompting you to validate the storage server appears. Click **OK** to refresh the file system and return to the Storage page.

### 4.9.1 Registering the File Server

To register a file server, follow these steps:

1. Click the **File Server** link on the Storage page. The Register File Server: General page appears.

**Figure 4–17 Register File Server: General**

The screenshot shows the 'Register File Server: General' configuration page. At the top, there are tabs for 'General' and 'Admin and Refresh Servers', with 'General' being the active tab. Below the tabs, the title 'Register File Server : General' is displayed. On the right side, there are navigation buttons: 'Back', 'Step 1 of 2', 'Next', and 'Cancel'. The main form area contains several fields: 'Name' (with a plus icon), 'Description' (a text area), 'Plugin Name' (a dropdown menu showing 'Oracle Generic Network File Sys'), 'Admin Host' (a text field), 'Admin Username' (a text field), 'Admin Password' (a text field), 'Access Host' (a text field), and 'Uniform Exports' (a checked checkbox). The form is organized into sections: 'Administration Information' for the first four fields and 'Access Information' for the last two.

2. Enter the following details:
  - **Name:** Name by which the storage server is to be identified.
  - **Plugin Name:** The Oracle VM Manager communicates with the external storage through a set of plug-ins. Generic plug-ins offer a limited set of standard storage operations on virtually all storage hardware, such as discovering and operating on existing storage resources. Vendor-specific plug-ins include a much larger set of operations (snapshot, clone, and so on).
  - **Admin Host:** Host name of the Administration Server.
  - **Admin Username and Password:** User name and password for the Administration Server.
  - **Access Host:** Host name or IP address of the server offering external storage.
  - **Uniform Exports:** Select this check box to indicate that the NFS exports are uniform. This ensures that all the VM Servers in the server pool will have access to the same exports on the NFS server. If this check box is not selected, the NFS exports can be accessed only by specific servers that are used to refresh the NFS file systems.
3. Click **Next**. The Register File Server: Admin and Refresh Servers page appears. Enter the following:
  - **Admin Servers:** VM Servers that have the required Storage Plugin installed and are responsible for information exchange with the Admin Host of the Storage Server are referred to as Admin Servers. Click Add to add one or more Admin Servers.
  - **Refresh Servers:** If you have not selected the Uniform Exports checkbox in the Register File Server: General page, click Add to specify the list of Refresh servers that will be used to refresh the NFS file systems.
4. Click **Finish** to add the file server. The new server appears in the table. At this point, the Storage Connect plugin has established a link to the storage location. No storage has been mounted yet.
5. Click **Discover** to start the discovery operation of the file systems available on the server. When the operation is complete, the available file systems appear in the File Systems tab in the Storage Details page.

## 4.9.2 Registering the Storage Array

To register a storage array, follow these steps:

1. Click the **Storage Array** link on the Storage page. The Register Storage Array: General page appears.

**Figure 4–18 Register Storage Array: General**

2. Enter the following details:
  - **Name:** Name by which the storage server is to be identified.
  - **Storage Type:** Select the type of storage array that can be defined. A storage array is additional network provided storage that can be specifically used to store virtual disks for virtual machines. You can define 2 types of storage arrays:
    - **Fibre Channel Storage Array:** A Fiber Channel SAN target configured or discovered on the Oracle VM Server.
    - **iSCSI Storage Array:** An iSCSI target configured or discovered on the Oracle VM Server.
  - **Plugin Name:** The Oracle VM Manager communicates with the external storage through a set of plug-ins. Generic plug-ins offer a limited set of standard storage operations on virtually all storage hardware, such as discovering and operating on existing storage resources. Vendor-specific plug-ins include a much larger set of operations (snapshot, clone, and so on).
  - **Plugin Private Data:** Any vendor-specific storage plug-in data that may be required. This is not available to generic storage.

---

**Note:** When registering a vendor-specific storage array, be sure to double-check the information you entered in the plug-in private data field. Once the storage array is registered, this field can no longer be modified. If you need to update the plug-in private data, you must unregister and re-register the storage array.

---

- **Admin Host:** Host name of the Administration Server.
  - **Admin Username and Password:** User name and password for the Administration Server.
3. Click **Next**. The Register Storage Array: Access Information page appears.

**Figure 4–19 Register Storage Array: Access Information**

General Access Information Admin Servers

Register Storage Array : Access Information

Back Step 2 of 3 Next Cancel

Access Host

\* Access Host

\* Access Port 3260

Access Username

Access Password

Use CHAP ☐

4. Enter the following details:

- **Access Host:** Host name or IP address of the server offering external storage.
- **Access Port:** Port number used to access the storage.
- **Access Username and Password:** User name and password for the Access Host.
- **Use Chap (Applies to all Access Hosts):** Select this check box to enable Chap authentication.

---

**Note:** If you have selected the Storage Type as iSCSI Storage Array and Plugin Name as Oracle Generic SCSI Plugin (1.2.1 or above) in the Register Storage Array: General page, you can add multiple Access Hosts as shown in [Figure 4–20](#).

---

**Figure 4–20 Register Storage Array (Multiple)**

OVM: Register Storage Array

General Access Information Admin Servers

Register Storage Array : Access Information

Back Step 2 of 3 Next Cancel

Access Host

View + Add Edit... Delete...

| Access Host | Access Port | Access Username |
|-------------|-------------|-----------------|
|             | 3260        |                 |

Use Chap (Applies to all Access Hosts) ☐

5. Click **Next**. The Register Storage Array: Admin Servers page appears. Click **Add** to add one or more Administration Servers on which the storage plug-in has been installed. The Administration Server is responsible for information exchange with the Administration Host of the storage server.
6. Click **Finish** to add the storage array. The new server appears in the table. At this point, the Storage Connect plugin has established a link to the storage location. No storage has been mounted yet.
7. Click **Discover** to start the discovery operation of the file systems available on the server. When the operation is complete, the available file systems appear in the General tab of the Storage Details page.

### 4.9.3 File Server Details

Click a File Server link to view the Details page. This page contains the following tabs:

- **General:** A list of Administration Servers is displayed. An Administration Server is a virtual server on which the storage plug-in has been installed. You can add an Administration Server or delete an existing one.

In the Configuration region, the storage server details are displayed. Click **Edit** to modify details such as Name, Description, and the name of the Access Host which provides the external storage.

In the Storage QOS Mapping region, map the storage type defined by the Cloud Administrator to the QOS defined by the Storage Server.

- **Refresh Servers:** For non-uniform NFS file system exports, you can specify the servers that need to be refreshed. Click **Add**. The Refresh Servers dialog box is displayed. Click Select Servers, choose one or more servers from the list and click Select. Click **OK** to refresh the servers.
- **File Systems:** A list of file systems is displayed. You can do the following:
  - Select a file system from the list and click **Edit**. You can modify the name and description of the file system.
  - Select a file system and click **Delete**. A confirmation message is displayed. Click **OK** to delete the file system.
  - Select one or more file systems and click **Refresh**. A confirmation message is displayed. Click **OK** to refresh the file system.
- **Access Groups:** Shows the mapping between Refresh Servers and File Systems (Exports). You can Create, Edit, or Delete an Access Group.

### 4.9.4 Storage Array Details

Click a Storage Array link to view the Details page. This page has four tabs:

- **General:** Provides general information on the storage array. A list of administration servers is displayed in the Administration region. An Administration Server is a virtual server on which the storage plug-in has been installed. You can add an Administration Server or delete an existing one.

In the Configuration region, the storage server details are displayed. Click Edit to modify the Name and Description.

- **Physical Disks:** Displays a list of all the physical disks on the storage array. Depending on the type of storage connector (You can Create, Clone, Edit, Resize (resize or extend the physical disk), Rescan (refreshes the physical disk), Delete, and Create File System (create OCFS2 file system on the physical disk).

---

**Note:** The operation allowed depends on the type of storage connect plug-in that has been installed. For example, if the generic storage connect plug-in has been installed, you cannot perform a Create operation on the physical disks.

---

- **Access Group:** Shows the mapping between Storage Initiators and Physical Disks. You can Create, Edit, Present / Unpresent Access Group (Add or Delete physical disk), and Delete access group.
- **Volume Group:** Displays the volume group available for this storage server.

## 4.9.5 Local File Server Details

Click on a Local File Server Servers link to view the Details page. This page contains the following tabs:

- **General:** The configuration details of the selected Local File Server including the Administration Host, Plug-in Name, and so on are displayed.
- **File Systems:** You can edit or refresh a file system in the list. Select a file system from the list and click **Edit**. You can modify the name and description of the file system.

To refresh, select a file system from the list and click **Refresh**. Select the name of the virtual server and click **OK** to refresh the file system.

## 4.10 Creating a Storage Repository

A storage repository is virtual disk space on top of physical storage hardware, made available to the Oracle VM Servers in a server pool or various server pools. It defines where Oracle VM resources may reside. Resources include virtual machines, templates for virtual machine creation, virtual machine assemblies, ISO images, shared virtual disks, and so on.

After you have created the storage server, you can assign these storage resources to the server pools in a zone by creating a storage repository.

To create a storage repository, follow these steps:

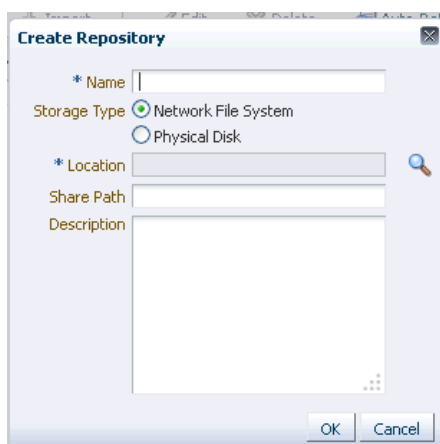
1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Right click on an OVM Manager target in the left panel and select **Manage Storage Repository**.

**Figure 4–21 Storage Repository**

| Name                | Owned | File System | File System Size | Size Used (%) | Disk | Share Path | Storage Server | Description |
|---------------------|-------|-------------|------------------|---------------|------|------------|----------------|-------------|
| leichRepos3         | Yes   |             | 5,609,545        | 20            |      |            | leichFS        |             |
| thin_supported_Rep3 | Yes   |             | 700              | 0             |      |            | Local FS       |             |
| thin_supported_Rep3 | Yes   |             | 700              | 0             |      |            | Local FS       |             |

3. The list of storage repositories that have been defined appears. Click **Create** to create a storage repository.
4. In the pop-up window, enter the name of the storage repository.



**Figure 4–22 Create Storage Repository**

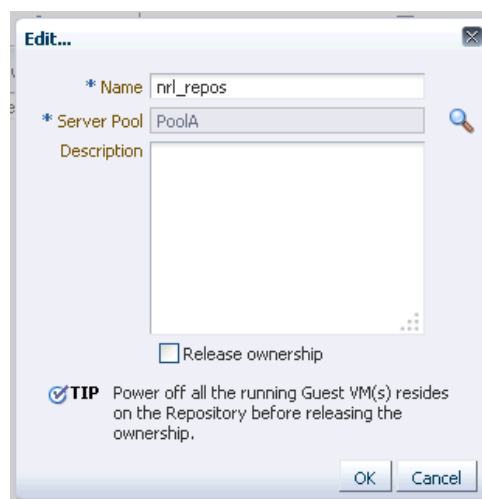
5. Click the **Search** icon in the Disk field to select the Storage Server and File System on which the storage repository is to be created and click **OK**.

If you select the NFS File System, you must select the location of the storage server and a virtual server. You can optionally specify the Share Path and a description. The storage repository has now been created and will be displayed in the Storage Repositories page.

#### 4.10.1 Performing Administrative Operations on a Storage Repository

After the storage repository has been created, you can perform a number of management operations on it. The listing of storage repositories depends on the server pool selection. If you select server pools, you see all the storage repositories. If you select an individual server pool, only the storage repositories related to that server pool appear. You can perform the following operations:

- **Edit:** Select a repository and click **Edit**. You can modify the name and description or release the ownership of the repository. Click **OK** to save the changes.

**Figure 4–23 Edit Storage Repository**

---

**Note:** If you are editing a storage repository that has been created on a physical disk, you can change the server pool with which the storage repository is associated. To change the server pool, in the Edit dialog box, click the **Search** icon next to the Server Pool field, select a different server pool and click **OK**.

---

- **Delete:** Select a repository and click **Delete**. All the contents of the selected repository along with its associations will be deleted.
- **Auto-Refresh:** Select this option to enable Auto-Refresh and periodically refresh and update the file system size of all storage repositories. In the dialog box that appears, select the **Enable** check box and specify the Interval at which the file system size is to be refreshed.

---

**Note:** As the Auto-Refresh option locks the repositories, you must use this option when no deployment operations are scheduled.

---


## 4.11 Presenting the Storage Repository

After you have created the storage repository, you can propagate it to one or more server pools. When the storage repository is prepared and created, it still must be made available for use by your virtual servers before it can be used. Typically you present the storage repository to all the virtual servers in the server pool.

To present a storage repository to server pool, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Right click on an OVM Manager target from the left panel and select **Manage Storage Repository**.
3. In the Storage Repository page, select a storage repository from the list. The list of server pools on which the storage repositories have been propagated appears.

**Figure 4–24 Presented Servers**



4. Select one or more server pools from the list and click **Present**. In the Present Servers dialog box, click the **Select Servers** icon. The list of members appears in the Select Targets dialog box. Expand the Members in this dialog box and check the boxes for zones, pools, or individual virtual servers and click **Select** to return to the previous dialog box. Click **Present** to mount the member on the server.

## 4.12 Importing Assemblies, Templates, and Virtual Disks into the Storage Repository


A storage repository associated with an Oracle VM server stores binary components associated with an Oracle VM servers. The following software components can be stored in the storage repository:

- Oracle VM Templates
- Oracle Virtual Assemblies
- ISOs
- Virtual Disks
- VM Files

The Oracle Virtual Assemblies, Oracle VM Templates and ISO images need to be imported into the storage repository before guest virtual machines can be created using them. Follow these steps to import software components into the storage repository:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Right click on an OVM Manager target in the left panel and select **Manage Storage Repository**.
3. Click on a Repository Name link. The Storage Repository Details page shows the list of server pools on which the selected storage repository has been propagated.
4. Click the **Assemblies** tab. An assembly can be imported as a single .ova file.

**Figure 4–25 Storage Repository - Assemblies**



- Click **Import**. You can import the assembly from the Software Library or from an external source. If you select:
  - **Software Library**: The list of assemblies present in the Software Library are displayed. Select the assemblies to be imported and click **OK**. After the assembly has been imported, a series of templates is created. This option unpacks the contents of the assembly, for example, configuration files, disk images, and so on.
  - **External**: Specify a external URL from which the assembly is to be imported and click **OK**. The URL you specify must be accessible from the OVS Host on which the repository is presented.
- 5. Click on the **ISOs** tab to import the ISO images. Virtual machines have no access to the physical DVD or CD-ROM drive. You can assign virtual drives to virtual machines by offering .iso files containing the image of a DVD or CD-ROM. These image files can be found in the ISOs tab of the storage repository.

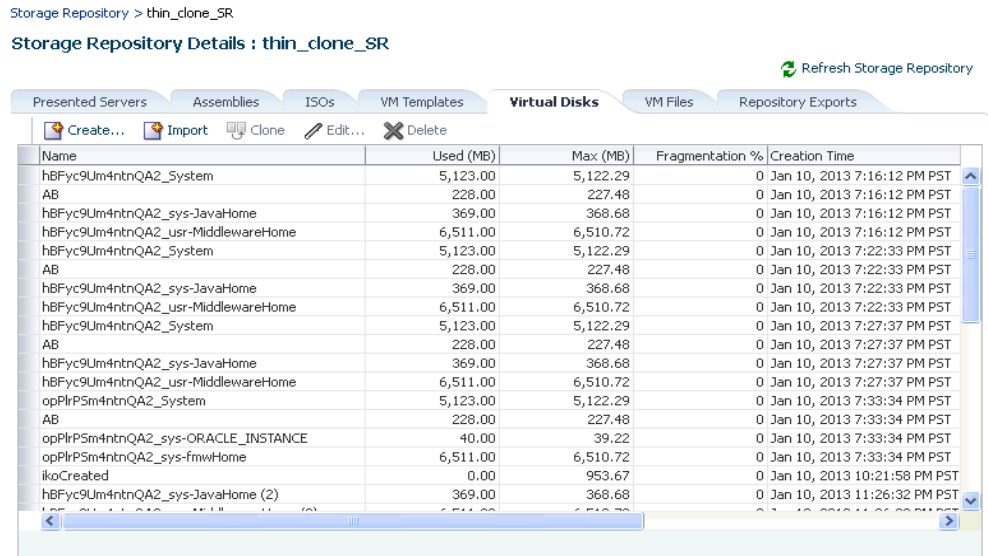
- Click **Import**. You can import the ISO image from the Software Library or from an external source. If you select:
  - **Software Library**: The list of ISO images present in the Software Library are displayed. Select the ISO images to be imported and click **OK**. Select the ISO image that has been imported and click **Assign** to assign this template to a server pool.
  - **External**: Specify a valid external URL from which the ISO image is to be imported and click **OK**. The URL you specify must be accessible from the OVS Host on which the repository is presented.
- 6. Click the **VM Templates** tab. Templates contain the configuration of a single virtual machine. When you create a new virtual machine, you choose to either create it from scratch or from a template. In a storage repository, templates are typically imported as an archive (.tgz, .tar or other). The archive contains a .cfg file with the virtual machine configuration, and at least one .img file which is a virtual disk image.

**Figure 4–26 Storage Repository - VM Templates**

Storage Repository Details : thin\_clone\_SR Refresh Storage Repository

| Presented Servers   Assemblies   ISOs <b>VM Templates</b> Virtual Disks   VM Files   Repository Exports |           |                             |                                 |             |  |
|---------------------------------------------------------------------------------------------------------|-----------|-----------------------------|---------------------------------|-------------|--|
| <span>Import</span> <span>Clone</span> <span>Edit...</span> <span>Delete</span>                         |           |                             |                                 |             |  |
| Name                                                                                                    | Size (MB) | Creation Time               | Parent Assembly                 | Description |  |
| fmw_V5_URN_oracle:defaultService:em:provisioning:1:ci                                                   | 12,231.00 | Jan 10, 2013 7:12:14 PM PST | fmw_V5_URN_oracle:Type: coherer |             |  |
| fmw_V5_URN_oracle:defaultService:em:provisioning:1:ci                                                   | 12,231.00 | Jan 10, 2013 7:16:16 PM PST | fmw_V5_URN_oracle:Type: wls Cat |             |  |
| fmw_V5_URN_oracle:defaultService:em:provisioning:1:ci                                                   | 12,231.00 | Jan 10, 2013 7:22:37 PM PST | fmw_V5_URN_oracle:Type: wls Cat |             |  |
| fmw_V5_URN_oracle:defaultService:em:provisioning:1:ci                                                   | 11,902.00 | Jan 10, 2013 7:27:42 PM PST | fmw_V5_URN_oracle:Type: ohs Cat |             |  |

- Click **Import**. You can import the VM Template from the Software Library or from an external source. If you select:
  - **Software Library**: The list of VM Templates present in the Software Library are displayed. Select the VM Template to be imported and click **OK**. After the template has been imported, select the template and click **Assign** to assign this template to a server pool.
  - **External**: Specify a valid external URL from which the VM template is to be imported and click **OK**. The URL you specify must be accessible from the OVS Host on which the repository is presented.
- 7. Click on the **Virtual Disks** tab.

**Figure 4–27 Storage Repository - Virtual Disks**

- Click **Create** and size of virtual disk to be created.
- Indicate if this is a shared disk by checking the **Shared** check box.
- Click **Import**. You can import the virtual disk from the Software Library or an external source. If you select:
  - **Software Library**: The list of virtual disks present in the Software Library are displayed. Select the virtual disks to be imported and click **OK**. Select the virtual disk has been imported, click **Assign** to assign this template to a server pool.
  - **External**: Specify a valid external URL from which the virtual disk is to be imported and click **OK**. The URL you specify must be accessible from the OVS Host on which the repository is presented. If you reference the virtual disk using `http://` or `ftp://` protocols, you must include `user/pw` in the reference.

For example, you must specify the URL as

`http://user:pw@host/virtualdisk_file_image_file` or

`ftp://user:pw@host/virtual_disk_image_file`. If the virtual disk is

accessible by way of NFS, you can also specify

`file:///net/host/virtual_disk_image_file` if the OVS is enabled for NFS automount.

8. Click **Return** to return to the Storage Repository page.

#### 4.12.1 Creating a Repository Export (Optional)

You can configure the virtual server to enable third party applications to back up the contents of the storage repository.

##### Prerequisites


- The Repository Exports feature is available only for storage repositories based on physical disks and at least one server must have been presented to the repository.

- The virtual server must be configured to provide an NFS share to access the contents of the repository.
- The virtual server must be in a clustered server pool and have the OCFS2-based storage repository presented to it.

To create a repository export, follow these steps:

1. Click the **Repository Exports** tab in the Storage Repository Details page.
2. Click **Create**. The Create Repository Export window appears.

**Figure 4–28 Storage Repository - Repository Exports**



3. Enter a name of the host which is used to mount the contents of the storage repository to the Export Server.
4. Specify the options to restrict access to the repository. For Linux systems, the default is `rw, async, no_root_squash`.
5. Click **Add** in the Export Server(s) region. Select one or more virtual servers on which the storage repository is to be presented and click **Select**.
6. Click **OK** to submit the job and return to the Storage Repository Details page.

## 4.13 Using Self Update to Download Oracle VM Templates and Virtual Assemblies

You can download the preconfigured assemblies and templates from the Self Update Console by following these steps:

1. Log in to Enterprise Manager as a user with the `EM_CLOUD_ADMINISTRATOR` role.
2. From the **Setup** menu, select **Extensibility** and then select **Self Update**. The Self Update Console with the list of entity types that can be downloaded appears.

**Figure 4–29 Self Update Console**

**Self Update** Page Refreshed Jan 25, 2012 11:44:23 AM PST

Oracle periodically provides new functionality and updates for existing features in Enterprise Manager. The Self Update home allows administrators to receive notifications and view, download, and apply such updates. While these updates are retrieved automatically, a manual check can be made at any time.

**Status** Informational Updates

Connection Mode: **Online** Last Download Time: Jan 15, 2012 9:30:53 PM PST Last Apply Time: Jan 15  
 Most Recent Refresh Time: ✔ Jan 25, 2012 10:25:31 AM PST Last Download Type: Plug-in Last Apply Type: Plug-in

Actions: Open Check Updates Agent Software

| Type                                              | Available Updates | Downloaded Updates | Applied Updates | Description                                                                                                                                    |
|---------------------------------------------------|-------------------|--------------------|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------|
| Agent Software                                    | 5                 | 0                  | 2               | Agent software has to be installed on hosts for managing the host.                                                                             |
| Compliance Content                                | 0                 | 0                  | 0               | Compliance Content contains Framework, Standard, Rules with support of add and delete on these entities.                                       |
| Diagnostics Checks                                | 0                 | 0                  | 0               | Target side policy checks that identify conditions that may require the attention of target administrators.                                    |
| EM Deployment Prerequisite Resources              | 0                 | 0                  | 0               | EM Deployment Pre-requisite Checks are the metadata used for checking prerequisites for Install, Upgrade, Patching of EM Platform and Plugins. |
| Exadata Configuration Template                    | 0                 | 0                  | 0               | Configuration Templates for Oracle Exadata Database Machines                                                                                   |
| Management Connector                              | 1                 | 0                  | 1               | Management Connectors are components that integrate different enterprise frameworks into the Enterprise Manager Console                        |
| Middleware Profiles and Gold Images               | 0                 | 0                  | 0               | A collection of Software Components used for provisioning of Oracle Application Server homes.                                                  |
| Oracle Database Provisioning Profiles             | 0                 | 0                  | 0               | A collection of Software Components used for provisioning of Oracle Database, Clusterware and Grid Infrastructure homes.                       |
| Oracle VM Templates and Oracle Virtual Assemblies | 0                 | 0                  | 0               | Preinstalled and preconfigured software packages to deploy Oracle Guest VM                                                                     |
|                                                   |                   |                    |                 | Plug-in extends Enterprise Manager to manage newer target type as well as to bring                                                             |

**Past Activities: Agent Software**

| Actions | Status    | OS Platform  | Version    | Administrator | Start Time | Elapsed Time(Sec) |
|---------|-----------|--------------|------------|---------------|------------|-------------------|
| Apply   | Succeeded | Linux x86-64 | 12.1.0.1.0 | SYSMAN        |            | 0.23              |

3. Select the Oracle VM Templates and Oracle Virtual Assemblies folder and click **Open**. The Oracle VM Templates and Oracle Virtual Assemblies Updates page appears.
4. Select an update and click **Download**. The Schedule Download dialog appears.
5. Specify if the update is to be downloaded immediately or at a later date.
6. Click **Select**. An Enterprise Manager job is created to download the update to the Software Library. Enterprise Manager starts downloading the archive from the Oracle Enterprise Manager store. Wait for the download to complete. (When in offline mode the system starts reading from the specified location.) When the download is complete, the status indicates that the update has been downloaded.
7. Click **Apply**. A job is submitted to apply the update. The downloaded plug-in is applied and appears on the Oracle VM Templates and Oracle Virtual Assemblies Updates page.

## 4.14 Creating a Virtual Server Pool

This section guides you through the ways of designing a server pool to meet your requirements.

Before creating a server pool, you need to consider how many physical servers will be contained in the server pool, and what functions each physical server will perform. The more virtual machines you will run in the server pool, the more resources these virtual machines will consume, therefore the more physical servers are needed to provide sufficient resources for the server pool.

A server pool is scalable. If you find a server pool does not have sufficient resources, such as CPU and memory, to run the virtual machines and the applications inside, you can expand the server pool by adding more Oracle virtual servers.

A virtual server pool contains one or more virtual servers and guest virtual machines. A virtual server can belong to one and only one virtual server pool at a time. Guest virtual machines and resources are also associated with the server pools.

### Prerequisites

- A non-clustered virtual server pool can contain a maximum of 64 Oracle VM servers and a clustered virtual server pool can contain a maximum of 32 Oracle VM servers.
- Virtual servers that can be deployed as the Master Server, the Utility Server, and/or the Guest Virtual Machine. At least one virtual server is required in a server pool.
- All virtual servers in a server pool should have CPUs in the same CPU family. If they are not in the same CPU family, some operations such as live migration may fail. Though the CPUs should be in the same CPU family, they may have differing configurations, such as different number of cores. Other hardware components on the host computer may also differ, such as the amount of RAM, the number and size of disk drives, and so on.

---

**Note:** Although the host computers may have differing configurations, Oracle recommends that all virtual servers in a server pool are identical.

---

- A dedicated file system (either a NAS export, or a LUN) to use as the server pool file system is available.
- IP addresses for the virtual servers must be available. The IP address to be used as the Master Server's virtual IP address must be an unused and available IP address.
- To create a virtual server pool, you must be logged in as an **Administrator** with **Add Any Target** privilege.
- To delete a virtual server pool, you must have **Full** privilege on the server pool and appropriate privileges on the virtual servers and guest virtual machines belonging to the virtual server pool.

To create a virtual server pool:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Right click on an OVM Manager target in the left panel and select **Create Virtual Server Pool**.



Figure 4–30 Create Virtual Server Pool

VM Manager

Page Refreshed Jan 11, 2013 8:17:09 PM PST

### Create Virtual Server Pool

OK Cancel

\* Virtual Server Pool Name

Description

Keymap en-us (English, United States)

VM Start Policy Start on Best Server

Secure VM Migrate ☐

Activate Cluster ☒

Cluster Timeout (sec) 120

Type of Pool File System ☒ Network File System ☐ Physical Disk

\* Location

☒ TIP File System must be refreshed before they can be used.

\* Virtual IP

☒ TIP Virtual IP should be an unused IP

#### Virtual Servers

+ Add... - Remove

| Name                    |
|-------------------------|
| No Virtual Server Added |

#### Concepts

- Secure VM Migrate: Select whether to enable encrypted migration of virtual machines. When Secure VM Migrate is checked, virtual machines are migrated using SSL to protect the data during the migration process.
- Activate Cluster: Select whether to enable clustering of the Oracle VM Servers in the server pool to enable HA.
- Cluster Timeout: When enable Activate Cluster, allow to set the timeout in seconds for cluster. Disk heartbeat and network heartbeat are derived from the cluster timeout value.
- Pool File System: The server pool file system is used to hold the server pool and cluster data, and is also used for cluster heartbeat. The size of Pool File System should be at least 12GB.
- For Clustered ServerPool, if you are using Network File System as Pool File System, Admin Server must be set for the Storage Server which exposes the Network File System. The Network File System should be refreshed before it can be used as Pool File System.
- The Pool File System(NFS or Physical Disk) should be accessible on all the virtual servers in the ServerPool.
- Virtual IP: An IP address used to identify the master Oracle VM Server, which controls the other Oracle VM Server in the server pool.

### 3. Enter the server pool information:

- **Virtual Server Pool Name:** The name of the server pool. A server pool name must consist of alphanumeric characters, and must not contain spaces or special characters, except the underscore (\_) or hyphen (-) characters. The maximum length of a server pool name is 200 characters.
- **Description:** A description of the server pool.
- **Keymap:** Specify the keyboard type for the guest virtual machines created in this server pool.
- **VM Start Policy:** Select the policy to be used to start the virtual machines. You can choose:
  - **Start on Best Server:** The best server in the server pool is used to start the virtual machine.
  - **Start on Current Server:** The virtual machine will be started on the same server on which it was created.
- **Secure VM Migrate:** Specify whether encrypted migration of virtual machines is to be enabled. If this checkbox is selected, virtual machines are migrated using SSL to protect data during the migration process.
- **Activate Cluster:** Select this check box to enable high availability on the server pool.
- **Cluster Timeout:** If the Activate Cluster is enabled, you can set the cluster timeout period in seconds. The disk heartbeat and network heartbeat are derived from the cluster timeout value.
- **Type of Pool File System:** Select the pool file system to use for the server pool. The pool file system is used to hold the server pool, cluster data, and is also used for the cluster heartbeat.

The type can either be Network File System or Physical Disk. If you are using a file server for the server pool, select Network File System. If you are using a physical disk on a storage array for the server pool, select the type as Physical Disk.

- **Location:** Click the **Search** icon to search for and select the pool file system. This must be accessible by all virtual servers in the server pool.
  - **Virtual IP:** An IP address used to locate the Master Server in the server pool. If the Master Server changes to another Oracle VM Server, this IP address is then assigned to the new host.
4. Click **Add** in the Virtual Servers section. Select one or more virtual servers from the Search window and click **OK**.

---

**Note:** Before you add a virtual server to the server pool, it must have been discovered. See [Section 4.5, "Discovering a Virtual Server"](#) for details.

---

5. Click **OK** to create a virtual server pool.

### 4.14.1 Editing a Virtual Server Pool

You can edit the configuration information of a virtual server pool, including the description, key map, start policy, and so on. You can also change the master server, which controls the cluster, as well as whether the virtual machines are migrated securely. You cannot change the Virtual IP or the Pool File System used for the server pool. To edit a virtual server pool, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. From the left panel, select the virtual server pool target that you want to edit.
3. From the **VM Server Pool** menu, select **Target Setup**, then select **Edit VM Server Pool**. The Edit Virtual Server Pool page appears.
4. You can also edit the following fields:
  - **Keymap:** The keyboard type for the guest virtual machines created in this server pool.
  - **VM Start Policy:** The policy to be used to start the virtual machines. You can choose to start on the current server or on the best server.
  - **Secure VM Migrate:** Indicates whether encrypted migration of virtual machines is to be enabled. If this checkbox is selected, virtual machines are migrated using SSL to protect data during the migration process.
5. In the Virtual Servers region, you can add or remove virtual servers.
6. Click **Change Agent Password** to change the password for the Management Agent running on the Master Server. The modified password will be applied to the Management Agents running on all virtual servers in the server pool.

---

**Note:** Before you add a virtual server to the server pool, it must have been discovered. See [Section 4.5, "Discovering a Virtual Server"](#) for details.

---

7. You can **Add**, **Edit**, or **Delete** Anti-Affinity Groups. Anti-affinity groups specify that a group of virtual machines cannot run on the same server. In the Add Anti-Affinity Group window, enter a name for the group and select one or more virtual machines that should be part of this group. The virtual machines that are part of this group cannot run on the same server.

8. Click **OK** to apply the changes to the virtual server pool.

## 4.15 Creating a Zone

A zone is used to group related cloud resources together. Cloud zones can be created based on location, software lifecycle status, for grouping resources according to a cost center or for metering and chargeback purposes. Typically, a zone can cover a department or possibly a small data center. The cloud environment can comprise of one or more zones. Each zone has a set of metrics that show the aggregate utilization of the zone.

### Prerequisites

- To create a zone, you must be logged in as an Administrator with Add Any Target privilege.
- To remove a zone, you must have Operator / Full privilege on the zone.

Follow these steps to create a zone:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Right click the OVM Manager target for which the zone is to be created in the left panel and select **Create Zone**.

**Figure 4–31 Create Zone**

3. In the Create Zone page, specify the zone name and description.
4. Indicate if this zone will be used by the self service users by checking the **Infrastructure Cloud Self Service Zone** check box.
5. Click **Add** in the Virtual Server Pool section to add one or more virtual server pools. See [Section 4.14, "Creating a Virtual Server Pool"](#) for details. You can add the virtual server pools for which you have the privileges and that are not part of any other zone.
6. Click **OK**. The virtual server pools and the storage arrays appears in the newly created zone on the Home page.

## 4.16 Setting Up the Storage Quality of Service (QoS)

Storage QoS refers to quality of service that is defined at the storage server level. Setting up the storage QoS involves defining various storage properties. Before


deploying a guest virtual machine, you can define the storage QoS, map this QoS to a storage server, and specify this QoS while deploying a guest virtual machine.

This involves defining the machine size, network type, and the storage QoS. After you have set up the storage QoS, you can set up the storage servers, network repositories, and storage arrays. This storage QoS can then be used during deployment of guest virtual machines.

To define the storage QoS, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Right click on the **Infrastructure Cloud** menu in the left panel and select **Setup**.

**Figure 4–32 Storage Quality of Service**




3. Click **Create** in the Storage Quality of Service page.
4. In the pop-up window, specify the name and description of the storage QoS and click **OK**. The newly added storage QOS appears on the Storage Quality of Service page.
5. Navigate to the OVM Manager Home page and select **Storage** from the **OVM Manager** menu.
6. Click the **Storage** link under File Servers. In the General tab, scroll down to the Storage QOS Mapping region.
7. Click **Edit**. The Storage QOS Mapping window appears. Click **Add** and select that you had created earlier. The Storage QOS you have added appears in the Storage QOS Mapping table.
8. To see how this mapping is used when you deploy an assembly, navigate to the Home page of a Zone under the OVM Manager. From the **Zone** menu, select **Deploy**, then select **Assembly**.
9. In the Deployment Configuration page, click **Override** in the Disk section and select the Storage QOS you have created. This storage QOS will be used for deploying the assembly.

### 4.16.1 Setting Up the Network Type

You can define the QoS for network by defining some network types like routable, internet-facing, and so on. After the network QoS has been defined, you can map this to any network that is being created. You can then use this network type while deploying a guest virtual machine.

To define the network QoS, follow these steps:

1. From the **Enterprise** menu, select **Infrastructure Cloud**, then select **Home**.
2. Select **Setup** from the **Infrastructure Cloud** menu.
3. Click the **Network Types** link on the Infrastructure Cloud page.

**Figure 4–33 Network Types**

4. Click **Create** in the Network Types page.
5. Enter a name and description of the network type and click **OK**.

### 4.16.2 Setting Up the Machine Sizes

You can define the machine size that will be available to the SSA users and during deployment. To define the machine size, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Right click on the **Infrastructure Cloud** menu in the left panel and select **Setup**.
3. Click the **Machine Sizes** link on the Infrastructure Cloud page.
4. Click **Add Machine Size**.
5. Enter the name (internal name of the machine size), description, number of VCPUs, amount of memory and local storage assigned to machines of this size. Click **Save** to add this machine size.

### 4.16.3 Configuring the Software Library User

While setting up the storage repository, you can import VM Templates, Assemblies, and other software components into the storage repository. To configure the software library user who has the privileges to perform these tasks, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Right click on the **Infrastructure Cloud** menu and select **Setup**.
3. Click the **Software Library User Configuration** link on the Infrastructure Cloud page.
4. Enter the password for the Cloud Software Library user who is allowed to import components from the Software Library. This password must be specified when the user accesses the Software Library and download the required components.

### 4.16.4 Configuring the Software Library User with External Management Agent

If the Management Agent is running on a remote machine, when importing the virtualization components, you must ensure that:

- The Software Library user must have Execute command privileges on the remote host.
- If Named Credentials are required to access the Management Agent on the remote host, the Software Library user must have View privileges while creating the Reference File locations for the Agent Storage Type.

To set up View privileges, follow these steps:

1. From the **Setup** menu, select **Security**, then select **Named Credentials**.

2. Select the Named Credentials to be used while creating the Software Library Reference Location of Storage Type: Agent.
3. Click **Add Grant**. Select the `CLOUD_SWLIB_USER` and click **Save**.

To set up Execute privileges on the remote host, follow these steps:

1. From the **Setup** menu, select **Security**, then select **Administrators**.
2. Select the `CLOUD_SWLIB_USER` and click **Edit**.
3. Navigate to the Target Privileges step. In the Target Privileges section, click **Add** and select the host target on which the Management Agent is running. The host target is added to the table.
4. Select the host target added in the Step 3 and click the **Edit** icon in the Manage Privilege Grants column. Select **Execute Command** in addition to the already existing View option and click **Continue**.
5. Click **Save** to save this configuration.

#### 4.16.5 Defining the Request Purging Policy

All deployment requests can be retained for a certain period and then purged. To specify the number of days for which the requests should be available, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Right click on the **Infrastructure Cloud** menu and select **Setup**.
3. Click the **Request Archiving Policy** link on the Infrastructure Cloud page.
4. Select the **Enable Archiving** check box, specify the Archiving Duration, and click **Apply**. The requests will be archived for the specified period and will then be purged.

### 4.17 Patching the Oracle VM Server

Oracle provides a publicly available YUM repository which includes repositories for Oracle VM. These repositories include the full releases of Oracle VM and mirror those releases on the Oracle VM DVDs/ISOs. The Oracle VM Server must be patched to ensure that the latest software updates and enhancements are applied. Patching the Oracle VM Server involves the following:

- Configuring the YUM Repository
- Upgrading the Virtual Server

#### 4.17.1 Configuring the YUM Repository

You can configure the YUM repository that is to be used to upgrade the virtual server. The configured YUM repository can be used for to upgrade all discovered virtual servers.

##### Prerequisites

- The YUM repository being configured must be present at a *http* location. You can configure an internal or public YUM repository.

To configure the YUM repository, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.

2. Right click on an OVM Manager target in the left panel, then select **Manage YUM Repository**.
3. Enter the following details in the YUM Repository page:
  - **Name:** Name of the YUM repository.
  - **YUM Base URL:** HURL for the YUM repository.
  - **Enable GPG Key:** Select this check box if you want to enable the GPG key. The GPG key checks the validity of the YUM repository and packages downloaded from the repository.
  - **YUM GPG Key:** If the Enable GPG Key field has been checked, specify the YUM GPG Key here.
4. Click **OK**. The YUM repository will be configured for the Oracle VM Manager.

### 4.17.2 Upgrading the Virtual Server

After you have configured the YUM repository, you must upgrade the virtual server to ensure that it has the latest updates. The virtual server being upgraded must be in maintenance mode. To upgrade the virtual server, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Right click on a VM Server in the left panel and select **Upgrade**.
3. A confirmation message appears. Click **OK** to proceed with the upgrade. After the virtual server has been upgraded, the virtual server will be restarted.

## 4.18 Acknowledging OVM Events

OVM events are generated for all operations performed on a server, server pool, storage, or network. The event can be Informational, Minor, or Critical. When a target is in an error state, a Critical event is generated. Targets with a Critical status cannot be used to perform any operations. The `EM_CLOUD_ADMINISTRATOR` can acknowledge such events as the associated targets cannot be used. To acknowledge an OVM event, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Click on a VM Server in the left panel to navigate to the VM Server Home page.
3. From the **VM Server** menu, select **OVM Events**. The list of unacknowledged events appears. Select an event from the list and click **Acknowledge** to acknowledge the event.

## 4.19 Managing the NTP Configuration

---

### Note:

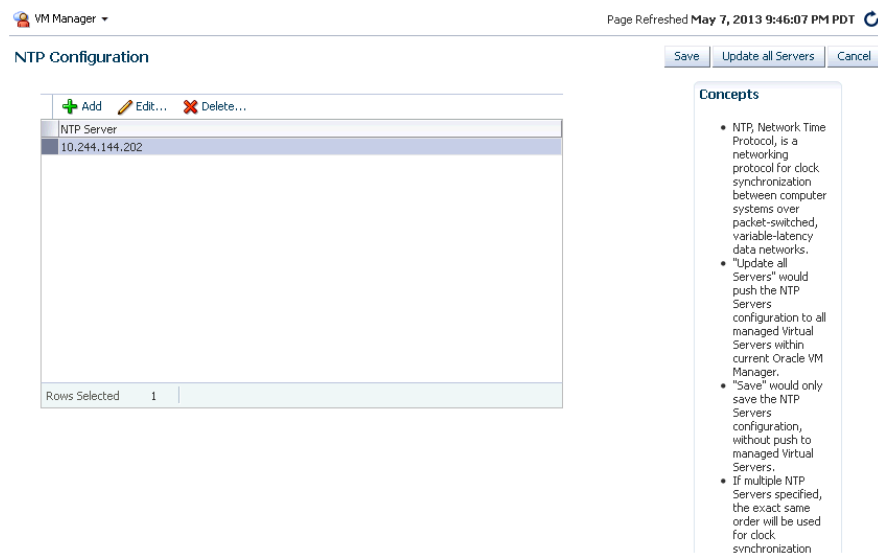
- This feature is available only with Oracle VM Manager 3.2.1 or later.
  - The NTP Server to be added must already be configured to provide the NTP service.
- 

Network Time Protocol (NTP) is a networking protocol that enables clock synchronization across all Oracle VM Servers managed by an Oracle VM Manager. By

default, when an Oracle VM Server is discovered, it is configured to use the Oracle VM Manager host computer as the NTP time source. You can configure a custom list of NTP servers that will be used to synchronize the clock when an Oracle VM Server is discovered. To create a custom list of NTP servers, follow these steps:

1. From the Enterprise menu, select Cloud, then select Infrastructure Home. The Infrastructure Cloud Home page appears.

**Figure 4–34 NTP Configuration**



2. Right click on an Oracle VM Manager target from the left panel and select **Manage NTP Configuration**.
3. To add an NTP server, click **Add**. In the Add NTP Server dialog box, enter the IP address or hostname of the NTP server in the IP Address/DNS Hostname field and click **OK**.
4. To edit an NTP server, select the server in the table and click **Edit**. In the Edit NTP Server dialog box, you are prompted for the name of the New NTP Server. Enter the IP address or hostname of the new NTP Server and click **OK**.
5. To delete an NTP server, select the server in the table and click **Delete**. A confirmation message is displayed. Click **OK** to delete the server or **Cancel** to cancel the operation.
6. To update the NTP server configuration on all Oracle VM Servers managed by the VM Manager, click **Update All Servers**. Any previous NTP configuration is overwritten.

Alternatively, you can save the NTP server configuration that you have added by clicking **Save**. This option saves the configuration but does not update the Oracle VM Servers. You can push the configuration to the VM Servers at a later date.

After any changes have been made the NTP server configuration, you must restart the VM Server for the changes to be effective.

---

**Note:** If multiple NTP configurations have been defined, the order in which have they have been defined will be used for clock synchronization.

---



## 4.20 Importing Virtual Machines

---

**Note:** This feature is available only with Oracle VM Manager 3.2.1 or later.

---

You can import a virtual machine into Oracle VM Manager from an external location. The imported virtual machine is placed under the VM Zone, VM Server Pool, or the VM Server you have selected. The virtual machine must be located on an FTP or web server, either as separate files, or compressed into a single archive file (for example, a .tgz or .zip file).

To import a virtual machine, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**. The Infrastructure Cloud Home page appears.
2. Right click on a VM Zone, VM Server Pool, or a VM Server target from the left panel and select **Import Virtual Machine**. The Import Virtual Machine: Select Target and Source page appears.

**Figure 4–35 Import Virtual Machine: Select Target and Source**

The screenshot shows the 'Import Virtual Machine: Select Target and Source' page. At the top, there is a breadcrumb trail: 'Select Target and Source' > 'Instance Details' > 'Schedule' > 'Review'. Below this is a progress bar indicating 'Step 1 of 4' with buttons for 'Back', 'Next', 'Review', 'Submit', and 'Cancel'. The main content area is divided into three sections: 'Destination', 'Source', and 'Request Details'. The 'Destination' section shows 'Target: NON\_CLUSTER\_POOL' and 'Target Type: Oracle VM Server Pool'. The 'Source' section has a text area for the URL of the virtual machine compress file. The 'Request Details' section has a text input for 'Request Name' with the value 'IMPORT\_GVM\_2013-05-08\_11-06-45-917'.

3. The destination target on which the virtual machine is to be imported and the name of the target is displayed.
4. Enter the following details:
  - **Source:** Enter the URL of the archive (tgz or gzip) that is to be imported as a virtual machine. The URL can be in format HTTP, HTTPS, or FTP.
  - **Request Name:** Enter a name for the virtual machine import request.
5. Click **Next**. The Import Virtual Machine: Instance Details page appears.

**Figure 4–36 Import Virtual Machine: instance Details**

6. Enter the following details:

- **Enable High Availability:** If you want to enable high availability for this guest virtual machine, select this checkbox. This option allows the guest virtual machine to be automatically restarted on other virtual servers in the pool in case the virtual server hosting the guest virtual machine fails.

---

**Note:** High availability must be enabled both on the virtual server pool and on the guest virtual machine. If it is not enabled on both, high availability is disabled.

---

- **Start VM After Creation:** Specify whether the guest virtual machine should be started automatically after it is created.
- **CPU Scheduling Priority:** Specify the priority of a guest virtual machine to access the physical CPUs, while competing with other guest virtual machines on the same virtual server. Higher the priority, higher is the stake of the guest virtual machine in claiming CPU cycles
- **CPU Scheduling Cap (%):** Restricts the amount of physical CPU that a guest virtual machine can use. Use this to constrain guest virtual machine resource allocation.
- **Keymap:** The keyboard mapping that will be used by the guest virtual machine being cloned.
- **Tags:** Specify one or more tags for the virtual machine. These tags are used to logical group the virtual machines and can be used to search for one or more virtual machines that meet a certain criteria.
- **VM Size:** The amount of memory to be allocated to the guest virtual machine. The virtual server should have sufficient memory to run the guest virtual machine. Otherwise, the guest virtual machine will be created, but will not be started. The size of the guest virtual machine (memory and CPU values). Select the VM Size which can be Small, Medium, Large, or Custom.

---

**Note:** The Root Password, Network, and Storage configuration cannot be modified for an imported virtual machine.

---

7. Click **Next**. In the Import Virtual machine: Schedule page, specify the Start and End Date. The Start Date is the date on which the request is being submitted and the End Date is the date on which the virtual machine is retired. Click **Next**.

8. Review the details of the virtual machine being imported and click **Submit**. From the Enterprise menu, select Job, then select Activity to view the Job Activity page. This page shows the details of the submitted job. You can click on the Name link to view more details.

## 4.21 Acknowledging the Repository Events

When the file system on the storage repository gets corrupted, a critical event is raised on the storage repository. This page displays all the critical events for the storage repository. To acknowledge repository events, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**. The Infrastructure Cloud Home page appears. Select an OVM Manager target from the left panel.
2. From the **VM Manager** menu, select **Administration**, then select **Storage Repository**.
3. A list of storage repositories that have been defined is displayed. Click on a storage repository name and click the **Repository Events** tab.
4. A list of unacknowledged critical events is displayed. Select an event from list and click **Acknowledge Events**. The file system associated with the acknowledged event cannot be used



---

## Defining and Managing Cloud Policies

This chapter describes the process to define policies for different target types and perform specific actions on virtualization targets. It contains the following sections:

- [Managing Cloud Policies](#)
- [Viewing Cloud Policies](#)
- [Viewing Policies for a Target](#)
- [Defining a Performance Policy](#)
- [Defining a Schedule-Based Policy](#)
- [Activating and Deactivating Policies](#)
- [Viewing Policy Groups](#)
- [Creating a Policy Group](#)
- [DRS and DPM Policies](#)

### 5.1 Managing Cloud Policies

Cloud management policies are rules that a `EM_CLOUD_ADMINISTRATOR` can define to take certain actions based on either performance metric thresholds or based on a schedule. You can create your own policies or use the ready-to-use policies. The ready-to-use policies provided by Oracle cannot be edited but you can do a *create like* cloning of the policy and then edit it as required.

---

**Note:** A self service user (`EM_SSA_USER`) can also create cloud policies for the servers for which he has access.

---

There are two kinds of cloud management policies: Performance-based policies and Schedule-based policies. Performance policies are based on the performance metrics of a specific type of target. Schedule based policies are executed based on the schedule you have defined. You can group one or more policies together to form policy groups. You can:

- View and Monitor Cloud Policies
- Define a new policy which can be of the following types:
  - Performance Policy
  - Schedule Policy
- Activate or Deactivate a Policy

- Define Policy Groups

## 5.2 Viewing Cloud Policies

You can view all the cloud policies and policy groups from the Cloud Policy Home page. You can view details like the status of the policy, type of policy, action type, policy associations and the schedule for the policy.

1. Log in to Enterprise Manager as a user with the `EM_CLOUD_ADMINISTRATOR` role and from the **Enterprise** menu, select **Cloud**, then select **Infrastructure Policies**.

**Note:** if you have logged into as a user with the `EM_SSA_USER` role, click the **Policies** tab in the Infrastructure Cloud Self Service Portal.

2. On the Policy Home page, you will see a list of a policies that you own and for which you have administration privileges.

**Figure 5–1 Cloud Policies**

The screenshot displays the 'Cloud Policies' interface. At the top, there's a header with 'Cloud Policies' and a refresh timestamp. Below this is a toolbar with actions like 'Create', 'View', 'Create Like...', 'Edit...', 'Delete', 'Activate', 'Deactivate', and a search bar. The main table lists several policies. The 'scale\_up\_tier\_instance' policy is highlighted. Below the table, a detailed view of the selected policy is shown, including its 'General' information (Type, Associations, Description, Status, Owner, Date Created, Last Modified, Target Type) and its 'Rule and Action' (Rule: At least one of Oracle Tier Instance/Oracle VM Guest/VM Guest/Physical CPU Utilization).

| Name                     | Type        | Associations | Target Type          | Status | Owner         | Date Created     | Last Modified    |
|--------------------------|-------------|--------------|----------------------|--------|---------------|------------------|------------------|
| scale_up_tier_instance   | Performance | 3            | Oracle Tier Instance | Active | CLOUD_ADMIN_1 | 2013-06-04 10:41 | 2013-06-04 10:41 |
| scale_down_tier_instance | Performance | 3            | Oracle Tier Instance | Active | CLOUD_ADMIN_1 | 2013-06-04 10:42 | 2013-06-04 10:42 |
| notification_perf_host   | Performance | 0            | Host                 | Active | CLOUD_ADMIN_1 | 2013-06-04 10:48 | 2013-06-04 10:48 |
| shutdown_vm              | Schedule    | 2            | Oracle VM Guest      | Active | CLOUD_ADMIN_1 | 2013-06-04 10:30 | 2013-06-04 10:30 |
| start_vm                 | Schedule    | 2            | Oracle VM Guest      | Active | CLOUD_ADMIN_1 | 2013-06-04 10:31 | 2013-06-04 10:31 |
| db_instance_perf_policy  | Performance | 1            | Database Instance    | Active | CLOUD_ADMIN_1 | 2013-06-04 10:34 | 2013-06-04 10:34 |

**scale\_up\_tier\_instance**

**General**

- Type: Performance
- Associations: 3
- Description: Policy to scale up tier instance by 1 node if VM CPU exceed 80%
- Status: Active
- Owner: CLOUD\_ADMIN\_1
- Date Created: 6/4/2013
- Last Modified: CLOUD\_ADMIN\_1
- Target Type: Oracle Tier Instance

**Rule and Action**

Rule: At least one of Oracle Tier Instance/Oracle VM Guest/VM Guest/Physical CPU Utilization

**Associations**

| Target               | Created By    | Last Updated On            |
|----------------------|---------------|----------------------------|
| VtService_tier_27_1  | CLOUD_ADMIN_1 | 2013-06-04 10:42:55 AM PDT |
| VtService_tier_99_1  | CLOUD_ADMIN_1 | 2013-06-04 10:43:18 AM PDT |
| VtService_tier_100_1 | CLOUD_ADMIN_1 | 2013-06-04 10:43:36 AM PDT |

3. Select a policy from the table. The policy details including the type of policy, performance (rule, action, action description), or schedule (schedule, action and action description), the status, date on which it was created, and the target type are displayed.
4. You can perform the following actions on this page:
  - **Create:** To create a new policy, from the **Create** menu, select either of the following:
    - Schedule Policy: See [Section 5.5, "Defining a Schedule-Based Policy"](#).
    - Performance Policy: See [Section 5.4, "Defining a Performance Policy"](#).
  - **Create Like:** Select a policy and click **Create Like** to make a copy of an existing policy.
  - **Edit:** Select a policy and click **Edit** to launch the Edit Policy wizard. You can modify the description of the policy and add metrics to the policy. You cannot

change the name of the policy, the type of policy, the target with which the policy is associated, and the actions that can be performed.

- **Activate or Deactivate a Policy:** See [Section 5.6, "Activating and Deactivating Policies"](#).
- **Create Policy Associations:** In the Associations region, you can view a list of targets that have been associated with a selected policy. A policy can be associated with one or more targets. Click **Add**. In the Search and Select: Targets window, choose a target and click **Select** to associate the target with the policy. After you have associated one or more targets with the policy, you can:
  - Click on a target to drill down to the Policies page for the target. See [Section 5.3, "Viewing Policies for a Target"](#).
  - Select a target and click **View History** to view the Policy Action History.
  - Select a policy and click **Remove** to delete the associations for the policy.

---



---

**Note:**

- Associations cannot be added for inactive policies.
  - The **View History** option is not available for users with the EM\_SSA\_USER role.
- 
- 

## 5.3 Viewing Policies for a Target

After you have associated a target with a policy, you can view all policies that can be associated with a target on the Policies page.

To view policies that can be associated with a selected target, follow these steps:

1. Navigate to the Cloud Policies page.
2. Select a policy from the list of active policies in the list. In the Associations region, click on a target link to drill down to the Policies page for the target.

**Figure 5–2 Policies - Target**

Policies > Policies for VtService\_assembly\_vm\_1\_1

**Policies for VtService\_assembly\_vm\_1\_1**

This page shows all policies that can be associated with target instance 'VtService\_assembly\_vm\_1\_1'. For each policy, you can see if the association is via the standalone method, template collection, policy group, or a combination of the three. In this page, you can manipulate (add or remove) only standalone associations.

View ▾ Associate Disassociate

| Name            | Description                                                         | Type     | Association |                     |              | Owner         | Created On       | Last Modified    |
|-----------------|---------------------------------------------------------------------|----------|-------------|---------------------|--------------|---------------|------------------|------------------|
|                 |                                                                     |          | Standalone  | Template Collection | Policy Group |               |                  |                  |
| stop_vm_policy  | Schedule policy to shutdown Oracle VM Guest on every Monday morning | Schedule | ✓           |                     |              | CLOUD_ADMIN_1 | 2013-06-12 09:51 | 2013-06-12 09:51 |
| start_vm_policy | Schedule policy to start Oracle VM Guest on every Monday morning    | Schedule | ✓           |                     |              | CLOUD_ADMIN_1 | 2013-06-12 09:51 | 2013-06-12 09:51 |

**start\_vm\_policy**

**General**

Type: Schedule

Associations: 4

Description: Schedule policy to start Oracle VM Guest on every Monday morning

Status: Active

Owner: CLOUD\_ADMIN\_1

Date Created: 2013-06-12 09:51:31 AM PDT

Last Modified: 2013-06-12 09:51:31 AM PDT

Target Type: Oracle VM Guest

**Schedule and Action**

Weekly on Monday beginning Jun 17, 2013 6:00:00 AM

**Policy Action History: start\_vm\_policy**

View ▾

| Action | Status    | Created On             | Last Updated On        |
|--------|-----------|------------------------|------------------------|
| Start  | Submitted | 2013-06-12 09:53:36 AM | 2013-06-12 09:53:36 AM |

3. All policies that can be associated with the selected target are listed on this page. For each policy, the type of policy association that has been defined is displayed. This can be:

- Standalone
- Template Collection
- Policy Group

For each policy, the following details are displayed:

- Name and description of the policy.
- The type of policy which can be Schedule or Performance.
- The policy association method. A policy can be associated with a standalone association, template collection, a policy group or a combination of the three.
- The General section displays details like status of the policy, the number of associations, the date on which the policy was created, and the type of policy.
- The Policy Action History region shows the history of associations for this target.

---

**Note:**

- The history of policy associations is displayed only for users with the `EM_CLOUD_ADMINISTRATOR` role.
  - The Policy Action History region is blank if the action has a **Failed** status.
- 

4. For policies that are active, you can do the following:
  - Select a policy and click **Associate** to associate it with the target.
  - Select a policy that has already been associated and click **Disassociate** to delete the policy association.

## 5.4 Defining a Performance Policy

Performance policies are formed based on the performance metrics of a specific target type and its associated targets. You can create conditions and rules for the policy based on the metrics.

To define a performance policy, follow these steps:

1. Depending on your user role, use the menu options to navigate to the Infrastructure Policies page.
2. On the Policy Home page, you will see a list of policies that you own and for which you have administration privileges.
3. From the **Create** menu, select **Performance Policy**.
4. In the General Information page, enter the name and description of the policy. Select the Target Type with which the policy is to be associated.



Figure 5–3 Create Performance Policy - General

Cloud Policies

General

Rule

Action

Review

Create Performance Policy : General

BackStep 1 of 4NextCancel

\* Name

perf\_scale\_up\_policy

Description

Performance policy to scale up 1 node Oracle Tier Instance if Oracle VM Guest CPU utilization is more than 90% for 5 times.

Type

Performance

\* Target Type

Oracle Tier Instance

5. Click **Next**. In the Rule page, select a metric from the right panel and click **Add to Definition**, or drag and drop metric from the Metrics column to the Definition column. Specify the condition, value, and number of occurrences. In the Match radio button, you can select
- All: All conditions must be met before the policy is executed on the target.
  - Any: Any of the conditions must be met for the policy to be executed.

Figure 5–4 Create Performance Policy - Rule

Cloud Policies

General

Rule

Action

Review

Create Performance Policy : Rule

BackStep 2 of 4NextCancel

A rule can have one or more metrics. Select the required metric and add it to the definition table.  
( If Cluster.CPU utilization above 90% for 5 occurrences OR Service1.responseTime above 7 seconds for 2 occurrences then "Scale out" cluster. )

Definition

View

Remove

Match

All

Any

| Condition      | Target Type     | Metric                | * Operator | * Value |
|----------------|-----------------|-----------------------|------------|---------|
| At least one o | Oracle VM Guest | OVIM Guest Load,Phys: | >          | 90      |

Metrics

View

Add to Definition

Oracle Tier Instance

Oracle VM Guest

OVIM Guest Load,Physical CPU Utilization

OVIM Guest DiskActivity Summary.Total Throughput (kb/s)

OVIM Guest Network Activity Summary.Total Throughput (kb/s)

6. Click **Next**. Specify the action to be performed when the policy condition is met. The list of available actions can vary depending on the selected target type.
- If notifications are disabled, you will see the following message:
- Notification action is currently disabled. Contact your system administrator to set up the email server information.

**Figure 5–5 Create Performance Policy - Action**

Cloud Policies

General Rule **Action** Review

Create Performance Policy : Action

Back Step 3 of 4 Next Cancel

Rule At least one of Oracle Tier Instance/Oracle VM Guest/OVM Guest Load.Physical CPU Utilization is greater than 90 for 5 consecutive occurrence(s)

\* Applies to Target Oracle Tier Instance

Action Scale Up

Description Notification action is currently disabled. Contact your system administrator to set up the email server information.

| Parameter       | Value |
|-----------------|-------|
| Number of Nodes | 1     |

- Review the policy and click **Save** to define this policy and return to the Policy Home page.

---

**Note:** After you have defined the policy, you can associate one or more targets with the policy.

---

## 5.5 Defining a Schedule-Based Policy

Schedules based on policies are based on a predefined schedule. To define a schedule-based policy, follow these steps:

- Depending on your user role, use the menu options to navigate to the Infrastructure Policies page.
- On the Policy Home page, you will see a list of a policies that you own and for which you have administration privileges.
- From the Create menu, select **Schedule Policy**.

**Figure 5–6 Create Schedule Policy**

Cloud Policies

Create Schedule Policy

OK Cancel

\* Name

Description

Type Schedule

\* Target Type Oracle VM Guest

\* Policy Action Start

Start 6/2/2011 4:26:55 PM (UTC+00:00) Universal Time

Repeat Every N minutes

Frequency Every 1 minute(s)

Grace Period Do not run if it cannot start within 1 hours of the scheduled start time

Duration Indefinitely For 1 hours

- In the Create Schedule Policy page, enter the name, description of the policy, and the Target Type with which the policy is to be associated.
- Specify the Grace Period, frequency, and schedule for the policy.
- When the policy conditions are met, the policy is executed as per the schedule on the selected target. The list of available actions can vary depending on the target type with which the policy is associated.
- Click **Save** to define this policy.

---

**Note:** After you have defined the policy, you can associate one or more targets with the policy.

---

## 5.6 Activating and Deactivating Policies

Once a policy is defined, it becomes active and can be applied to one or more targets. You can deactivate an active policy by following these steps:

1. Depending on your user role, use the menu options to navigate to the Infrastructure Policies page.
2. On the Policy Home page, you will see a list of a policies that you own and have administration privileges.
3. Select a policy from the table and click **Deactivate**. The policy is deactivated and the status of the policy is marked as Inactive. You cannot associate this policy to any target until it is activated and its current associations will be suspended for further evaluations.

Once a policy is deactivated, it is marked as Inactive. To activate an Inactive Policy, select the policy from the table and click **Activate**. In the confirmation window, click **Activate** again. The policy is now marked as Active.

## 5.7 Viewing Policy Groups

A policy group is a collection of policies that can be applied together and evaluated to generate a compliance score. To view the policy groups, follow these steps:

1. Log in to Enterprise Manager as a user with the EM\_CLOUD\_ADMINISTRATOR role and from the **Enterprise** menu, select **Cloud**, then select **Infrastructure Policies**.

---

**Note:** if you have logged into as a user with the EM\_SSA\_USER role, click the **Policies** tab in the Infrastructure Cloud Self Service Portal.

---

2. The Cloud Policies page appears. Click on the **Policy Groups** tab.

**Figure 5–7 Policy Groups**

Cloud Policies Page Refreshed Jun 12, 2013 11:43:47 PM PDT

---

| Policy Groups    |             |          |              |               |                     |                  |                     |
|------------------|-------------|----------|--------------|---------------|---------------------|------------------|---------------------|
| Group Name       | Description | Policies | Associations | Owner         | Date Created        | Last Modified By | Last Modified       |
| New Policy Group |             | 3        | 2            | CLOUD_ADMIN_1 | 2013-06-12 23:00:00 | CLOUD_ADMIN_1    | 2013-06-12 23:00:00 |

---

**New Policy Group**

| Associations                |             | Policies Included   |        |               |                        |                  |                        |
|-----------------------------|-------------|---------------------|--------|---------------|------------------------|------------------|------------------------|
| Name                        | Type        | Target Type         | Status | Owner         | Date Created           | Last Modified By | Last Modified          |
| > start_vm_policy           | Schedule    | Oracle VM Guest     | Active | CLOUD_ADMIN_1 | 2013-06-12 09:51:31 AM | PE CLOUD_ADMIN_1 | 2013-06-12 09:51:31 AM |
| > vm_perf_scale_up_policy   | Performance | Oracle Tier Instanc | Active | CLOUD_ADMIN_1 | 2013-06-12 10:01:02 AM | PE CLOUD_ADMIN_1 | 2013-06-12 10:02:59 AM |
| > vm_perf_scale_down_policy | Performance | Oracle Tier Instanc | Active | CLOUD_ADMIN_1 | 2013-06-12 10:02:45 AM | PE CLOUD_ADMIN_1 | 2013-06-12 10:02:45 AM |

3. The list of policy groups that have already been defined are displayed. You can do the following:
  - **Create a Policy Group:** See [Section 5.8, "Creating a Policy Group"](#).
  - **Create Like:** Select a policy group and click **Create Like** to make a copy of an existing policy group.
  - **Edit:** Select a policy group and click **Edit** to launch the Edit Policy Group wizard. You can add or remove policies from the policy group. You cannot modify the name of the policy group.
  - **Delete:** Select a policy group and click **Delete**. When a policy group is deleted, all associations defined for the policy group will also be deleted.

---

**Note:** You can edit or delete policy groups that you have created or for which you have the privileges.

---

- **Create Policy Group Associations:** Click the **Associations** tab in the Policy Group region, to view a list of targets that have been associated with the selected policy group.

**Figure 5–8 Policy Group: Associations**

Cloud Policies Page Refreshed Jun 13, 2013 8:36:12 AM PDT

Policies

Policy Groups

View ▾

Create... Create Like... Edit... Delete Search Group Name ▾

| Group Name       | Description | Policies | Associations | Owner         | Date Created        | Last Modified By | Last Modified       |
|------------------|-------------|----------|--------------|---------------|---------------------|------------------|---------------------|
| New Policy Group |             | 3        | 2            | CLOUD_ADMIN_1 | 2013-06-12 23:00:00 | CLOUD_ADMIN_1    | 2013-06-12 23:00:00 |
| Policy Group II  |             | 1        | 0            | CLOUD_ADMIN_1 | 2013-06-13 00:00:00 | CLOUD_ADMIN_1    | 2013-06-13 00:00:00 |

New Policy Group

Associations

Policies Included

View ▾

| Name                      | Type        | Target Type          | Status | Owner         | Date Created           | Last Modified By | Last Modified          |
|---------------------------|-------------|----------------------|--------|---------------|------------------------|------------------|------------------------|
| start_vm_policy           | Schedule    | Oracle VM Guest      | Active | CLOUD_ADMIN_1 | 2013-06-12 09:51:31 AM | PC CLOUD_ADMIN_1 | 2013-06-12 09:51:31 AM |
| vm_perf_scale_up_policy   | Performance | Oracle Tier Instance | Active | CLOUD_ADMIN_1 | 2013-06-12 10:01:02 AM | PC CLOUD_ADMIN_1 | 2013-06-12 10:02:59 AM |
| vm_perf_scale_down_policy | Performance | Oracle Tier Instance | Active | CLOUD_ADMIN_1 | 2013-06-12 10:02:45 AM | PC CLOUD_ADMIN_1 | 2013-06-12 10:02:45 AM |

A policy group can be associated with one or more targets. Click **Add**. In the Search and Select: Targets window, choose a target and click **Select** to associate the target with the policy group. After you have associated one or more targets with the policy group, you can:

- Click on a target to drill down to the Policies page for the target. See [Section 5.3, "Viewing Policies for a Target"](#).
- Select a target and click **View History** to view the Policy Group Action History.
- Select a target and click **Remove** to delete the policy group association.
- **View Policies:** Click the **Policies Included** tab in the Policy Group region to view the list of policies included in the policy group.

**Figure 5–9 Policy Group: Policies Included**

Cloud Policies Page Refreshed Jun 13, 2013 8:36:12 AM PDT

Policy Groups

View + Create... Create Like... Edit... Delete Search Group Name

| Group Name       | Description | Policies | Associations | Owner         | Date Created        | Last Modified By | Last Modified       |
|------------------|-------------|----------|--------------|---------------|---------------------|------------------|---------------------|
| New Policy Group |             | 3        | 2            | CLOUD_ADMIN_1 | 2013-06-12 23:00:00 | CLOUD_ADMIN_1    | 2013-06-12 23:00:00 |
| Policy Group II  |             | 1        | 0            | CLOUD_ADMIN_1 | 2013-06-13 00:00:00 | CLOUD_ADMIN_1    | 2013-06-13 00:00:00 |

New Policy Group

Associations Policies Included

View +

| Name                      | Type        | Target Type          | Status | Owner         | Date Created              | Last Modified By | Last Modified             |
|---------------------------|-------------|----------------------|--------|---------------|---------------------------|------------------|---------------------------|
| start_vm_policy           | Schedule    | Oracle VM Guest      | Active | CLOUD_ADMIN_1 | 2013-06-12 09:51:31 AM PT | CLOUD_ADMIN_1    | 2013-06-12 09:51:31 AM PT |
| vm_perf_scale_up_policy   | Performance | Oracle Tier Instance | Active | CLOUD_ADMIN_1 | 2013-06-12 10:01:02 AM PT | CLOUD_ADMIN_1    | 2013-06-12 10:02:59 AM PT |
| vm_perf_scale_down_policy | Performance | Oracle Tier Instance | Active | CLOUD_ADMIN_1 | 2013-06-12 10:02:45 AM PT | CLOUD_ADMIN_1    | 2013-06-12 10:02:45 AM PT |

## 5.8 Creating a Policy Group

To define a policy group, follow these steps:

1. Depending on your user role, use the menu options to navigate to the Infrastructure Policies page.
2. Click on the **Policy Groups** tab. On the Policy Groups page, you will see a list of a policy groups that have been defined.
3. Click **Create**. The Create Policy Group page appears.

**Figure 5–10 Policy Group: Create**

Create Policy Group OK Cancel

\* Group Name Policy Group III

Description

Policies Included

View + Add... Remove

There are no policies included in this policy group.

4. Click **Add**. A list of policies that can be added to the policy group are displayed. Select one or more policies from the list and click **Select**.
5. The selected policies will be added to the policy group. Click **OK** to save this policy group.
6. Click **OK** to save this policy group.

### 5.8.1 Example 1

If you want to shut down a guest virtual machine every Friday at 6:00 pm and restart it every Monday at 9:00 pm, you can create policies as follows:

- First, create a schedule based policy where the guest virtual machine shuts down every Friday at 6:00 pm. See [Section 5.8.1.1, "Policy for Shutting Down Guest Virtual Machine"](#).

- Secondly, create another schedule based policy where the guest virtual machine is restarted every Monday at 6:00. See [Section 5.8.1.2, "Policy for Starting Up Guest Virtual Machine"](#).
- Combine these two policies into a Policy Group and associate these policies with one or more guest virtual machine targets. See [Section 5.8.1.3, "Creating a Policy Group \(Start Up and Shut Down\)"](#).

### 5.8.1.1 Policy for Shutting Down Guest Virtual Machine

To shut down a guest virtual machine every Friday, you can set up the policy as follows:

1. Depending on your user role, use the menu options to navigate to the Infrastructure Policies page.
2. In the Cloud Policies page, select **Schedule Policy** from the **Create** menu. In the Create Schedule Policy page, enter the following details:
  - Enter a name and description for the policy.
  - Select the Target Type as **Oracle VM Guest**.
  - Select **Stop** in the Policy Action field.
  - In the Start field, specify the date and time from which this policy should be effective.
  - In the Repeat field, select **Specific Days of the Week** and select **Friday** in the Days of Week checkbox.

**Figure 5–11 Create Schedule Policy (Example 1)**

The screenshot shows the 'Create Schedule Policy' form in the 'Cloud Policies' section. The form has the following fields and values:

- Name:** guestvm\_policy1
- Description:** (empty text area)
- Type:** Schedule
- Target Type:** Oracle VM Guest
- Policy Action:** Stop
- Description:** Stop a Guest Virtual Machine
- Start:** 8/10/2012 6:00:00 PM (GMT+00:00) GMT
- Repeat:** Specific days of the week
- Days of Week:** Monday, Tuesday, Wednesday, Thursday, ☒ Friday, Saturday, Sunday
- Grace Period:** Do not run if it cannot start within 1 hours of the scheduled start time

3. Click **OK** to save this policy and return to the Cloud Policies page.

### 5.8.1.2 Policy for Starting Up Guest Virtual Machine

To start a guest virtual machine every Monday, navigate to the Cloud Policy page and select the policy you just created. Perform the following steps:

1. Click **Create Like** to make a copy of the policy.
2. In the Cloud Policies page, select **Schedule Policy** from the **Create** menu. In the Create Schedule Policy page, enter the following details:
  - Enter a name and description for the policy.
  - Select the Target Type as **Oracle VM Guest**.

- Select **Start** in the Policy Action field.
  - In the Start field, specify the date and time from which this policy should be effective.
  - In the Repeat field, select **Specific Days of the Week** and select **Monday** in the Days of Week checkbox.
3. Click **OK** to save this policy and return to the Cloud Policies page.

### 5.8.1.3 Creating a Policy Group (Start Up and Shut Down)

After defining the two policies for the guest virtual machine target type, you can combine these policies into a policy group and associate these policies to one or more targets. To create a policy group, follow these steps:

1. Depending on your user role, use the menu options to navigate to the Infrastructure Policies page.
2. On the Cloud Policies page, click the **Policy Groups** tab.
3. Click **Create**. The Cloud Policy Group page appears.
4. Enter a name and description for the group.
5. Click **Add** and select the two schedule policies that you have created.
6. Click **OK** to return to the Cloud Policies page and click on the Policy Group that you have created.
7. The details of the policy group that you have created such as the name of the policy group, number of policies associated with this group, number of targets associated, and so on are displayed.
8. Click the **Policies Included** tab in the bottom region to view the policies associated with this policy group.
9. Click the **Associations** tab to view the list of guest virtual machines associated with this policy. Click **Add** to associate more guest virtual machines with this policy group.
10. When this policy group is active, the guest virtual machines associated with the policies are shut down and started up as per the schedule.

## 5.8.2 Example 2

Suppose you want to define a policy group for an Oracle Tier Instance target where the tier instance is scaled down to 1 every Friday and scaled up to 2 every Monday. You can create the following policies:

- First, create a Schedule Policy to scale down the tier instance to 1. See [Section 5.8.2.1, "Policy to Scale Down a Tier Instance"](#).
- Next, create another policy to scale up the tier instance to 2 every Monday. See [Section 5.8.2.2, "Policy to Scale Up a Tier Instance"](#).
- Combine these policies and create a Policy Group. See [Section 5.8.2.3, "Creating a Policy Group \(Scale Up and Scale Down\)"](#).

### 5.8.2.1 Policy to Scale Down a Tier Instance

To scale down a tier instance every Friday, you can set up a performance policy as follows:

1. Depending on your user role, use the menu options to navigate to the Infrastructure Policies page.
2. In the Cloud Policies page, select **Schedule Policy** from the **Create** menu. The Create Schedule Policy page appears.

**Figure 5–12 Schedule Policy Example 1**

The screenshot shows the 'Create Schedule Policy' form in the 'Cloud Policies' section. The form has the following fields and values:

- Name:** (empty text field)
- Description:** (empty text area)
- Type:** Schedule
- Target Type:** Oracle Tier Instance
- Policy Action:** Scale Down
- Action Parameters:**

| Parameter       | Value |
|-----------------|-------|
| Number of Nodes | 1     |
- Start:** 6/13/2013 9:49:42 AM (GMT-08:00) Los Angeles - Pacific Time (PT)
- Repeat:** Specific days of the week
- Days of Week:** Monday ☐ Tuesday ☐ Wednesday ☐ Thursday ☐ Friday ☒ Saturday ☐ Sunday
- Grace Period:** ☐ Do not run if it cannot start within 1 hours of the scheduled start time

3. Enter a name and description for the policy and specify the following details:
  - **Target Type:** Select the **Tier Instance** target type.
  - **Policy Action:** Select the **Scale Down** policy action.
  - **Repeat:** In the Repeat field, select **Specify Days of the Week**.
  - **Days of the Week:** Select **Friday** here.
4. Click **OK** to save the policy.

### 5.8.2.2 Policy to Scale Up a Tier Instance

To scale down a tier instance every Friday, you can set up a performance policy as follows:

1. Depending on your user role, use the menu options to navigate to the Infrastructure Policies page.
2. In the Cloud Policies page, select **Schedule Policy** from the **Create** menu. The Create Schedule Policy page appears.



**Figure 5–13 Schedule Policy Example II**

**Cloud Policies**

**Create Schedule Policy** [OK] [Cancel]

\* Name: Scale Up Policy

Description: [Empty text area]

Type: Schedule

\* Target Type: Oracle Tier Instance

\* Policy Action: Scale Up

Description: [Empty text area]

Action Parameters:

| Parameter       | Value |
|-----------------|-------|
| Number of Nodes | 2     |

Start: 6/13/2013 10:14:07 AM (GMT-08:00) Los Angeles - Pacific Time (PT)

Repeat: Specific days of the week

Days of Week: ☒ Monday ☐ Tuesday ☐ Wednesday ☐ Thursday ☐ Friday ☐ Saturday ☐ Sunday

Grace Period: ☐ Do not run if it cannot start within 1 hours of the scheduled start time

3. Enter a name and description for the policy and specify the following details:
  - **Target Type:** Select the **Tier Instance** target type.
  - **Policy Action:** Select the **Scale Up** policy action.
  - **Action Parameters:** Enter 2 for the **Number of Nodes** parameter.
  - **Repeat:** In the Repeat field, select **Specify Days of the Week**.
  - **Days of the Week:** Select **Monday** here.
4. Click **OK** to save the policy.

### 5.8.2.3 Creating a Policy Group (Scale Up and Scale Down)

After defining the two policies for the Oracle Tier Instance target type, you can combine these policies into a policy group and associate these policies to one or more targets. To create a policy group, follow these steps:

1. Depending on your user role, use the menu options to navigate to the Infrastructure Policies page.
2. On the Cloud Policies page, click the **Policy Groups** tab.
3. Click **Create**. The Cloud Policy Group page appears.
4. Enter a name and description for the group.
5. Click **Add** and select the two schedule policies that you have created.
6. Click **OK** to return to the Cloud Policies page and click on the Policy Group that you have created.

## 5.9 DRS and DPM Policies

The Dynamic Resource Scheduling (DRS) and Distributed Power Management (DPM) policies are automated policies that improve quality of service and reduce power consumption by dynamically relocating running virtual machines within a server pool based on the load.

Dynamic Resource Scheduling (DRS) policies provide real-time monitoring of Oracle VM Server utilization to rebalance a server pool and provide consistent resources to the running virtual machines. DRS migrates VMs away from heavily loaded Oracle VM Servers to those servers with more resources available.

Distributed Power Management (DPM) policies optimize the server pool for minimum power consumption. DPM complements DRS by reducing the number of powered-on servers in the pool when there are periods of low resource utilization. It can automatically power-on capacity as needed when resource utilization ramps up.

To set up a DRS / DPM policy, follow these steps:

1. Log in as a user with the `EM_CLOUD_ADMINISTRATOR` role, and from the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Right click on a VM Server Pool target in the left panel and select **Manage DRS/DPM Policies**. Click **Configure**. The Configure DRS/DPM General page appears.

**Figure 5–14 Configure DRS/DPM: General**

The screenshot shows the 'Configure DRS/DPM: General' configuration page. At the top, there are tabs for 'General' and 'Network', with 'General' active. Below the tabs, the title is 'Configure DRS/DPM: General' and 'Policy Control'. On the right, there are 'Back', 'Step 1 of 2', 'Next', and 'Cancel' buttons. The 'Policy Type' is set to 'Distributed Resource Scheduler' (indicated by a dropdown arrow), and the 'Time Period (Minutes)' is set to '10' (indicated by a dropdown arrow). Below this is the 'Virtual Servers' section, which includes a 'View' dropdown, 'Add...' and 'Remove' buttons, and a table with a 'Name' column. At the bottom is the 'Constraints' section, specifically for 'Virtual Server CPU'. It has an 'Enable' checkbox checked and a 'Threshold (%)' dropdown set to '90'.

3. Select the **Policy Type**. This can be:
  - **Distributed Resource Scheduler:** Select this policy type to optimize virtual machine resource utilization in a server pool.
  - **Distributed Power Management:** Select this policy type to increase the consolidation ratio to use fewer Oracle VM Servers during periods of relative low resource utilization.
4. Specify the **Time Period (Minutes)**. This is the time period for the policy job to run. This sets the policy job to run every n minutes, for example, 10 sets the policy job to run every 10 minutes. You can enter a number between 1 and 60.
5. Click **Add** in the Virtual Servers region to add one or more virtual server targets for which the policy is to be enabled.
6. Specify the Constraints for the Virtual Server CPU as follows:
  - **Enable:** Set this field to enable or disable logging of CPU performance and utilization.
  - **Threshold:** The maximum amount of CPU percentage usage allowed before the policy must be enacted. You can enter between 0 and 99.
7. Click **Next**. The Configure DRS/DPM Network page appears.

**Figure 5–15 Configure DRS/DPM Policies: Network**

| Network       | Enable                              | Threshold (%) |
|---------------|-------------------------------------|---------------|
| Network       | <input checked="" type="checkbox"/> | 51            |
| cloud_network | <input type="checkbox"/>            | 0             |

8. Select the **Enable** checkbox to indicate whether the policy is to be enabled on the network, and select the threshold at which the policy is to be enacted for the network.
9. Click **Apply**.

---

**Note:** Once the policy is created, it is enabled and applicable to all the virtual server pool targets that have been included in the policy. To disable the DRS/DPM policy, create a schedule based policy for the Oracle VM Server Pool target type and in the Policy Action field, select **Disable DRS/DPM** policies. See [Section 5.5, "Defining a Schedule-Based Policy"](#) for more details on creating a schedule based policy. You can enable the policy by editing the Schedule Policy and select **Enable DRS** or **Enable DPM** as necessary.

---

### 5.9.1 Creating Automated DRS and DPM Policies

You can create automated DRS or DPM policies by following these steps:

1. Log in as a user with the `EM_CLOUD_ADMINISTRATOR` role and from the **Enterprise** menu, select **Cloud**, then select **Infrastructure Policies**.
2. From the **Create** menu, select **Schedule Policy**.
3. In the Create Schedule Policy page, enter a name and description for the policy.
4. Select Oracle VM Server Pool as the Target Type.
5. Select **Enable DRS** or **Enable DPM** Policy from the Policy Action field.
6. Specify the schedule and frequency for the policy and click **OK** to save the policy.
7. In the Cloud Policies page, select the policy you have created and click **Add** in the Associations region.
8. Select the Oracle VM Server Pool target with which the policy is to be associated and click **Select**.
9. The DRM or DPM policy is now automatically enabled for the selected target.



## Setting Up the IaaS Self Service Portal

This chapter describes the process of setting up the Self Service Portal. It contains the following sections:

- [Setting Up the Self Service Portal](#)
- [Uploading Large Files on the Self Service Portal](#)

### 6.1 Setting Up the Self Service Portal

The Self Service Portal allows administrators to configure the cloud and perform various operations such as deciding the total amount of resources each user can reserve, quota for each user role, and so on. Users can log in to the Self Service Portal and identify the images they need, reserve resources (for example, compute power, storage, memory, and so on) and define policies specific to their system.

Before the Self Service Portal can be used, the user quota, roles for zone access, and software components must be defined.

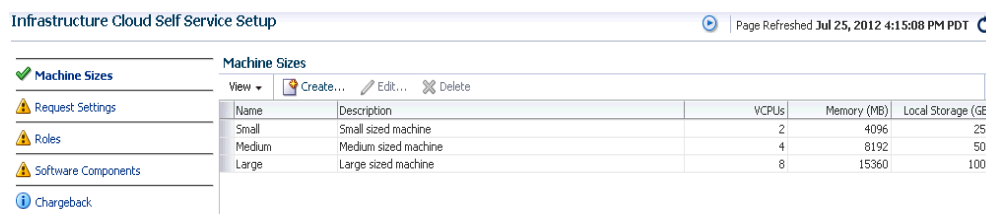
#### Prerequisites

- Enterprise Manager Cloud Control is set up to manage the cloud environment.
- The Enterprise Manager administrator has the privileges to set up the self service application. This user must have the EM\_SSA\_ADMINISTRATOR role. See [Section 3.4, "Defining Roles and Assigning Users"](#) for details.

Follow this process to set up and configure the Self Service Portal.

1. From the **Setup** menu, select **Cloud**, then select **Infrastructure**. Click **Machine Sizes** in the left panel. Small, Medium, and Large machine sizes are available out-of-the-box.

**Figure 6–1 Self Service Portal Setup - Machine Sizes**



The screenshot shows the 'Infrastructure Cloud Self Service Setup' page. On the left is a navigation menu with 'Machine Sizes' selected. The main area displays a table of machine sizes. Above the table are buttons for 'Create...', 'Edit...', and 'Delete'. The table has columns for Name, Description, VCPUs, Memory (MB), and Local Storage (GB). Three machine sizes are listed: Small, Medium, and Large.

| Name   | Description          | VCPUs | Memory (MB) | Local Storage (GB) |
|--------|----------------------|-------|-------------|--------------------|
| Small  | Small sized machine  | 2     | 4096        | 250                |
| Medium | Medium sized machine | 4     | 8192        | 500                |
| Large  | Large sized machine  | 8     | 15360       | 1000               |

- To view a demo on setting up the self service portal, click the video link next to the Page Refreshed section of the page. Click **Create** to add a machine size. You are prompted for the Name (internal name of the machine size), number

of VCPUs, amount of memory and local storage to be assigned to machines of this size. Click **Save** to add this machine size.

- Select a machine size from the list and click **Edit**. You can configure the machine sizes by specifying the amount of memory, number of CPUs, and the local storage for each machine. While requesting a server, the SSA user can specify any of the machine sizes defined here.
2. Click **Request Settings** in the left panel to define request related policies. On this page, you can enter the following details:
- **Future Reservation:** Indicate how far in advance you can make a request.
  - **Request Duration:** The maximum duration for which requests should be made.
  - **Configure EM Agent:** Select the **Enable EM Agent Configuration** check box to configure the Management Agent on the servers while the self service user's request is being processed.

---

**Note:** If you have chosen to configure the Management Agent, you must ensure the following:

- Write permissions to the directory specified in the **Installation Base Directory** field must be available.
  - Write permissions to the `/home/oracle/agent` directory specified during assembly or template deployment must be available.
  - Necessary software components must be configured.
- 

- **Allow Partial Deployment:** Select this option to enable partial deployment. By default, when a deployment fails, the assembly instance, tier instance, and the Guest VMs that have been created are automatically deleted. But if partial deployment is enabled, this clean up activity is disabled.
- **Software Library Top Folder:** Select the top level folder in the Software Library in which user specific folders need to be created. This folder is used by the SSA users to save their environment and store their assemblies and templates.

**Figure 6–2 Self Service Portal Setup - Request Settings**

Infrastructure Cloud Self Service Setup Page Refreshed Jan 11, 2013 9:29:46 PM PST

**Request Settings** Apply Revert

**Machine Sizes**

**Request Settings**

**Roles**

**Software Components**

**Chargeback**

**Future Reservation**  
How far in advance a user can make a request  
Future Reservation ☒ No Restriction ☐ Restricted

**Request Duration**  
Maximum duration for which requests can be made  
Request Duration ☒ No Restriction ☐ Restricted  
Default Duration

**Configure EM Agent**  
EM Agent will be configured on the servers while processing the SSA user's request. The EM Agent configuration properties should be set for assemblies and templates using the 'Software Components' page under SSA Setup.  
☐ Enable EM Agent Configuration

**Allow Partial Deployment**  
During assembly deployment if the step that does product configuration fails, then the assembly instance, tier instance and Guest VMs that were created are cleaned up. The administrator can disable this clean up activity by checking the "Allow Partial Deployment" check box.  
☒ Allow partial deployment in case of the product configuration step while deploying an assembly fails.

**Software Library top level folder**  
Choose the top level folder in 'Software Library' under which user specific folders will be created.  
Folder Name

- Click **Roles** in the left panel to define the mappings between roles and other entities like zones, quota, and network profiles. The SSA administrator can control the quota that the SSA user can request which ensures that there is no over provisioning of cloud resources.

**Figure 6–3 Self Service Portal Setup - Roles**

Infrastructure Cloud Self Service Setup Page Refreshed Feb 21, 2012 6:40:37 PM PST

**Roles**

View

| Role Name | No of Servers | VCPUs | Memory (GB) | Local Storage (GB) | Archive to Library                  | Save Server on Expiry               | Zones | Network |
|-----------|---------------|-------|-------------|--------------------|-------------------------------------|-------------------------------------|-------|---------|
|           | 50            | 70    | 100         | 1000               | <input checked="" type="checkbox"/> | <input checked="" type="checkbox"/> | Zone  | NW/NW   |

- In the Roles page, click **Assign Quota to Role** to define quota at the role level. The quota assigned to a role is available to all users with that role. For example, if the Number of Servers for the SSA\_USER\_ROLE1 is set to 5, all users with this role can have a maximum of 5 servers. The quota is the aggregate amount of resources that can be granted to each self service user belonging to a certain role. If the user has more than one role, his quota is calculated as the absolute maximum across the roles for each attribute value.
  - Select Roles:** Click the torch icon to select a role for which the mapping is to be defined. The list of all the SSA roles appears.
  - Select Zones:** Click the torch icon. All zones marked as Cloud Self Service Zones during zone creation are displayed here. Select a zone from the list and click **OK**.
  - Number of Servers:** Maximum number of servers that can be reserved at any time. This number is across the zones. For example if the quota for a user is 5. Then, the total number of instances from all zones is 5. So the user can have 3 instances in zone 1 and 2 instances in zone 2, or all 5 instances in the same zone.
  - Number of VCPUs:** Maximum number of VCPUs that can be allocated.

- **Memory:** Maximum amount of memory that can be allocated.
- **Local Disk:** Maximum amount of local storage that can be allocated.
- **Additional Storage:** Maximum amount of additional storage that can be allocated.
- **Allow Archiving to Software Library:** Indicate whether the machine images can be stored in the Software Library after expiry. You can also select this option to save the machine image as a template when the guest virtual machine is stopped.
- **Select Network Profiles:** Select the network profile to assigned to the selected role. All deployments are done in a specific network defined by the network profile.

---

**Note:** Network profiles that have been associated with networks in the selected zone are displayed. For more details on network profiles, see [Section 4.8, "Creating a Network Profile"](#)

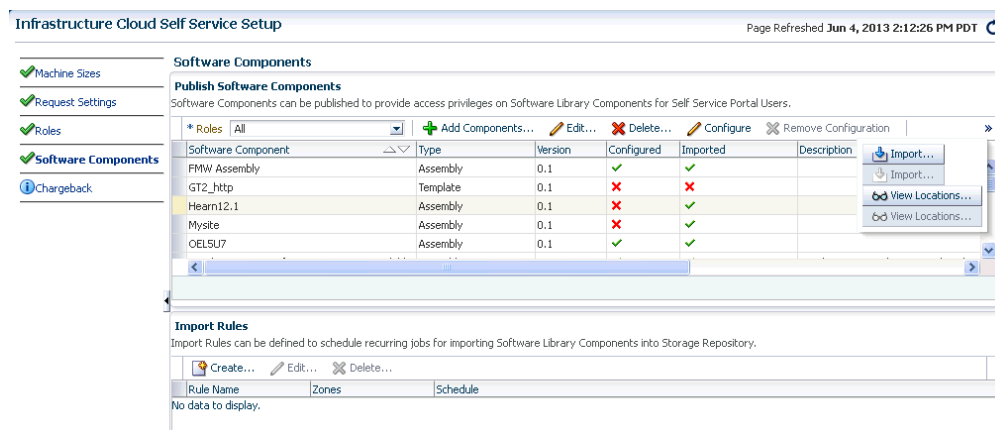
---

- **Override Global Request Settings:** Select this checkbox to override the request settings for the selected role. You can modify the Future Reservation, Request Duration, and configure the Management Agent.

Select a role and click **Edit**. You can modify the role mappings in the **Assign Quotas, Zones & Network Profiles to Role** page.

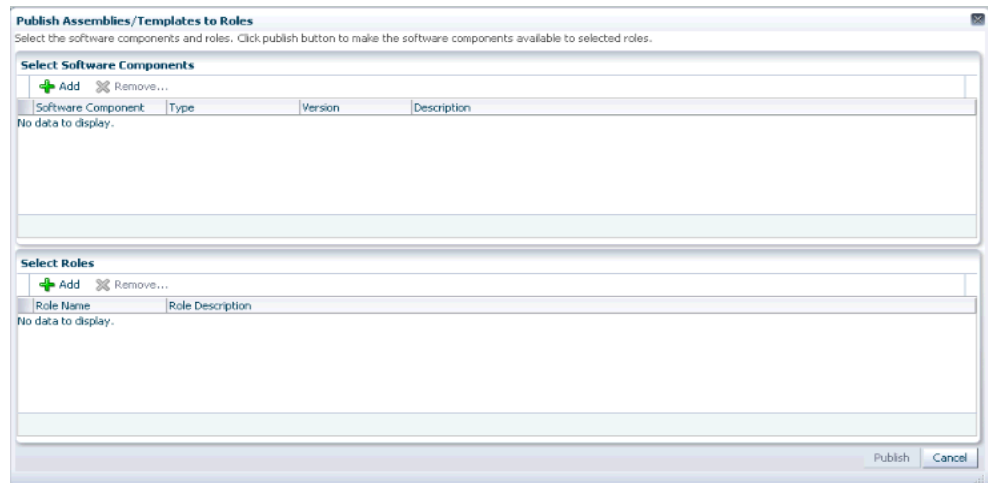
5. Click **Software Components** in the left panel to select the set of templates or assemblies that are accessible to the SSA users. You can publish the assemblies and templates to the SSA Portal and provide SSA users with access to these software components.

**Figure 6–4 Self Service Portal Setup - Software Components**



6. You can do the following on this page:
  - **Add Components:** Click **Add Components**. The Publish Assemblies / Templates to Roles page appears.



**Figure 6–5 Publish Assemblies / Templates to Roles**

On this page, do the following:

1. Click **Add**. Select the assembly or template from the list and click **Select**.
2. Click **Add** in the Select Roles section. Select the role from the list and click **Select**.
3. Click **Publish** to publish these software components to the SSA Portal. These software components are now available to the SSA users and can be deployed.

---

**Note:** You will see a list of assemblies or templates if they have been uploaded to the Software Library. See the Oracle Enterprise Manager Cloud Control Administration Guide for details on uploading components to the Software Library.

---

- **Import:** You can import the published assembly to a selected set of server pools and zones. Select the assembly and click **Import**. In the Select Zones page, select one or more zones and pools and click **Save** to import the assembly to the selected zones.
- **Configure:** Select a software component and click **Configure** to configure the Management Agent.

---

**Note:** The Management Agent can be deployed only if the **Enable Agent Configuration** check box has been selected in the Request Settings page.

---

The Management Agent configuration details are displayed:

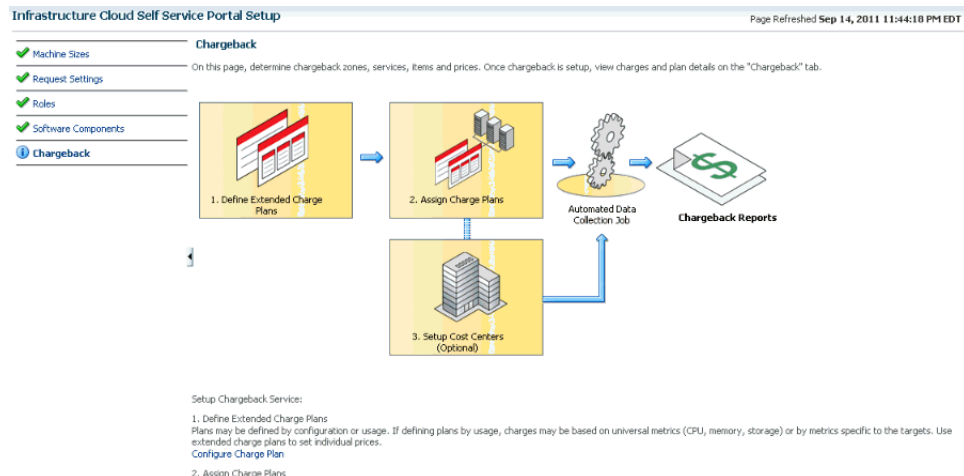
- **Installation Base Directory** (Mandatory): Specify the full path to the directory in which the Management Agent is to be installed. Ensure that the Installation Base Directory you specify in the response file is empty and has write permissions. This is a mandatory field.
- **Platform** (Mandatory): The platform for which you want to download the Management Agent software. This is a mandatory field.

- **Password (Mandatory):** Specify a password for registering the new Management Agents. By default, the communication between the Oracle Management Service and the Management Agent is secure and locked. Any new Management Agents that are registered must be authenticated. The password you specify here will be used to authenticate the new Management Agents.
- **Port (Mandatory):** Enter the port used for the Management Agent process. If no value is specified, then either 3872 or any free port between 1830 and 1849 will be used.
- **Additional Parameters:** You can add some options supported by agent installation. See the *Enterprise Manager Basic Installation Guide* for details.
- **SSH Timeout:** The timeout minutes for SSH.
- **Privilege Delegation Settings:** Run as root.

Specify the details and click **OK** to configure the Management Agent.

- **Import Rules:** You can define rules to import published assemblies and templates to server pools and zones on a periodic basis based on a predefined schedule. Click **Create** in the Import Rules section. Enter a Rule Name, select a set of zones and server pools, a schedule for the import and click **Save**. The published assembly or template is imported to the selected server pools and zones on a regular basis.
7. Click **Chargeback** tab in the left panel. The Chargeback application allows the SSA Administrator to define a charge plan and assign the plan to a specific zone or target.

**Figure 6–6 Chargeback**



8. Setting up the Chargeback Plan involves:

- **Defining Extended Charge Plans:** Click **Configure Charge Plan**. You can edit the Base Charge Plan or create an Extended Charge Plan to customize the charge plan for a specified target type.
- **Assigning Charge Plans:** Click **Configure Targets**. You can assign charge plans to a set of targets.

- **Setting Up the Cost Center:** Click **Configure Cost Center**. You can setup cost centers to aggregate costs among groups of users.

For more details on setting up the chargeback plan, see [Chapter 22, "Chargeback Administration"](#).

## 6.2 Uploading Large Files on the Self Service Portal

The EM\_SSA\_USER can upload software components to a library location. See [Section 7.1.5, "My Library"](#) for details. If the file to be uploaded is large, the EM\_SSA\_USER will see the following error message.

Error: The file is too large. Please add a smaller file,

To resolve this issue, you must modify the configuration settings in the web.xml file as follows:

1. Locate the web.xml configuration file by navigating to the following location in the Oracle Management Service Home

```
./oms/sysman/archives/emgc/deployments/EMGC_
DOMAIN/emgc.ear/em.war/WEB-INF/web.xml
```

If this directory structure is not present, navigate to the following location in the Oracle Management Service Home

```
./oms/sysman/archives/emgc/deployments/GCDomain/emgc.ear/em.war/WEB-INF
/web.xml
```

2. Open the web.xml and locate the following four lines. If these lines are not present in the web.xml, you must add them:

```
<context-param>

<param-name>org.apache.myfaces.trinidad.UPLOAD_MAX_DISK_
SPACE</param-name>

<param-value>40960000</param-value>

</context-param>

<param-name>org.apache.myfaces.trinidad.UPLOAD_MAX_FILE_
SIZE</param-name>

<param-value>40960000</param-value>

</context-param>
```

3. Increase the size of the <param\_value> parameter to a value larger than the size of the file to be uploaded.
4. Save, exit, and restart the Oracle Management Service.

If the file to be uploaded is huge, the page may expire while the upload is still in progress (usually after 10 minutes). To avoid this, the OMS property oracle.sysman.eml.maxInactiveTime must be set to a value (in seconds) greater than the expected time for the upload to complete. For example, set \$ORACLE\_HOME/bin/emctl set property -name oracle.sysman.eml.maxInactiveTime -value 3600

While the file upload is in progress, the uploaded chunks are written in a default location on the Oracle Management Service (OMS). This default location is usually a subdirectory of where the application is deployed on the WLS server. This location can be configured by setting the parameter org.apache.myfaces.trinidad.UPLOAD\_TEMP\_

DIR in the `web.xml` file to a custom location. If there are multiple users uploading files concurrently on the OMS, you must ensure that this location has enough disk space to accommodate all of this data. It is recommended that this location is on a different volume than the location on which the OMS is installed. This ensures that if volume runs out of space, the OMS performance is not impacted.

---

## Using the IaaS Self Service Portal

This chapter provides instructions on using the Self Service Portal for IaaS. It contains the following sections:

- [Using the Self Service Portal to Request Server Resources](#)
- [Requesting Servers for a Specific Duration](#)

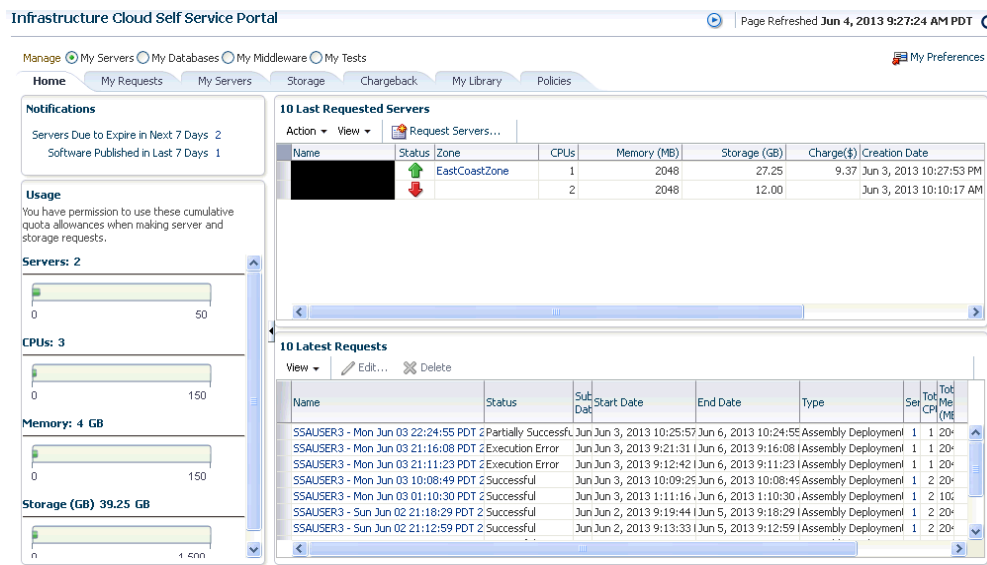
### 7.1 Using the Self Service Portal to Request Server Resources

The Self Service Portal for IaaS provides a dashboard for business users and cloud consumers to request resources. It allows SSA users to request and manage resources, view the metering and usage of these resources, and define policies to manage the workload and deployment automatically. The SSA Portal allows users to:

- Create new deployment requests
- Manage requested deployments
- Track metering and chargeback information
- Author and enforce cloud policies.

To use the Self Service Portal, follow these steps:

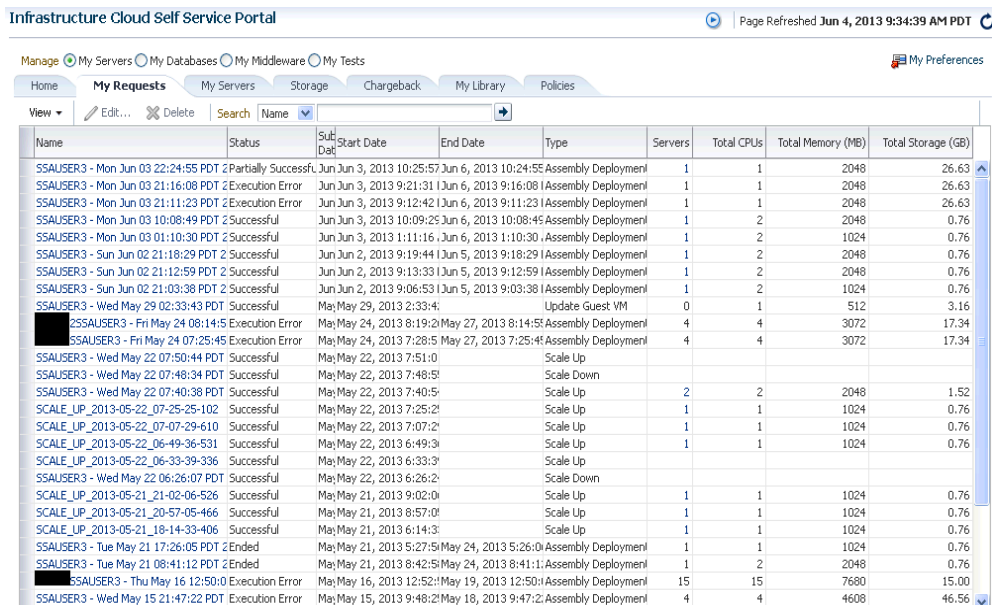
1. Log in as a user with the `EM_SSA_USER` role.
2. The Cloud Self Service Portal Home page appears. Click on the video link next to the Page Refreshed section to view a demo on using the self service portal. The following details are displayed:
  - **Notifications:** Lists notifications related to machine expiry and new assemblies, errors during requests, and quota used.
  - **Your Usage:** Lists the total quota allowances available for server and storage requests and quota that has been used.
  - **Servers Expiring Soon:** Lists the top 10 machines that you own in descending order of creation date. Click a Server Name to view the server details. See [Section 7.1.2.1, "Server Details Page"](#) for more information.
  - **Latest Requests:** This section lists the top 10 requests which are sorted as follows:
    - Descending order of the End Date.
    - Descending order of the Creation Date.Click a Request Name to view the request details.

**Figure 7–1 Self Service Portal**

- Click **Request Servers**. The New Server Request wizard appears. See [Section 7.2, "Requesting Servers for a Specific Duration"](#) for details.

### 7.1.1 Viewing My Requests

The My Requests tab lists all the requests made by the SSA user.

**Figure 7–2 My Requests**

The name of the request, status, start and end date, type of request, number of servers, total CPU, and memory appears. Click a Name to drill down to the Request Details page which shows complete details about the request.

## 7.1.2 Viewing My Servers

Click the **My Servers** tab to view a list of servers grouped according to the zone. On this page, the SSA user can view the specific assembly and template deployments, the tiers, and the servers in each tier.

**Figure 7–3 My Servers**

| Name          | Type              | Operating System | Server Name |
|---------------|-------------------|------------------|-------------|
| EastCoastZone | Zone              |                  |             |
| SIDBAgt1      | Assembly Instance |                  |             |
|               | Server            | None             |             |
| SandZone      | Zone              |                  |             |
| OELS7         | Assembly Instance |                  |             |
|               | Tier Instance     |                  |             |

You can drill down to the Assembly or Tier Instance Home page or click a server name to drill down to the Server Details page. From this page, you can perform the following actions:

- **Request Servers:** You can request or reserve servers for a specific duration. See [Section 7.2, "Requesting Servers for a Specific Duration"](#) for details.
- **Assemblies:** You can start, stop, delete an assembly instance, and add tiers. Select an assembly instance from the list and from the **Actions** menu, select an appropriate option.
- **Tiers:** You can start, stop, scale up, and scale down a tier instance.
- **Servers:** You can modify the configuration of the server, stop, start or restart a server, clone a server, save as template, suspend, resume, and launch a VNC Console.

### 7.1.2.1 Server Details Page

The Server Details page displays detailed information about the server. From this page, you can modify configuration, clone, stop, restart, suspend, resume, and launch VNC console. This page contains the following sections:

- **General:** Displays information like status of server, the associated zone, date from which it is available, and the expiry date.
- **Configuration:** Displays configuration details of the server such as number of CPUs, memory, storage, operating system, and other software installed on the server.
- **Availability:** Shows the availability status of the server over the last 24 hours.
- **Charts:** You can see the Availability chart that displays the status of the server for past 24 hours, and the CPU, Memory, and Storage Utilization charts.
- **Chargeback Details:** This section displays all the charge back break-up data for the server. The charge incurred will be calculated from the server requested date to the current date.

### 7.1.2.2 Scaling Up a Tier Instance

A scalable tier allows creation of multiple servers and has a minimum and maximum limit. The scalability of the tier depends on how it is defined in the Assembly. See [Section 8.8.4, "Provisioning Guest Virtual Machines Using Oracle Virtual Assemblies \(OVA\)"](#) for details. To scale up a tier instance, follow these steps:

1. Log in to Enterprise Manager as an user with the EM\_SSA\_USER role.
2. Click the **My Servers** tab, expand an assembly and click on a tier of the assembly.
3. In the Tier Instance Home page, click the **Scale Up** option from the Tier menu.
4. The Scale Up: General page appears. The name of the request, the assembly, zone, and number of instances appears. Click **Next** to continue.
5. In the Server Selection page, select the servers for the scale up operation. As a part of the Scale Up operation, either the existing shutdown servers of the tier are brought up and/or new servers are created to the Max Instance Number Limit if required.

If the Max Instance Number Limit of the tier is 4 and only 1 server has been created during the assembly deployment, the tier can be scaled up by 3 new servers. Click **Next** to continue.

6. If new servers are to be provisioned, you can modify their configuration in the Server Configuration page. The values are pre-filled with the values that specified during the deployment of the assembly. Modify the details if required and click **Next** to continue.

---

---

**Note:** If the tier is auto scalable, the Server Configuration page cannot be modified.

---

---

7. Specify the schedule for this operation and click **Next**.
8. Review the details and click **Finish** to scale up the tier.

### 7.1.2.3 Scaling Down a Tier Instance

You can scale down the servers belonging to a tier. To scale down a tier, follow these steps:

1. From the **Enterprise** menu, select **Infrastructure Cloud**, then select **Self Service Portal**.
2. Click the **My Servers** tab, expand an assembly and click on a tier of the assembly.
3. In the Tier Instance Home page, click the **Scale Down** option from the Tier menu.
4. In the Scale Down page, select the servers to be included in the Scale Down operation and click **OK**. The selected servers will be stopped and scaled down.

## 7.1.3 Viewing Storage

Click the **Storage** tab to view all the storage elements grouped by the Guest VMs with which they are associated. The Storage tab lists all the Servers owned by the EM\_SSA\_USER along with all the storage resources in tree format. The top level nodes represent the servers owned by the user and the sub nodes represent the storage disks that belong to the server. You can do the following:

- **Request Storage:** This action allows the user to add a shared or non-shared storage disk to a server. Select one or more servers and click **Request Storage**.



Enter the Disk Name, Size, and Mode. Select the Shared Disk check box to add a shared disk. If you create a non-shared storage disk, it is available only for the server on which it is created. A shared storage disk can be assigned to other servers. If you select multiple servers, you can create only non-shared storage disks.

- **Assign Storage:** This action allows the user to assign shared storage disk to any other server owned by the EM\_SSA\_USER. Select a server and click **Assign Storage**. The list of available shared storage disks are displayed. Select a disk from the list and click **OK** to assign the shared storage disk to the server.
- **Release Storage:** Select a server and click **Release Storage**. This action allows the user to release an existing shared or non-shared storage disks.

### 7.1.4 Viewing Chargeback Details

Click the **Chargeback** tab to view the chargeback details. The tab contains the following sections:

- **Charge Trend:** The Charge Trend shows the charges across a date range for various metrics. You can set a date range and click **Go** to recalculate charges.
- **Details:** This section shows charge details for the specified date range. You can reorder the data by selecting an option from the Detail Level drop-down menu. Click **Export** to export the details of the charge plan to a .csv file.
- **Charge Plans:** This section allows you to select a zone to see the charge plan in effect for the target type.
- **Oracle VM Guest (Default):** If you are using the Extended Charge Plan, the base plan for the Oracle Guest VM target type appears here.

### 7.1.5 My Library

Click the **My Library** tab to view the list of assemblies, templates, and deployment plans that you can access. You can choose to view:

- **Public:** The list published by the SSA Administrator for your role. The assemblies, templates and deployment plans on which you have view privileges. You will see the assemblies, templates, or deployment plans if the EM\_SSA\_ADMINISTRATOR has published them or if another EM\_SSA\_USER has made them shareable.
- **Private:** These include the assemblies, templates, and deployment plans that you own. Generally, an SSA user may not own assemblies but may own templates and deployment plans.

You can request servers from this page. See [Section 7.2, "Requesting Servers for a Specific Duration"](#) for details.

You can upload software, share, delete, or stop sharing a software component that you own. Select a component from the list and click **Share**. Select one or more users with whom the entity is to be shared and click **Share**.

To unshare an entity, select a shared entity from the list and click **Stop Sharing**. The list of users with whom this component has been shared appears. Select one or users from the list and click **Stop Sharing**.

To upload software, click **Upload**. In the dialog box that appears, select the source file to be upload, enter a name and description, and specify the type (Assembly or Template) of component being uploaded. Click **Upload** to upload to software to My Library.

## 7.1.6 Viewing Policies

Click the **Policies** tab to view the list of policies that you have defined. You can create a new Performance or Schedule Based Policy or a Policy Group. For example, you can define a policy to start an automatic scale up for a tier if the CPU Utilization is greater than 90 for 5 consecutive occurrences.

For more details on creating and using policies, see [Section 5.1, "Managing Cloud Policies"](#).

## 7.1.7 Viewing My Preferences

Click **Apply** to save the preferences. The values specified here will be used when you request resources using Infrastructure Self Service Portal.

Click **My Preferences**. The following page appears:

**Figure 7–4 My Preferences**

**General Preferences**

☒ **General Preferences**

Email Address

Locale

☐ Request Notifications

Request Status

☐ Instance Expiry Notifications

Days before first notification

Days before second notification

☐ Quota Notifications

Quota Threshold (%)

Default Service Type

☒ **My Servers Preferences**

Default Server Root Password

Confirm Default Server Root Password

Default Zone for New Requests

Default Source Software

☒ **My Database Preferences**

Default PaaS Infrastructure Zone

☒ **My Middleware Preferences**

Default PaaS Infrastructure Zone

☒ **My Tests Preferences**

Default Instance Root Password

Before you request resources using SSA, you can set up your user profile by specifying the following details:

- **General Preferences**
  - **Email Address:** The default email address to be used to receive notifications.
  - **Locale:** The default locale for receiving cloud notifications.
  - **Request Notifications:** Select this checkbox to receive notifications of different statuses of the request.

- **Request Status:** If you select the **Request Notifications** checkbox, select the status for each you wish to receive notifications from the drop down list.
- **Instance Expiry Notifications:** Select this checkbox to receive notifications of service instances that are due to expire based on the notification period.
- **Days Before the First / Second Notification:** If you select the Instance Expiry Notifications checkbox, specify the number of days before the first and second notifications are sent.
- **Quota Notification:** If you select this checkbox, notifications are sent when the quota reaches a specified threshold.
- **Quota Threshold (%):** If you have selected the **Quota Notification** checkbox, specify the quota threshold in this field.
- **Default Service Type:** Select the default service page that should appear when you log in to the Self Service Portal.
- **My Servers Preferences**
  - **Default Server Root Password:** The default root password that will be used when a machine is created.
  - **Default Zone for New Requests:** The default zone for new requests.
  - **Default Source Software:** The default location in which the assemblies, templates, and deployment plans are stored.
- **My Database Preferences**
  - **Default PaaS Infrastructure Zone:** Select the default PaaS Infrastructure Zone for database requests.
- **My Middleware Preferences**
  - **Default PaaS Infrastructure Zone:** Select the default PaaS Infrastructure Zone for middleware requests.
- **My Tests Preferences**
  - **Default Instance Root Password:** The default password for the new environment and test driver instance.

## 7.2 Requesting Servers for a Specific Duration

Follow this process if you need to request or reserve servers for a specific duration:

1. From the **Enterprise** menu, select **Infrastructure Cloud**, then select **Self Service Portal**.
2. On the Cloud Self Service Portal page, click **Request Servers** to create a new request for a virtual machine.
3. In the General page, enter a name for the request and select a zone in which the server is to be created.

Figure 7-5 New Server Request: General

**New Server Request : General**

Back Step 1 of 4 Next Save As Deployment Plan... Finish Cancel

**Request Name**  
Specify a name for your request. This will help you track it later.  
\* Name SSAUSER3 - Tue Jun 04 09:49:16 PDT 2013

**Destination**  
Select the zone in which servers will get created.  
\* Zone SandZone  
Description SandZone

**Source**  
Select the source software (assembly or template) to be used for this request.  
\* Source SIDB Assembly

**Assembly Instance Name**  
\* Name SIDB Assembly

**Deployment Plan(Optional)**  
Deployment plans allow you to record inputs and use them later while creating new requests using standard values.  
Name  
Description

4. Click the **Search** icon in the **Source** field. Select an assembly or template to be used for this request and enter the server name. For details on how to create an assembly component, see [Section 8.8.3.1.1, "Creating an Assembly Component"](#). For details on how to create a template component, see [Section 8.8.3.1.2, "Creating a Template Component"](#).
5. You can optionally select a Deployment Plan to pre-fill values in the deployment flow from an existing saved or imported deployment plan.
6. Click Next. The Deployment Configuration page appears.

Figure 7-6 New Server Request: Deployment Configuration

**New Server Request : Deployment Configuration**

Back Step 2 of 4 Next Save As Deployment Plan... Finish Cancel

Configure Networks

| Name                               | Number of Servers |         |         |         | Auto Scalable | Deploy                              | Server Size | Deployment Option | Server Name Prefix |
|------------------------------------|-------------------|---------|---------|---------|---------------|-------------------------------------|-------------|-------------------|--------------------|
|                                    | Default           | Minimum | Maximum | Initial |               |                                     |             |                   |                    |
| SIDB Assembly                      |                   |         |         |         |               |                                     |             |                   |                    |
| sidbA_linuxx64_11203psu1/sidb:SIDB | 1                 | 1       | 1       | 1       |               | <input checked="" type="checkbox"/> | Custom      |                   |                    |

Selected Row sidbA\_linuxx64\_11203psu1/sidb:SIDB Assembly

**sidbA\_linuxx64\_11203psu1/sidb:SIDB Assembly**

Deployment Option

**Server Configuration** Product Configuration

**General**

Enable High Availability ☐

\* Root Password \*\*\*\*\*

\* Confirm Root Password \*\*\*\*\*

Keymap en-us (English, United States)

**Server Size : Custom**

\* Maximum Memory(MB) 2048 \* Maximum Number of CPUs 1

\* Memory(MB) 2048 \* Number of CPUs 1

**Network**

View Add Edit Remove

| Name | Assembly Network | Quality Of Service | Backend Network | IP Assignment | Network Profile Name |
|------|------------------|--------------------|-----------------|---------------|----------------------|
| eth0 | network-1        | Any Network        | T System Assign | DHCP          |                      |

**Storage**

View Quality Of Service System Defined Custom

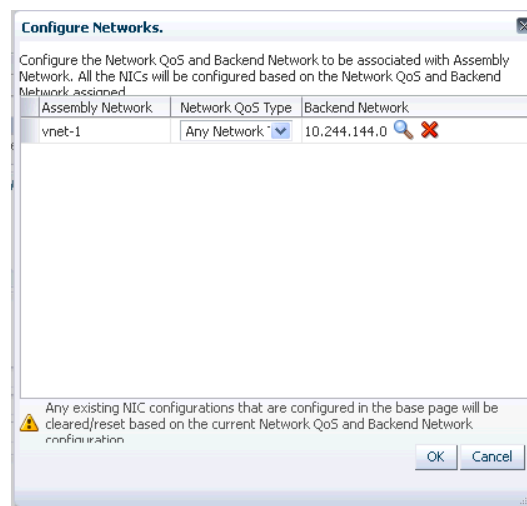
| Disk Name                 | Size (GB) | Mode      |
|---------------------------|-----------|-----------|
| nirYE45aVZYW_System       | 5.00      | ReadWrite |
| AB                        | 0.00      | ReadWrite |
| nirYE45aVZYW_sys-db_base  | 0.49      | ReadWrite |
| nirYE45aVZYW_sys-db_home  | 9.94      | ReadWrite |
| nirYE45aVZYW_sys-oradata  | 3.15      | ReadWrite |
| nirYE45aVZYW_sys-recovery | 8.05      | ReadWrite |

Total Size: 26.63 GB

7. The configuration properties defined for the assembly or template are displayed. Click on a tier that you wish to configure.

8. Click on the **Server Configuration** tab. The following regions can be configured:
  - **General**
    - Enable High Availability: If you want to enable high availability for this server, select this check box.
    - Root Password: The root password to access the server.
  - **Network:** The network interfaces for the server. Click Add to add a network interface card. In the Add a Network Interface Card, specify the name of the NIC card and click **OK**. To modify the QoS, Backend Network, and IP Assignment, click the **Configure Networks** in the top right hand corner of the page.

**Figure 7–7 SSA: Configure Networks**



On this page, you can configure the Network QoS and Backend Network to be associated with Assembly Network. All the NICs will be configured based on the Network QoS and Backend Network that have been assigned. Click **OK** to return to the previous page.

- **Server Size:** If the server size is large, medium, or small, you can modify the maximum memory and maximum number of CPUs that can be allocated to the server. If the server size is custom, you can specify the memory that can be allocated and the number of CPU cores to be assigned to the server.
  - **Storage:** The storage disks assigned to the server. This can be system defined or custom.
9. Click on the **Product Configuration** tab. This tab shows the properties of the products in the tier. You can choose to view All Properties or Required Properties. For each property, the default value and the current value is displayed. Select a property to modify the value if applicable.
  10. If the assembly has more than one server, you can To configure servers individually in the
  11. If a tier has more than one server, you can configure each server separately. Click the **Configure Individual Servers** link on the top right corner of the lower region. The Deployment Configuration: Configure Individual Servers page appears. You

can modify the configuration of the individual server and click **Continue** to return to the New Server Request: Deployment Configuration page.

12. Specify the schedule for the request and click **Next**.

---

**Note:** The schedule for the new server request is dependent on the schedule defined by the EM\_SSA\_ADMIN in the Self Service Portal Setup pages. See [Section 6.1, "Setting Up the Self Service Portal"](#)

---

13. In the Review page, you can click **Finish** to submit the request or **Save as Deployment Plan**. If you select the **Save as Deployment Plan** option, the request will be saved as a Deployment Plan that can be used to record inputs that can be used to create new request.

---

# Administering and Monitoring an IaaS Cloud

This chapter describes the monitoring and administrative tasks that you can perform on the zones, virtual server pools, virtual servers and guest virtual machines within an IaaS Cloud. It includes the following sections:

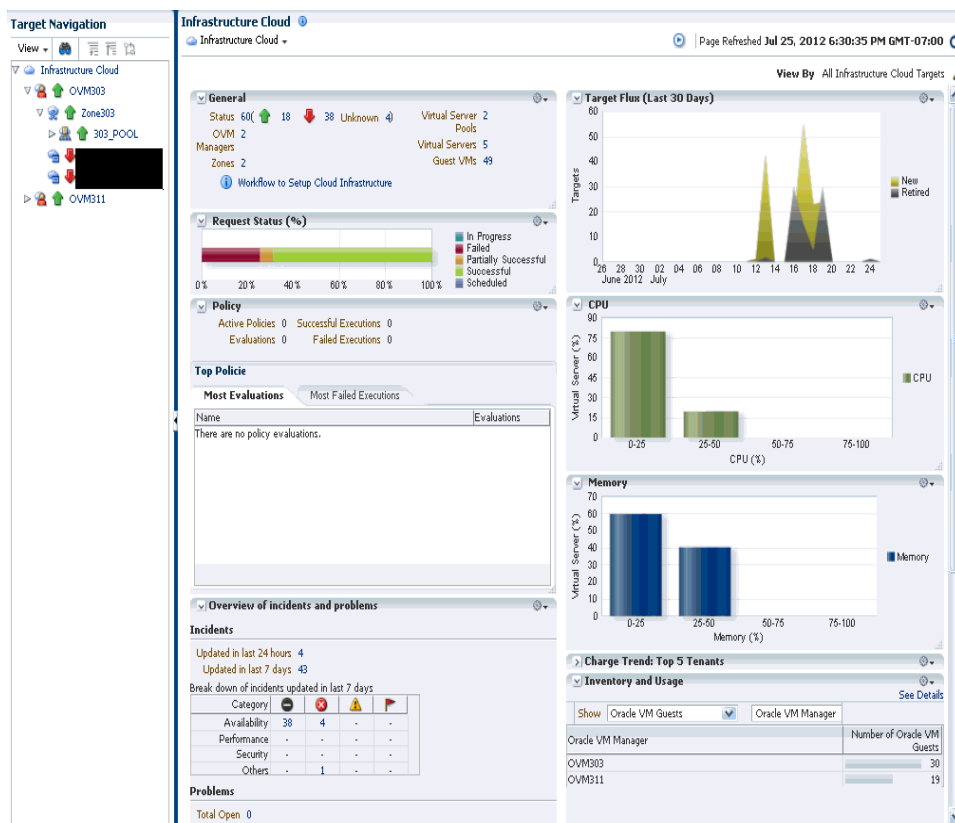
- [Viewing the Infrastructure Cloud Home Page](#)
- [Viewing the OVM Manager Home Page](#)
- [Viewing and Managing Members](#)
- [Managing Zones](#)
- [Viewing the Virtual Server Pool Home Page](#)
- [Managing Virtual Servers](#)
- [Managing Guest Virtual Machines](#)
- [Deploying Guest Virtual Machines](#)
- [Viewing the Infrastructure Request Dashboard](#)

## 8.1 Viewing the Infrastructure Cloud Home Page

The Cloud target is available out of the box and represents the entire virtualization infrastructure monitored by Enterprise Manager Cloud Control. You can view and monitor the various targets in the data center from the Infrastructure Cloud Home page. To view the Infrastructure Cloud Home page, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**. The Infrastructure Cloud Home page shows the various targets such as zones, pools, and virtual servers, request status, policies, and so on. Click the video link to view a demo on setting up the cloud infrastructure.

Figure 8–1 Infrastructure Cloud Home



The Infrastructure Cloud Home page shows the following details:

- **General:** Shows the number of OVM Managers, zones, and servers, and guest virtual machines in the cloud and the status of each target. Select a link to drill down to the Members page.
- **Target Flux:** Shows the number of targets (Guest VMs, Oracle VM Managers, Zones, server pools, virtual servers) created and deleted over the last 30 days.
- **Request Status:** Shows the status of various requests and reservations. Click the graph to drill down to the Requests page.
- **CPU & Memory:** Shows the CPU and memory utilization by the virtual servers in the data center.
- **Software Library:** Shows the number of assemblies and templates available in the software library and how many have been deployed.
- **Policy:** Shows the status of the various policies that have been defined for the cloud. It shows the number of corrective actions that were executed and whether they were successful. Click the graph to drill down to the Policy Home page.
- **Overview of Incidents and Problems:** Shows the various incidents and policy violations that have occurred in the last 24 hours.
- **Charge Trend:** Shows the charge back break-up data across the cloud. It lists the top 5 requests that have incurred the highest cost.
- **Inventory and Usage:** Shows the inventory summaries for virtual servers, guest virtual machines, and assembly instances. Select an option from the Show drop



down list to view the details. You can view the summary information in the context of different dimensions. For example, for Oracle VM Servers, you can view the number of virtual servers under each Oracle VM Manager, zone, virtual server pool, and so on.

Click on the bar charts to drill down multiple levels of inventory details. Click **See Details** to drill down to the Inventory and Usage Details page which provides detailed information on the selected target and associated inventory charts.


Before you can start monitoring a virtualization target, you must register the OVM Manager target. See [Section 4.3, "Registering the Oracle VM Manager"](#) for details.

## 8.2 Viewing the OVM Manager Home Page

The OVM Manager provides the interface to monitor zones, virtual servers, virtual server pools, and guest virtual machines. A zone is a set of virtual server pools. Each virtual server pool contains one or more virtual servers and guest virtual machines. A virtual server pool cannot belong to more than one zone. To view the OVM Manager Home page, follow these steps:

1. From the Enterprise menu, select **Cloud**, then select **Infrastructure Home**.
2. Click on an OVM Manager target from the left panel. The OVM Manager Home page shows the following details:

**Figure 8–2 OVM Manager Home**



- **General:** Shows the number of zones, and servers, and guest virtual machines under the OVM Manager and the status of each target. Click on a link to drill down to the Members page.
- **Request Status:** Shows the status of all requests for this OVM Manager. Click on the graph to drill down to the All Requests page.
- **Job Activity:** Shows the list of jobs started within the last 7 days.
- **Policy:** Shows the status of all the policies that have been defined for the OVM Manager. It shows the number of corrective actions that were executed and whether they were successful.
- **Overview of Incidents and Problems:** Shows all the incidents and problems reported on the OVM Manager and member targets.
- **Inventory and Usage:** Shows the inventory summaries for virtual servers, guest virtual machines, and assembly instances. Select an option from the Show drop down list to view the details. You can view the summary information in the context of different dimensions. For example, for Oracle VM Servers, you can view the number of virtual servers under each Oracle VM Manager, zone, virtual server pool, and so on.

Click on the bar charts to drill down multiple levels of inventory details. Click **See Details** to drill down to the Inventory and Usage Details page which provides detailed information on the selected target and associated inventory charts.

- **CPU & Memory:** Shows the CPU and memory utilization by the virtual servers under the OVM Manager.
  - **Software Library:** Shows the templates and assemblies that have been deployed.
  - **Charge Trend: Top 5 Tenants:** Shows the charge back break-up data across the OVM Manager. It lists the top 5 requests that have incurred the highest cost.
3. From the OVM Manager Home page, you can do the following:
- Edit OVM Manager
  - Synchronize OVM Manager
  - Deregister OVM Manager
  - Create a Zone
  - Create a Virtual Server Pool
  - Discover a Virtual Server
  - Manage Unowned Virtual Server
  - Manage Networks
  - Manage Storage
  - Manage Storage Repository
  - Manage YUM Repository
  - Manage NTP Configurations

## 8.3 Viewing and Managing Members

You can view and manage members for one or more virtualization targets on this page.

1. From the **Enterprise** menu select **Cloud**, then select **Infrastructure Home**.
2. Right click on the **Infrastructure Cloud** menu, select **Members**, then select **Show All**.

| Target Name     | Status | Type        | Operating System | Memory (MB) | CPU | Utilization |
|-----------------|--------|-------------|------------------|-------------|-----|-------------|
| Exalogic_OVMM   |        | OVM Manager |                  |             |     |             |
| OVM Manager     |        | OVM Manager |                  |             |     |             |
| OVM311_Ken      |        | OVM Manager |                  |             |     |             |
| OVM36069_wentao |        | OVM Manager |                  |             |     |             |
| OVM_slc01qps    |        | OVM Manager |                  |             |     |             |

3. The list of OVM Managers, Zones, Virtual Server Pools, Virtual Servers, and Guest Virtual Machines that can be managed from the Infrastructure Cloud appear.
4. You can:
  1. Click on a target to drill down to its Home page.
  2. Select a target and click on an option from the **Actions** menu. The actions you can perform vary based on the target you have selected.

---

**Note:** You can view and manage the members under a OVM Manager target, zone, virtual server pool, or virtual server. From the Infrastructure Cloud Home page, select the appropriate target (such as an OVM Manager, zone, virtual server pool, or virtual server) from the left panel. From the Infrastructure Cloud menu, select **Members**, then select **Show All**. The Members page for the selected target appears and you can perform various administrative activities on these targets.

---

## 8.4 Managing Zones

This section describes the following:

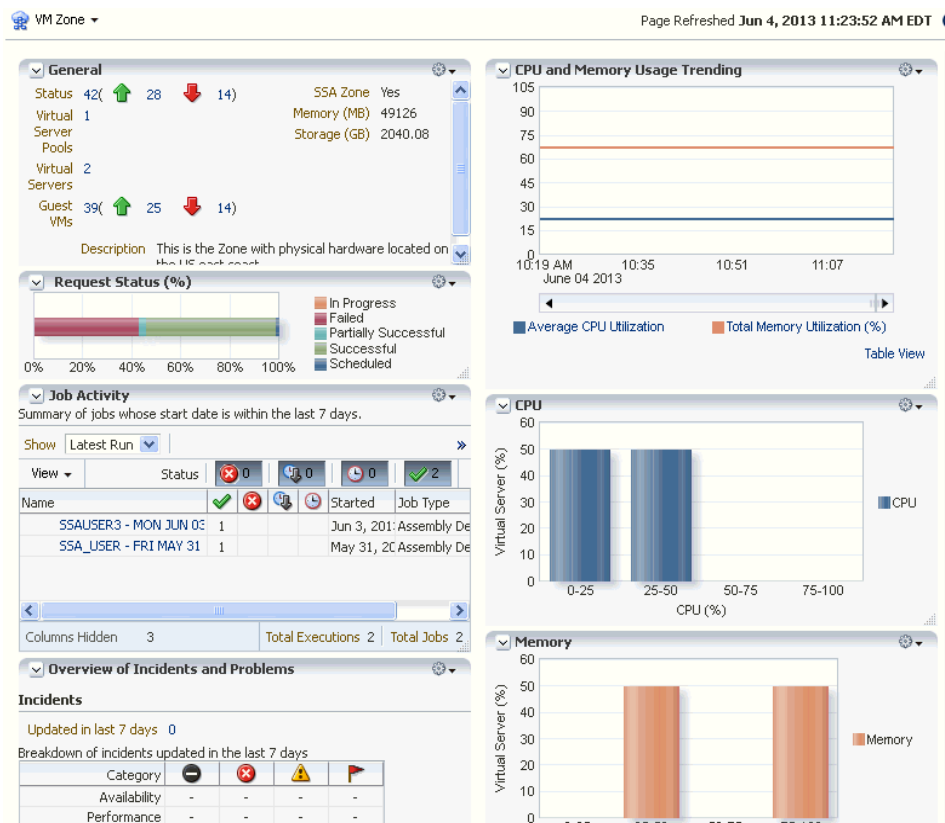
- [Viewing the Zone Home Page](#)
- [Creating or Editing a Zone](#)
- [Deleting a Zone](#)

### 8.4.1 Viewing the Zone Home Page

A zone is a logical grouping of resources like servers and storage entities. It is a logical entity used to organize the guest virtual machines exposed to the self service application user. Zones have metrics that show the aggregate utilization of the zone. To view the Zone Home page, follow these steps:

1. From the **Enterprise** menu select **Cloud**, then select **Infrastructure Home**. In the left panel, click on the zone to be viewed. The following details are displayed:

Figure 8–3 Zone Home



- **General:** Shows the number of virtual servers, virtual server pools, and guest virtual machines under this zone and the status of each target. Click on a link to drill down to the Members page.
- **Job Activity:** Shows the list of jobs started within the last 7 days.
- **Overview of Incidents and Problems:** Shows all the incidents and problems reported on the OVM Manager and member targets.
- **CPU and Memory Usage Trending:** Shows the average CPU and memory utilization over the last 30 days.
- **CPU and Memory:** Shows the CPU and memory utilization by the virtual servers and guest virtual machines in this zone.

## 8.4.2 Creating or Editing a Zone

A zone is a collection of virtual server pools which share storage. Follow these steps to create a zone:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Select the OVM Manager target for which the zone is to be created, and click **Create Zone** from the OVM Manager menu.
3. In the Create Zone page, specify the zone name and description.
4. Indicate if this zone will be used by the self service users by checking the **Infrastructure Cloud Self Service Zone** check box.

5. Click **Add** in the Virtual Server Pool section to add one or more virtual server pools. See [Section 4.14, "Creating a Virtual Server Pool"](#) for details.
6. Click **Create Zone** and click **OK**. After the successful completion of this job, the virtual server pools appear in the newly created zone on the Virtualization Central Home page.

---

**Note:** To edit an existing zone, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
  2. Select a zone target from the left panel.
  3. From the **Zone** menu, select **Target Setup**, then select **Edit VM Zone**.
- 

### 8.4.3 Deleting a Zone

To delete a zone, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Click on a zone target from the left panel to navigate to the Zone Home page. From the **Zone** menu, select **Target Setup**, then select **Delete**. A confirmation message appears.
3. Click **OK** to delete the zone. Any virtual server pools that are members of this zone will now be moved under the OVM Manager target.

---

**Note:** If Chargeback has been enabled for the zone (or virtual server pools / guest virtual machines within the zone) being deleted, the metric data is automatically disabled when the target is deleted.

---

### 8.4.4 Deleting a Zone from Enterprise Manager

To delete a zone from Enterprise Manager, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Click on a zone target from the left panel to navigate to the Zone Home page. From the **Zone** menu, select **Target Setup**, then select **Delete from Enterprise Manager**. A confirmation message appears.
3. Click **OK** to delete the zone. The zone and the associated targets under this zone will be removed from Enterprise Manager. After synchronization, these targets will be added back to Enterprise Manager.

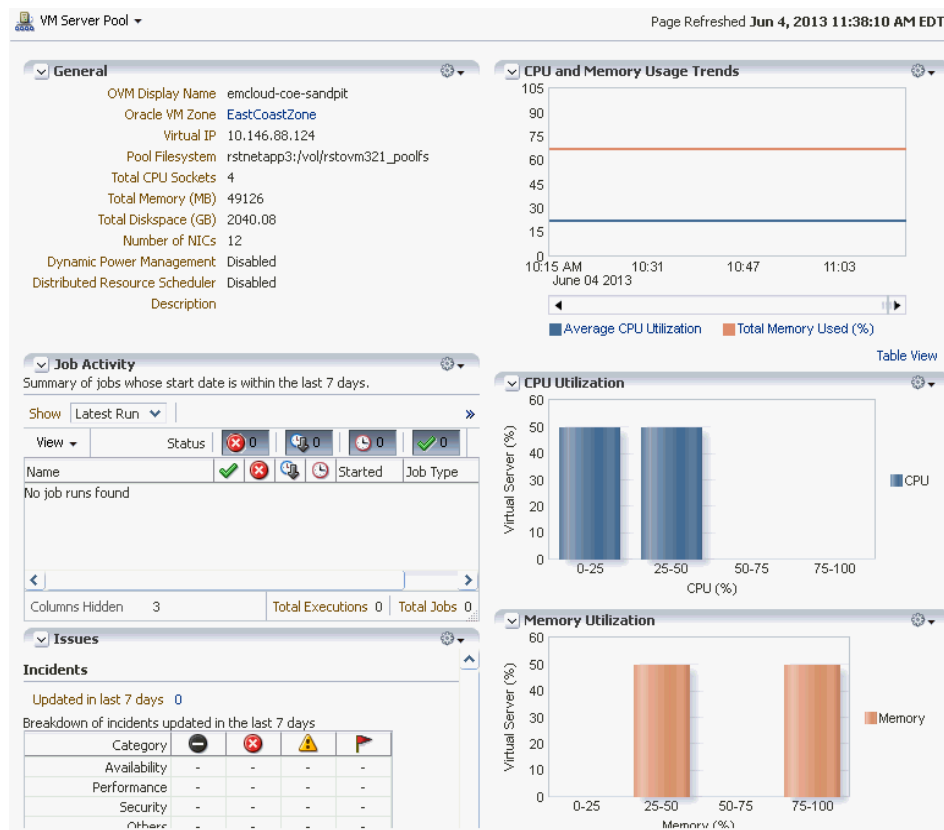
---

**Note:** If Chargeback has been enabled for the zone (or virtual server pools / guest virtual machines within the zone) being deleted, the metric data is automatically disabled when the target is deleted.

---

## 8.5 Viewing the Virtual Server Pool Home Page

A virtual server pool is logically an autonomous region that contains one, or more virtual servers. A virtual server pool may or may not belong to a zone. However, it can belong to only one zone at a time. The following details are displayed:

**Figure 8–4 Virtual Server Pool Home**

- **General:** Shows details of the virtual server pool such as the zone to which it belongs, the number of CPUs, memory and storage, and number of NICs.
- **Job Activity:** Shows the list of jobs started within the last 7 days.
- **Issues:** Shows all the incidents and problems reported on the virtual server pool and member targets.
- **CPU and Memory Usage Trends:** Shows the average CPU and memory utilization by the virtual servers in the zone over the last 30 days.
- **CPU and Memory Utilization:** These charts show the CPU and memory utilization by all the virtual servers in the server pool.

## 8.6 Managing Virtual Servers

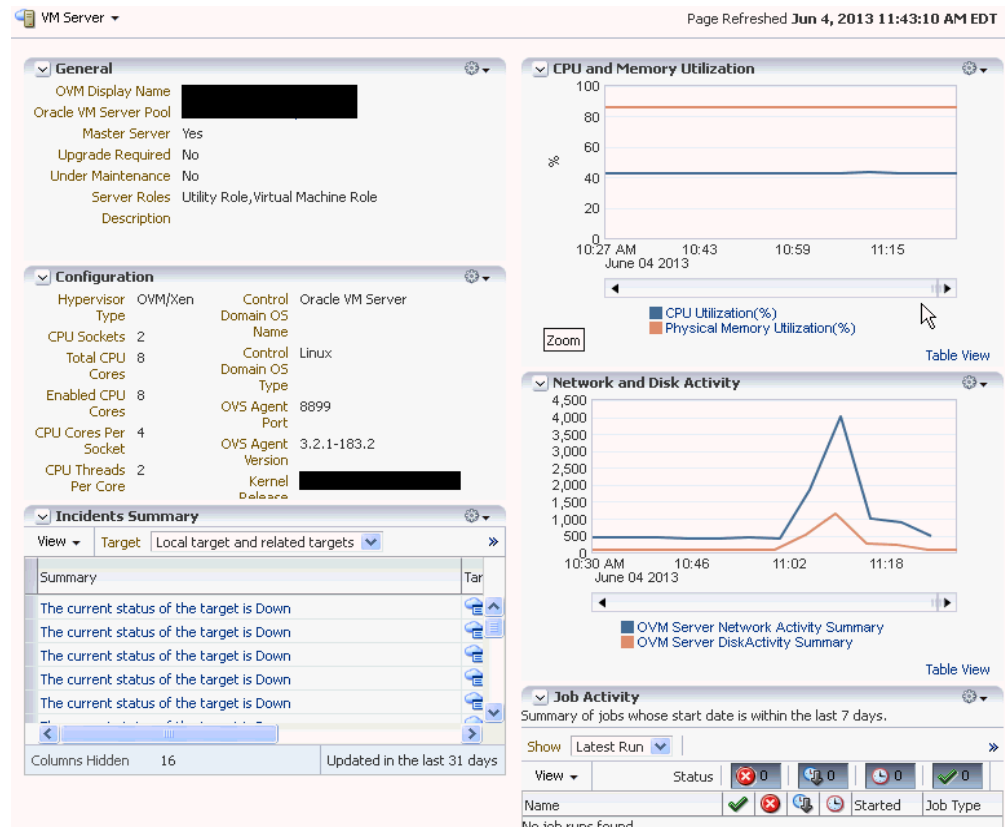
This section describes the following:

- [Virtual Server Home Page](#)
- [Editing a Virtual Server](#)
- [Upgrading a Virtual Server](#)
- [Starting and Stopping a Virtual Server](#)
- [Performing Maintenance on the Virtual Server](#)

## 8.6.1 Virtual Server Home Page

A virtual server is a generic term used to describe a physical machine which has virtualization software (hypervisor) running on it. A virtual server can belong to one and only one virtual server pool at a time. Guest virtual machines and resources are also associated with the server pools.

**Figure 8–5 Virtual Server Home**



The following regions are displayed:

- **General:** Shows details of the virtual server such as the virtual server pool it belongs to, whether the server needs to be upgraded, if it is under maintenance, and so on.
- **Job Activity:** Shows the list of jobs started within the last 7 days.
- **Configuration:** The virtual server configuration such as number of enabled CPU cores and processors, memory, disk space, OVS Agent Port, and so on.
- **Network and Disk Activity:** The network and disk activity of the virtual server.
- **CPU and Memory Utilization:** The CPU and memory utilization of the virtual server.
- **Incident Summary:** This region shows a list of incidents and problems reported on all the Guest VMs and related targets.

## 8.6.2 Editing a Virtual Server

To edit a virtual server, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. In the left panel, select the virtual server that is to be edited. The Virtual Server Home page appears.
3. From the **Virtual Server** menu, select **Target Setup**, then select **Edit VM Server**. You can modify the description of the virtual server and the following fields:
  - **Configure IPMI:** Select the **Enable Server IPMI** checkbox to enable the Intelligent Platform Management Interface (IPMI). Selecting this option allows you to remotely power off a virtual server and send a wake on lan message to power on a virtual server without having to physically press the power button. If this checkbox is selected, you must enter the user name, password, and IP address of the IPMI.
  - **Server Role:** A server can be designated to perform utility or virtual machine functions by specifying the role which can be:
    - **Utility Role:** Servers with this role are used to perform non-virtual machine operations such as cloning, importing templates, creating repository, deleting assemblies and so on.
    - **Virtual Machine Role:** Servers with this role are used to run virtual machines. This role cannot be deselected if the virtual server already contains Guest VMs.
4. Click **OK** to confirm the changes.

### 8.6.3 Upgrading a Virtual Server

After the YUM Repository has been set up (see [Section 4.17.1, "Configuring the YUM Repository"](#)), you must upgrade the virtual server to ensure that it has the latest updates. The virtual server must be in maintenance mode if it is to be upgraded.

To upgrade the virtual server, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Right click on a VM Server from the left panel and select **Upgrade**.
3. After the virtual server has been upgraded, the virtual server will be restarted.

### 8.6.4 Starting and Stopping a Virtual Server

You can start a virtual server that is in a **Stopped** state. To start a virtual server, follow these steps:

---

---

**Note:** Before you start the virtual server, you must ensure that the IPMI has been configured. See [Section 8.6.2, "Editing a Virtual Server"](#).

---

---

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Right click on a virtual server target in the left panel and select **Start**. A confirmation message appears.
3. Click **OK** to start the virtual server.

You can stop a virtual server that is **Up**. To stop a virtual server, follow these steps:

1. Log in to Enterprise Manager Cloud Control.
2. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.



3. Right click on a virtual server target in the left panel and select **Stop**. A confirmation message appears.
4. Click **OK** to stop the virtual server.

### 8.6.5 Performing Maintenance on the Virtual Server

You can place a virtual server in maintenance mode to perform hardware and software maintenance. When a virtual server is placed in maintenance mode, any virtual machines running on the virtual server are automatically migrated to other virtual servers in the server pool, if they are available, otherwise they are stopped. The server pool roles being performed by the virtual server are also moved to other virtual servers in the server pool, if available.

To place an Oracle VM Server into maintenance mode:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Select a virtual server target from left panel.
3. Right click on a virtual server target from the left panel and select the **Start Maintenance**. The virtual server target is placed in blackout mode during maintenance.

When you have finished performing maintenance, click **Stop Maintenance** for it to rejoin the server pool.

## 8.7 Managing Guest Virtual Machines

This section describes the following:

- [Creating a Guest Virtual Machine](#)
- [Guest Virtual Machine Home Page](#)
- [Migrating a Guest Virtual Machine](#)
- [Starting and Shutting Down a Guest Virtual Machine](#)
- [Cloning a Guest Virtual Machine](#)
- [Editing a Guest Virtual Machine](#)
- [Saving a Guest Virtual Machine as a Template](#)

### 8.7.1 Creating a Guest Virtual Machine

A guest virtual machine runs on a virtual server. You can create a guest virtual machine by using one of the following methods:

- Assemblies
- Templates
- Installation Media (ISO or PXE)

For more details on creating a Guest VMs, see [Section 8.8, "Deploying Guest Virtual Machines"](#).

Before you create a guest virtual machine, you must have:

- Discovered a virtual server.
- Created a virtual server pool.

- Created the necessary components in the Software Library
- Deployed assemblies, templates, ISO images, or PXEs for creating virtual machines based on these.

## 8.7.2 Guest Virtual Machine Home Page

A guest virtual machine is the container running on top of a virtual server. Multiple guest virtual machines can run on a single virtual server. The following regions are displayed:

- **General:** Shows details of the guest virtual machine such as the virtual server pool, and virtual server it belongs to, and a description.
- **Networks:** Displays the network configuration for the guest virtual machine.
- **CPU and Memory Utilization:** The CPU and memory utilization of the guest virtual machine.
- **Networks and Network Activity:** The networks in the guest virtual machine and the network throughput of the guest virtual machine.
- **Configuration:** The virtual server configuration such as number of virtual CPUs, CPU priority, allocated storage and memory, whether it is HA enabled, and the domain type.
- **Job Activity:** Shows the list of jobs started within the last 7 days.
- **Incidents Summary:** Shows the various alerts and policy violations that have occurred in the last 31 days.

From the Guest Virtual Machine Home page, you can do the following:

- Edit a Guest VM
- Migrate a Guest VM
- Clone a Guest VM
- Start and Stop
- Reboot, Restart, Suspend, and Resume

## 8.7.3 Migrating a Guest Virtual Machine

To migrate one or more guest virtual machines, do the following

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. From the **Infrastructure Cloud** menu, select **Members**, then select **Show All**.
3. Expand the list of targets displayed and select one or more guest virtual machines from the list. From the **Actions** menu, select **Guest VMs**, then select **Migrate**. The Migrate page appears. You can select the following:
  - Running Guest VMs
  - Stopped / Suspended Guest VMs.
4. Select one or more Guest VMs from these regions, and click **Submit**.

---

**Note:** For running Guest VMs, the virtual server to which the guest virtual machine is to be migrated is auto selected by default. For stopped or suspended Guest VMs, the target server can be modified if required.

---

## 8.7.4 Starting and Shutting Down a Guest Virtual Machine

After a guest virtual machine has been created, you can start it or shut it down. When a guest virtual machine is running, you can pause or suspend it temporarily:

- [Starting a Guest Virtual Machine](#)
- [Restarting a Guest Virtual Machine](#)
- [Stopping a Guest Virtual Machine](#)

### 8.7.4.1 Starting a Guest Virtual Machine

When a Guest VM has not been used for a while, it may be shut down to release system resources. To use the Guest VM, you must start it:

#### Prerequisites

The guest virtual machine must be **Down**.

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Select the guest virtual machine target from the left panel.
3. Click **Start** from the guest virtual machine menu.

### 8.7.4.2 Restarting a Guest Virtual Machine

You may need to reboot or restart a virtual machine if operating system updates require you to restart the virtual machine, for example Microsoft Windows updates. To reboot a guest virtual machine:

#### Prerequisites

The guest virtual machine must be **Up**.

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Right click on a guest virtual machine target from the left panel and click **Restart**.

### 8.7.4.3 Stopping a Guest Virtual Machine

You may need to stop a Guest VM that is not being used to release system resources.

#### Prerequisites

The guest virtual machine must be **Up**.

To stop a guest virtual machine:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Right click on a guest virtual machine target from the left panel and click **Stop**.

### 8.7.4.4 Killing a Guest Virtual Machine

Killing a virtual machine is equivalent to performing a power off of a virtual machine, similar to unplugging the power cable from a physical machine. This is not the

recommended method of shutting down a virtual machine, but may be used if the **Stop** command fails to shut down the virtual machine.

To kill a guest virtual machine:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Right click on a guest virtual machine target from the left panel and click **Kill**.

## 8.7.5 Cloning a Guest Virtual Machine

### Prerequisites

- At least one virtual machine must exist and be in a state appropriate for cloning to proceed.
- You can clone a guest virtual machine to create one or more new virtual machines in the same server pool.

---

---

**Note:** If you are using the Enterprise Manager for Oracle Virtualization 12.1.0.3 plug-in or above, you can clone running guest virtual machines if the virtual disks associated with them are thin cloneable.

---

---

To clone a guest virtual machine, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Select a guest virtual machine target from the left panel.
3. Click **Clone** from the guest virtual machine menu.
4. In the Select Target and Source page, specify the Request Name.
5. In the Deployment Configuration page, you can modify the following:
  - **Guest VM:**
    - **Number of Guest VMs:** The number of guest virtual machine instances being cloned.
    - **Guest VM Naming:** Indicates whether the guest virtual machine names are to be automatically generated or if the names will be uploaded from a file.
    - **Guest VM Name:** Enter a name of the guest virtual machine being cloned.
  - **General:**
    - **Enable High Availability / Start VM After Creation:** Select the **Enable High Availability** and **Start VM After Creation** checkboxes to restart the guest virtual machine automatically after any outage.
    - **CPU Scheduling Priority:** Set the Priority of the virtual machine. You can select a high, medium, low, or custom priority for the virtual CPUs.
    - **CPU Scheduling CAP:** This parameter defines the percentage to which the virtual CPUs can receive scheduled time. You can select a high, medium, low, or a custom percentage of scheduled time for the virtual CPUs.
    - **Keymap:** The keyboard mapping that will be used by the guest virtual machine being cloned.

- **Tags:** Specify one or more comma separated tags. These tags are used to logical group the virtual machines and can be used to search for one or more virtual machines that meet a certain criteria.
  - **VM Size**
    - **VM Size:** The amount of memory to be allocated to the guest virtual machine. The virtual server should have sufficient memory to run the guest virtual machine. Otherwise, the guest virtual machine will be created, but will not be started. The size of the guest virtual machine (memory and CPU values).
    - **Memory:** Specify the memory allocated to guest virtual machine. For a running guest virtual machine, this can be increased up to the Max Memory.
    - **CPU Cores:** Specify the number of virtual CPUs for the guest virtual machine. For a running guest virtual machine, this can be increased upto the Max virtual CPU.
  - **Disk Configuration:** You can configure the storage type to specify the location of the disks of the new cloned Guest VM. You can use the Default type as specified by the server pool or select Override. If you select Override, you must specify the Storage QoS and click the Search icon in the Storage field to specify the Storage Type. In the Storage Type window, specify the following:
    - **Clone Target Type:** This can be Repository or Storage Array.
    - **Clone Target:** If you select Repository, select the storage location where the disk is to be created. If you select Storage Array, specify the type, which can be SAN or iSCSI storage server.
    - **Clone Type:** This can be Sparse or Non-Sparse Copy. A sparse copy is a disk image file of a physical disk, taking up only the amount of space actually in use; not the full specified disk size. A non-sparse copy is a disk image file of a physical disk, taking up the space equivalent to the full specified disk size, including empty blocks.
  - Apart from specifying the storage type, you can also add new disks for the Guest VM to be cloned. Click **Add New Disk** to create a new disk in repositories and connect it to the target Guest VM. Click **Attach Disk** to provide the existing disk in the repositories and connect it to the Guest VM.
6. Click **Next** to continue.
  7. On the Schedule page, specify whether the guest virtual machine is to be cloned immediately or later. Click **Next**.
  8. On the **Review** page, review the configuration details and click **Submit** to run the cloning procedure.

## 8.7.6 Editing a Guest Virtual Machine

### Note:

- If you edit a guest virtual machine that is **Up**, you cannot modify the server size, storage and network fields. To edit these values, you must stop or halt the guest virtual machine.
- You can create a new disk for a guest virtual machine which is **Down**, if Oracle VM Manager 3.1.1 and Enterprise Manager for Oracle Virtualization 12.1.0.2 plug-in have been installed. In other cases, existing disks are attached to the guest virtual machines.

To modify the configuration of a guest virtual machine, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Select a guest virtual machine target from the left panel. From the **Guest VM** menu, select **Target Setup**, then select **Edit VM Guest**. The Guest VM Configuration page appears.

**Figure 8–6 Edit Guest Virtual Machine: Hardware Configuration**

VM Guest Page Refreshed Aug 1, 2012 10:41:55 PM GMT

**Hardware Configuration** Other Configuration

**General**

Configure CPU and Memory

\* VM Size Custom

Max. Memory (MB) 1536

\* Memory (MB) 1536

Max. Virtual CPUs 1

\* Virtual CPUs 1

CPU Cap 10 40 70 100 Custom

☒ Oracle VM Server will limit the allocation of physical CPU cycles, based on the CPU Cap of Guest VM

CPU Scheduling Priority 0 25 50 75 100 Custom

☒ Oracle VM Server will schedule the physical CPU cycles, based on the CPU Schedule Priority of the Guest VM

**Network**

View Add Delete...

| MAC Address       | Network                    | Network Type |
|-------------------|----------------------------|--------------|
| 00:21:f6:4d:70:0e | cloud_network_setup_ovm303 |              |

Available Backend Network 1

**Storage**

View Add Delete...

| Disk Name                | Size (MB) | Storage            | Storage Type |
|--------------------------|-----------|--------------------|--------------|
| JSnLjyo5fsvnQA2_System ( | 5122      | cloud_repo_ovm_303 | Repository   |
| AB (5)                   | 219       | cloud_repo_ovm_303 | Repository   |
| JSnLjyo5fsvnQA2_sys-Java | 439       | cloud_repo_ovm_303 | Repository   |
| JSnLjyo5fsvnQA2_usr-Midd | 6565      | cloud_repo_ovm_303 | Repository   |

3. In the Hardware Configuration page, you can modify the following details:
  - **General:**
    - **VM Size:** The amount of memory to be allocated to the guest virtual machine. The virtual server should have sufficient memory to run the guest virtual machine. Otherwise, the guest virtual machine will be created, but will not be started. The size of the guest virtual machine (memory and CPU values).
    - **Memory:** Specify the memory allocated to guest virtual machine. For a running guest virtual machine, this can be increased up to the Max Memory.

- **CPU Cores:** Specify the number of virtual CPUs for the guest virtual machine. For a running guest virtual machine, this can be increased up to the Max virtual CPU.
  - **CPU Scheduling Priority:** Set the Priority of the virtual machine. You can select a high, medium, low, or custom priority for the virtual CPUs.
  - **CPU Scheduling CAP:** This parameter defines the percentage to which the virtual CPUs can receive scheduled time. You can select a high, medium, low, or a custom percentage of scheduled time for the virtual CPUs.
  - **Network:** You can add delete virtual network interfaces. For more details, see [Section 4.7, "Setting Up Networks"](#).
  - **Storage:** Select the desired storage configuration of your virtual machine, such as virtual disks and ISO file.
4. Click the **Other Configuration** tab.

**Figure 8–7 Edit Guest Virtual Machine: Other Configuration**

VM Guest Page Refreshed Jan 11, 2013 11:45:58 PM PST

Hardware Configuration Other Configuration

**Attribute**  
Set the basic attributes of Guest VM

Virtual Machine Name: assembly\_ohs\_1102/FPMW\_Wfs/Clu  
Description: Type: wls CatalogId: 4yTshoISqIaY  
Operating System Type: None

**High Availability**  
Guest VM can be startup automatically after outage if HA enabled  
High Availability ☐

**Boot Sequence**  
By changing the boot order, Guest VM can be forced to boot from Disk, CDROM, or Network.

| Sequence Number | Boot Type |
|-----------------|-----------|
| 1               | None      |
| 2               | None      |
| 3               | None      |

Network (PXE) boot path

**Advanced**  
Configure KeyMap Name, Mouse Type and Domain Type

KeyMap Name: en-us (English, United States)  
Mouse Type: Default  
Domain Type: Xen PVM  
VM Start Policy: Use server pools VM start policy  
Tags:

- **Attribute:** You can set the basic attributes of the guest virtual machine such as the description and the type of operating system to be used.
- **High Availability:** If the **High Availability** checkbox is selected, the guest virtual machine will automatically restart after any outage.
- **Boot Sequence:** Specify the boot sequence for the guest virtual machine. This is the boot order for the guest virtual machine and indicates if the guest virtual machine is forced to boot from disk, CDROM, or network.
- **Advanced:**
  - **Keymap Name:** The name of the keyboard mapping that will be used by the guest virtual machine.
  - **Mouse Type:** The mouse type to be used for the guest virtual machine. This can be Default, PS2 Mouse, USB Mouse, or USB Tablet.
  - **Domain Type:** The domain type of the virtual machine. This can be:

**Xen HVM:** Hardware virtualization, or fully virtualized. When you select this option you must supply an ISO file in a repository (in the Create Storage step of the wizard) from which to create the virtual machine.

**Xen HVM, PV Drivers:** Identical to Xen HVM, but with additional paravirtualized drivers for improved performance of the virtual machine. This Domain Type is used to run Microsoft Windows guest operating systems with an acceptable performance level.

**Xen PVM:** Enables you to select a location for the mounted ISO file from which to create the virtual machine. Before you create the virtual machine using the paravirtualized method, mount the ISO file on an NFS share, or HTTP or FTP server. You supply the location of the mounted ISO file in the Network Boot Path field in the Boot Options step of the wizard.

**OVM/SPARC:** This domain type should be selected if the server pool and hypervisors use Oracle VM Server for SPARC as the hypervisor instead of Oracle VM Server for x86.

**Unknown:** This hypervisor should be selected if the domain type is unknown.

- **VM Start Policy:** Select the policy to be used to start the virtual machines. You can choose:

**Start on Best Server:** The best server in the server pool is used to start the virtual machine.

**Start on Current Server:** The virtual machine will be started on the same server on which it was created.

- **Tags:** Specify one or more comma separated tags. These tags are used to logical group the virtual machines and can be used to search for one or more virtual machines that meet a certain criteria.

### 8.7.7 Saving a Guest Virtual Machine as a Template

You can save a guest virtual machine as a template and use the template for provisioning guest virtual machines.

To save a guest virtual machine as a template, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Select a Guest VM target from the left panel. From the **Guest VM** menu, select **Target Setup**, then select **Save Guest VM as a Template**. The Save Guest VM as Template: Select Target and Source page appears.
3. The name of the guest virtual machine target that is to be saved as template is displayed.
4. In the Target Repository field, select the repository on which the guest virtual machine template is to be stored.
5. Specify a name for the template and click **Next**.
6. In the Save Guest VM as Template: Deployment Configuration page, you can modify the following details:
  - **General Configuration:**
    - **Memory and No. of CPU Cores:** Specify the Memory and No of CPU Cores (virtual CPUs) for the virtual machine. The number of processors is expressed in number of physical CPU cores, and is limited to 32. The



memory and the number of processors that you can select is limited to the properties of the Virtual Machine Server. If you do not select a Virtual Machine Server however, you can set your desired value. You can modify the Max Virtual CPUs and Max Memory fields without restarting the Guest VM.

- **CPU Scheduling Priority:** Set the Priority of the virtual machine. You can select a high, medium, low, or custom priority for the virtual CPUs.
- **CPU Scheduling CAP:** This parameter defines the percentage to which the virtual CPUs can receive scheduled time. You can select a high, medium, low, or a custom percentage of scheduled time for the virtual CPUs.
- **Storage:** Select the desired storage configuration of your virtual machine, such as virtual disks and ISO file.
- **Tags:** Specify one or more comma separated tags. These tags are used to logical group the virtual machines and can be used to search for one or more virtual machines that meet a certain criteria.
- **Disk:** Specify the storage type. You can use the Default type as specified by the server pool or select Override. If you select Override, you must specify the Storage QoS and click the Search icon in the Storage field to specify the Storage Type. In the Storage Type window, specify the following:
  - **Clone Target Type:** This can be Repository or Storage Array.
  - **Clone Target:** If you select Repository, select the storage location where the disk is to be created. If you select Storage Array, specify the type, which can be SAN or iSCSI storage server.
  - **Clone Type:** This can be Sparse or Non-Sparse Copy. A sparse copy is a disk image file of a physical disk, taking up only the amount of space actually in use; not the full specified disk size. A non-sparse copy is a disk image file of a physical disk, taking up the space equivalent to the full specified disk size, including empty blocks.

## 8.8 Deploying Guest Virtual Machines

This section describes the various options that can be used to deploy guest virtual machines. It contains the following sections:

- [Getting Started](#)
- [Deployment Options](#)
- [Prerequisites](#)
- [Provisioning Guest Virtual Machines Using Oracle Virtual Assemblies \(OVA\)](#)
- [Provisioning a Guest Virtual Machine Using Oracle VM Templates](#)
- [Provisioning a Guest Virtual Machine Using an ISO Image](#)
- [Creating PXE Bootable Guest Virtual Machines](#)

### 8.8.1 Getting Started

This section helps you get started by providing an overview of the steps involved in provisioning virtualized systems. Consider this section to be a documentation map to understand the sequence of actions you must perform to successfully provision

virtualized systems. Click the reference links provided against the steps to reach the relevant sections that provide more information.

**Table 8–1 Getting Started with Oracle VM Provisioning**

| Step   | Description                                                                                                                                                                                                                              | Reference Link                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Step 1 | <b>Understanding the Deployment Options</b><br>Understand the deployment plans offered by Enterprise Manager for deploying virtualization systems.                                                                                       | To learn about the deployment options, see <a href="#">Section 8.8.2, "Deployment Options"</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| Step 2 | <b>Meeting the Prerequisites</b><br>Before you run the Deployment Procedure, you must meet the prerequisites, such as setting up the provisioning environment, applying mandatory patches, and creating the Software Library components. | To learn about the prerequisites for provisioning guest virtual machines, see <a href="#">Section 8.8.3, "Prerequisites"</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| Step 3 | <b>Selecting the Use Case</b><br>This section covers a few use cases for the provisioning of guest virtual machines. Select the use case that best matches your requirements.                                                            | To provision a guest virtual machine using: <ul style="list-style-type: none"> <li>■ Assemblies: See <a href="#">Section 8.8.4, "Provisioning Guest Virtual Machines Using Oracle Virtual Assemblies (OVA)"</a></li> <li>■ Templates: See <a href="#">Section 8.8.5, "Provisioning a Guest Virtual Machine Using Oracle VM Templates"</a></li> <li>■ ISO Images: See <a href="#">Section 8.8.6, "Provisioning a Guest Virtual Machine Using an ISO Image"</a></li> <li>■ PXE Boot: See <a href="#">Section 8.8.7, "Creating PXE Bootable Guest Virtual Machines"</a></li> </ul> |

## 8.8.2 Deployment Options

Enterprise Manager provides the following deployment options for provisioning guest virtual machines:

- **Virtual Machine Provisioning Using Assembly:** Creates new guest virtual machines using an assembly component.
- **Virtual Machine Provisioning Using Template:** Creates new guest virtual machines using a template component.
- **Virtual Machine Provisioning Using ISO:** Creates a new guest virtual machine using an ISO image.
- **Virtual Machine Provisioning Using PXE:** Creates a new guest virtual machine using Preboot Execution Environment. The new guest virtual machines will be PXE-booted and they will get provisioned by way of networks.
- **Cloning a Virtual Machine:** Clones a guest virtual machine to one or more guest virtual machines.

## 8.8.3 Prerequisites

Before you run any deployment procedures, meet the infrastructure requirements described in this section.

- Mandatory Infrastructure Requirements as defined in the Enterprise Manager Lifecycle Administrator's Guide.
- Optional Infrastructure Requirements as defined in the Enterprise Manager Lifecycle Administrator's Guide.
- The assembly, template, or the ISO image to be deployed must be present as a Software Library component, or at an external HTTP or NFS location, or imported to the server pool shared location.

### 8.8.3.1 Creating and Storing Virtualization Components in the Software Library

This section describes the procedure to create the virtualization components in the Software Library. It includes the following:

- [Creating an Assembly Component](#)
- [Creating a Template Component](#)
- [Creating an ISO Component](#)
- [Creating a Virtual Disk](#)

**8.8.3.1.1 Creating an Assembly Component** An assembly is a collection of virtual machine instances. When an assembly is deployed, a group of virtual machines or an **assembly instance** is created.

#### Prerequisites

- The Assembly Binary, a file with a .ova extension must be present on the Management Agent or on an external `http/ftp/nfs` location.

Follow this process to create a new assembly component:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Software Library**.
2. Select a folder on the Software Library page. From the Actions menu, select **Create Entity**, then select **Virtualization**. You can also right click on a folder and select the virtualization option from the context menu.
3. From the **Create Entity: Virtualization** pop-up window, choose **Assembly** from the Select Subtype drop down list.
4. In the Create Assembly: Describe page, enter the following details:
  - Name and description of the assembly.
  - In the Attachments section, click **Add** to attach any files relevant to the entity. Ensure that the file size is less than 2 MB.
  - In the Notes section, include information related to the entity like changes being made to the entity or modification history that you want to track and click **Add**.
5. Click **Next** to continue. In the Create Assembly: Upload Files page, select one or more files to be associated with the assembly. You can either:
  - **Upload Files:** You can upload files from a storage location in the Software Library. For Software Library to become usable, at least one upload file location must be configured. Select the **Upload Files** option. In the Specify Destination section, click the **Browse** button in the Upload Location field. Two storage options are supported:

- **OMS Shared File System:** An OMS Shared File System location is required to be shared (or mounted) across all the Oracle Management Server (OMS) hosts. This option is ideal for UNIX systems.

For single OMS environments, you can configure the Software Library either on the host where the OMS is running or in a shared location, so that it is accessible to all the OMS hosts. For multiple OMS environments, Oracle recommends that you configure the Software Library in a shared location so that the storage is accessible through NFS mount points to all Oracle Management Servers in the environment

- **OMS Agent File System:** An OMS Agent File System location is a location that is accessible to one of the OMS host's Agent. This option is ideal for OMS installed on Windows hosts. By selecting this option for uploading files, you can avoid sharing a location between all participating OMS hosts.

Credentials must be set before using an OMS Shared File System or OMS Agent File System. For an OMS Shared File System, normal host credential must be set before configuring a storage location. However, for OMS Agent File System location configuration, a credential (preferred or named) has to be specified.

- **Refer Files:** Select this option if you want to use an external file location where the software library files are staged. Referenced File Locations support three storage options:

- **HTTP:** An HTTP storage location represents a base URL which acts as the source of files that can be referenced. For example, the base URL `http://my.files.com/scripts` could be configured as an HTTP location for sourcing files such as `http://my.files.com/scripts/perl/installMyDB.pl` or `http://my.files.com/scripts/linux/stopMyDB.sh`.
- **Agent:** A storage location is similar to the OMS Agent File System option, but can be any host monitored by an Management Agent. The Management Agent can be configured to serve the files located on that host. For example, the directory `/u01/binaries` on the Enterprise Manager Host `my.em.file.server` could be configured as an Agent location for sourcing files such as `/u01/binaries/rpms/myCustomDB.rpm` or `/u01/binaries/templates/myTemplate.tar.gz`
- **NFS:** An NFS storage location represents an exported file system directory on a server. The server need not be an Enterprise Manager host target. For example, the directory `/exported/scripts` is exported on server `my.file.server` could be configured as an NFS location for sourcing files such as `/exported/scripts/generic/installMyDB.pl` or `/exported/scripts/linux/stopMyDB.sh` once mounted on a target host file system.

These locations require a named credential to be associated which will be used to access the files from the base location on the host through the Management Agent.

- **Specify Source:** You can add additional file storage locations.
  - To add a HTTP location that can be accessed through a HTTP URL, select HTTP from the Storage Type list and click **Add**. In the Add HTTP Location dialog box, enter a unique name and a HTTP location for the

storage that you want to reference, and click **OK**. A new entry for the storage location is created, with details like Name, Location, and Status.

- To add a Management Agent location that has read-only privileges set on it, select Agent from the Storage Type list and click **Add**.
6. Click **Next**. In the Assembly Details page, you will see a set of predefined deployment profiles. You will also see the structure of the assembly and all the appliances in the assembly. Click **Next**.
  7. On the Review page, click **Save** to save the assembly component. Click **Save and Upload** to save the assembly component and upload it to the Software Library.

**8.8.3.1.2 Creating a Template Component** Follow these steps to create a template component:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Software Library**.
2. Select a folder on the Software Library page. From the **Actions** menu, select **Create Entity**, then select **Virtualization**. You can also right click on a folder and select the virtualization option from the context menu.
3. Select the Subtype as **Template**.
4. In the Create Template: Describe page, enter the name and description of the template.
5. Select the Operating System Name, Type, and Virtualization Type. The Virtualization Type can be Para Virtualized, Hardware Virtualized or both. Select **Para Virtualized** if you want use machines that are not enabled for virtualization. Select **Hardware Virtualized** if you want to use machines that are enabled for virtualization.
  - **Para Virtualized:** A virtual machine with a kernel that is recompiled to be made aware of the virtual environment. Runs at near native speed, with memory, disk and network access optimized for maximum performance.
  - **Hardware Virtualized:** A virtual machine with an unmodified guest operating system. It is not recompiled for the virtual environment. There may be substantial performance penalties running as a hardware virtualized guest. Enables Microsoft Windows™ operating system to be run, as well as legacy operating systems. Hardware virtualization is only available on Intel VT or AMD SVM CPUs.
6. Click **Next** to continue. In the Create Template: Upload Files page, select one or more files to be associated with the assembly. You can:
  - **Upload Files:** You can upload files from a storage location in the Software Library. For the Software Library to become usable, at least one upload file location must be configured. Select the **Upload Files** option. In the Specify Destination section, click the **Browse** button in the Upload Location field. The following storage options are supported:

- **OMS Shared File System:** An OMS Shared File System location is required to be shared (or mounted) across all the Oracle Management Service (OMS) hosts. This option is ideal for UNIX systems.

For single OMS environments, you can configure the Software Library either on the host where the OMS is running or in a shared location, so that it is accessible to all the OMS hosts. For multiple OMS environments, Oracle recommends that you configure the Software Library in a shared

location so that the storage is accessible through NFS mount points to all Oracle Management Servers in the environment

- **OMS Agent File System:** An OMS Agent File System location is a location that is accessible to one of the OMS host's Agent. This option is ideal for OMS installed on Windows hosts. By selecting this option for uploading files, you can avoid sharing a location between all participating OMS hosts.

Credentials must be set before using an OMS Shared File System or OMS Agent File System. For an OMS Shared File System, normal host credential must be set before configuring a storage location. However, for OMS Agent File System location configuration, a credential (preferred or named) has to be specified.

- **Refer Files:** Select this option if you want to use a external file location where the software library files are staged. Referenced File Locations support three storage options:
  - **HTTP:** An HTTP storage location represents a base URL which acts as the source of files that can be referenced. For example, the base URL `http://my.files.com/scripts` could be configured as an HTTP location for sourcing files such as `http://my.files.com/scripts/perl/installMyDB.pl` or `http://my.files.com/scripts/linux/stopMyDB.sh`.
  - **Agent:** A storage location is similar to the OMS Agent File System option, but can be any host monitored by the Management Agent. The Management Agent can be configured to serve the files located on that host. For example, the directory `/u01/binaries` on the Enterprise Manager Host `my.em.file.server` could be configured as an Agent location for sourcing files such as `/u01/binaries/rpms/myCustomDB.rpm` or `/u01/binaries/templates/myTemplate.tar.gz`
  - **NFS:** An NFS storage location represents an exported file system directory on a server. The server need not be an Enterprise Manager host target. For example, the directory `/exported/scripts` is exported on server `my.file.server` could be configured as an NFS location for sourcing files such as `/exported/scripts/generic/installMyDB.pl` or `/exported/scripts/linux/stopMyDB.sh` once mounted on a target host file system.

These locations require a named credential to be associated which will be used to access the files from the base location on the host through the Management Agent.

- **Specify Source:** You can add additional file storage locations.
    - To add a HTTP location that can be accessed through a HTTP URL, select HTTP from the Storage Type list and click **Add**. In the Add HTTP Location dialog box, enter a unique name and a HTTP location for the storage that you want to reference, and click **OK**. A new entry for the storage location is created, with details like Name, Location, and Status.
    - To add the Management Agent location that has read-only privileges set on it, select Agent from the Storage Type list and click **Add** and enter the required information in the Add Agent Location dialog box.
7. Click **Next**. In the Review page, review the details of the template and click **Save**. Click **Save and Upload** to save the template component and upload it to the

Software Library. The template component will now appear as an entity in the Software Library page.

#### 8.8.3.1.3 Creating an ISO Component Follow these steps to create an ISO component:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Software Library**.
2. Select a folder on the Software Library page. From the **Actions** menu, select **Create Entity**, then select **Virtualization**. You can also right click on a folder and select the virtualization option from the context menu.
3. Select the Subtype as **ISO**.
4. In the Create ISO: Describe page, enter the name and description of the deployment plan.
5. Click **Next** to continue. In the Create ISO: Upload Files page, select one or more files to be associated with the assembly. You can either:

- **Upload Files:** You can upload files from a storage location in the Software Library. For Software Library to become usable, at least one upload file location must be configured. Select the **Upload Files** option. In the Specify Destination section, click the **Browse** button in the Upload Location field. The following storage options are supported:
  - **OMS Shared File System:** An OMS Shared File System location is required to be shared (or mounted) across all the Oracle Management Service (OMS) hosts. This option is ideal for UNIX systems.  
  
For single OMS environments, you can configure the Software Library either on the host where the OMS is running or in a shared location, so that it is accessible to all the OMS hosts. For multiple OMS environments, Oracle recommends that you configure the Software Library in a shared location so that the storage is accessible through NFS mount points to all Oracle Management Servers in the environment
  - **OMS Agent File System:** An OMS Agent File System location is a location that is accessible to one of the OMS host's Agent. This option is ideal for OMS installed on Windows hosts. By selecting this option for uploading files, you can avoid sharing a location between all participating OMS hosts.

Credentials must be set before using an OMS Shared File System or OMS Agent File System. For an OMS Shared File System, normal host credential must set before configuring a storage location. However, for OMS Agent File System location configuration, a credential (preferred or named) has to be specified.

- **Refer Files:** Select this option if you want to use a external file location where the software library files are staged. Referenced File Locations support three storage options:
  - **HTTP:** An HTTP storage location represents a base URL which acts as the source of files that can be referenced. For example, the base URL `http://my.files.com/scripts` could be configured as an HTTP location for sourcing files such as  
`http://my.files.com/scripts/perl/installMyDB.pl` or  
`http://my.files.com/scripts/linux/stopMyDB.sh`.
  - **Agent:** A storage location is similar to the OMS Agent File System option, but can be any host monitored by the Management Agent. The

Management Agent can be configured to serve the files located on that host. For example, the directory `/u01/binaries` on the Enterprise Manager Host `my.em.file.server` could be configured as an Agent location for sourcing files such as `/u01/binaries/rpms/myCustomDB.rpm` or `/u01/binaries/templates/myTemplate.tar.gz`

- **NFS:** An NFS storage location represents an exported file system directory on a server. The server need not be an Enterprise Manager host target. For example, the directory `/exported/scripts` is exported on server `my.file.server` could be configured as an NFS location for sourcing files such as `/exported/scripts/generic/installMyDB.pl` or `/exported/scripts/linux/stopMyDB.sh` once mounted on a target host file system.

These locations require a named credential to be associated which will be used to access the files from the base location on the host through the Management Agent.

- **Specify Source:** You can add additional file storage locations.
  - To add a HTTP location that can be accessed through a HTTP URL, select HTTP from the Storage Type list and click **Add**. In the Add HTTP Location dialog box, enter a unique name and a HTTP location for the storage that you want to reference, and click **OK**. A new entry for the storage location is created, with details like Name, Location, and Status.
  - To add an NFS shared location, select NFS from the Storage Type list and click **Add**. In the Add NFS Location dialog box, enter a unique name in the Name field for the storage. In NFS server field, provide a fully qualified domain name or the IP address of the hosted machine that has NFS services running on them. In the Location field, provide the shared location or directory path on the NFS server to define a storage location, then click **OK**. A new entry for the storage location is created in the table, with details like Name, Location, and Status.
  - To add a Management Agent location that has read-only privileges set on it, select Agent from the Storage Type list and click **Add**.
- 6. Click **Next**. In the Review page, review the details and click **Save**. Click **Save and Upload** to save the assembly component and upload it to the Software Library. The ISO image will now appear as an entity in the Software Library page.

**8.8.3.1.4 Creating a Virtual Disk** Follow these steps to create a virtual disk component:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Software Library**.
2. Select a folder on the Software Library page. From the **Actions** menu, select **Create Entity**, then select **Virtualization**. You can also right click on a folder and select the virtualization option from the context menu.
3. Select the Subtype as **Virtual Disk**.
4. In the Virtual Disk: Describe page, enter the name and description of the deployment plan.
5. Click **Next** to continue. In the Create Virtual Disk: Upload Files page, select one or more files to be associated with the assembly. You can either:
  - **Upload Files:** You can upload files from a storage location in the Software Library. For Software Library to become usable, at least one upload file location must be configured. Select the Upload Files option. In the Specify



Destination section, click the **Browse** button in the Upload Location field. The following storage options are supported:

- **OMS Shared File System:** An OMS Shared File System location is required to be shared (or mounted) across all the Oracle Management Server (OMS) hosts. This option is ideal for UNIX systems.

For single OMS environments, you can configure the Software Library either on the host where the OMS is running or in a shared location, so that it is accessible to all the OMS hosts. For multiple OMS environments, Oracle recommends that you configure the Software Library in a shared location so that the storage is accessible through NFS mount points to all Oracle Management Servers in the environment

- **OMS Agent File System:** An OMS Agent File System location is a location that is accessible to one of the OMS host's Agent. This option is ideal for OMS installed on Windows hosts. By selecting this option for uploading files, you can avoid sharing a location between all participating OMS hosts.

Credentials must be set before using an OMS Shared File System or OMS Agent File System. For an OMS Shared File System, normal host credential must set before configuring a storage location. However, for OMS Agent File System location configuration, a credential (preferred or named) has to be specified.

- **Refer Files:** Select this option if you want to use an external file location where the software library files are staged. Referenced File Locations support three storage options:
  - **HTTP:** An HTTP storage location represents a base URL which acts as the source of files that can be referenced. For example, the base URL `http://my.files.com/scripts` could be configured as an HTTP location for sourcing files such as `http://my.files.com/scripts/perl/installMyDB.pl` or `http://my.files.com/scripts/linux/stopMyDB.sh`.
  - **Agent:** A storage location is similar to the OMS Agent File System option, but can be any host monitored by the Management Agent. The Agent can be configured to serve the files located on that host. For example, the directory `/u01/binaries` on the Enterprise Manager Host `my.em.file.server` could be configured as an Agent location for sourcing files such as `/u01/binaries/rpms/myCustomDB.rpm` or `/u01/binaries/templates/myTemplate.tar.gz`
  - **NFS:** An NFS storage location represents an exported file system directory on a server. The server need not be an Enterprise Manager host target. For example, the directory `/exported/scripts` is exported on server `my.file.server` could be configured as an NFS location for sourcing files such as `/exported/scripts/generic/installMyDB.pl` or `/exported/scripts/linux/stopMyDB.sh` once mounted on a target host file system.

These locations require a named credential to be associated which will be used to access the files from the base location on the host through the Management Agent.

- **Specify Source:** You can add additional file storage locations.
  - To add a HTTP location that can be accessed through a HTTP URL, select HTTP from the Storage Type list and click **Add**. In the Add HTTP

- Location dialog box, enter a unique name and a HTTP location for the storage that you want to reference, and click **OK**. A new entry for the storage location is created, with details like Name, Location, and Status.
- To add an Agent location that has read-only privileges set on it, select Agent from the Storage Type list and click **Add**.
6. Click **Next**. In the Review page, review the details of the virtual disk and click **Save**. Click **Save and Upload** to save the virtual disk component and upload it to the Software Library. The virtual disk component will now appear as an entity in the Software Library page.

## 8.8.4 Provisioning Guest Virtual Machines Using Oracle Virtual Assemblies (OVA)

An assembly is a packaging of a multi-tier application stack, which when deployed results in the creation of a set of related virtual machines representing every tier of the application stack. For example, a Siebel assembly can contain the Siebel mid-tier and a database. It can be deployed such that it creates 2 virtual machines for a 2-node Siebel mid-tier and 1 virtual machine running the database.

The deployment of the assembly is based on the policies/constraints defined in the assembly definition. The initial placement of virtual machines will be based on the storage, network and computational resources. Once an assembly is deployed, additional instances can be added (scale up) or removed (scale down) based on application demand.

When an assembly is deployed, an assembly instance is created. An assembly instance is a collection of dependent virtual machines, which is created after an assembly is deployed. The composition of the assembly instance is dynamic in nature, as any subsequent scale-up or scale-down operations will change the membership. An assembly instance contains one or more tiers, each tier further comprising of virtual machines and or tiers.

Follow this process to deploy an assembly and create the virtual machines.

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Right click on a Zone, VM Server Pool, or VM Server target and select **Deploy Assembly**.

**Figure 8–8 Assembly Deployment: Instance Details**

The screenshot shows the 'Assembly Deployment: Instance Details' page. The top navigation bar includes 'Instance Details', 'Network and Agent', 'Deployment Configuration', 'Schedule', and 'Review'. The page is titled 'Assembly Deployment : Instance Details' and has a progress indicator 'Step 1 of 5' with 'Back', 'Next', and 'Cancel' buttons.

**Destination**  
Select the Destination target on which Guest VM(s) is/are to be deployed.  
Target: EastCoastZone  
Target Type: Oracle VM Zone

**Source**  
Select Source image for deployment.  
Deployment Plans allows the recording of inputs and using them later while performing deployments using standard source images and values.  
Selecting a Deployment Plan at this stage allows you to pre-populate the current interview with previously saved values and configuration.  
Use Deployment Plan: ☐  
Deployment Plan:   
Image: FFW Assembly

**Details**  
Request Name: ASSEMBLY\_DEPLOYMENT\_2013-06-  
Assembly Instance Name: example\_assembly  
Allow Partial Deployment: ☐  
Tip: If this option is enabled, and the deployment process fails, you must manually delete any newly created Guest VMs.

**Assembly Content**  
Select Tiers/VM Nodes in the Assembly that are to be deployed.  
A Tier/VM Node can be disabled for deployment only if the minimum number of instances for it is defined to be 0 in the OVF metadata.  
A Tier/VM Node cannot be disabled if a deployment plan is being used where this Tier/VM Node was earlier enabled for deployment.

| Node Name                              | Start Order | Deploy                              |
|----------------------------------------|-------------|-------------------------------------|
| assembly_ohs_1102/PMW_Wls/AdminServer  | 0           | <input checked="" type="checkbox"/> |
| assembly_ohs_1102/PMW_Wls/CacheCluster | 1           | <input checked="" type="checkbox"/> |
| assembly_ohs_1102/PMW_Wls/Cluster1     | 1           | <input checked="" type="checkbox"/> |
| assembly_ohs_1102/PMW_ohs              | 2           | <input checked="" type="checkbox"/> |

3. On the Assembly Deployment: Instance Details page, specify the following:

- **Target:** Select the server pool on which the guest virtual machines are to be deployed.
- **Source:** You can select either of the following options:
  - **Deployment Plan:** If you select the **Use Deployment Plan** check box, you can select a deployment plan (if available) with predefined configurations. The deployment plan will be used when the assembly is deployed. (image is selected automatically when the deployment plan is selected)
  - **Image:** The assembly component to be used for this deployment. This can either be a component present in the Software Library, or at an external location. If it is stored in an external location, specify the NFS or HTTP location where the component is stored. For example, `http://myvm_server/assembly` or `file://myvm_server/assembly`.
- **Details:**
  - **Request Name:** This is the name of the assembly deployment request.
  - **Assembly Instance Name:** The name to be assigned to the assembly instance.
  - **Allow Partial Deployment:** Select this option to enable partial deployment. By default, when a deployment plan fails, all virtual machines that have been created are automatically deleted. But if partial deployment is enabled, virtual machines are not deleted even if the deployment has failed.
- **Assembly Content:** You can disable an assembly tier or a node from being deployed. By default, all nodes and tiers in the selected component are automatically deployed. To disable a node or tier from being deployed, unselect the **Deploy** checkbox in the table.

---



---

**Note:**

- A tier or a node can be disabled for deployment if the minimum number of instances for the node or tier is set to **0** in the **ovf** metadata file.
  - If you choose to use a deployment plan where a tier or a node has been enabled for deployment, that node or tier cannot be disabled.
- 
- 

4. Click **Next**. In the Deploy Assembly: Network and Agent page, specify the physical public and private networks to be used for the assembly instance. A public network is used for external traffic. In a private network, the guest virtual machines can communicate only within a predefined network set up by the Cloud Administrator.

Assembly network refers to the logical network defined in the OVF metadata file and the network defined here is the actual physical network to which it is to be mapped. If you do not choose to associate the network explicitly, the mappings are automatically assigned.

**Figure 8–9 Assembly Deployment: Network and Agent**

**Assembly Deployment : Network and Agent**

▼ **Network configured in the table will be deployed in this request.**  
 Changing the Network QoS/Type to any value other than "Any Network Type" will reset back-end/physical network selection.  
 Network QoS/Type is used to filter the list of applicable networks that can be chosen in the backend/physical network selector. Choose "Any Network Type" if you want Network Selector to display all networks.

| Assembly Network | Type          | Network         |
|------------------|---------------|-----------------|
| vnet-1           | Any Network 1 | System Assigned |

**Configure EM Agent Push**

☒ Enable EM Agent Push Configuration

▼ **Nodes displayed in the table will be deployed in this request. Those which have EM agent install type set as "RPM Installation" cannot be marked for push installation.**  
 Nodes that have EM agent install type set as "No Agent Installation" will not have any agent installed after they are deployed.  
 VM level agent push selection can be done in the Miscellaneous tab of the respective Tier.

Assembly Node Selection for EM Agent Push Selection

| Node Name                               | Installation Type       | Enable Agent Push                   |
|-----------------------------------------|-------------------------|-------------------------------------|
| assembly_ohs_1102/FMW_Wlts/AdminServer  | Push Agent Installation | <input checked="" type="checkbox"/> |
| assembly_ohs_1102/FMW_Wlts/CacheCluster | Push Agent Installation | <input checked="" type="checkbox"/> |
| assembly_ohs_1102/FMW_Wlts/Cluster 1    | Push Agent Installation | <input checked="" type="checkbox"/> |
| assembly_ohs_1102/FMW_Ohs               | Push Agent Installation | <input checked="" type="checkbox"/> |

\* Installation Base Directory: /home/oracle/agent

\* Platform: linux

\* Username: oracle

\* Password:

\* Verify Password:

Port: 3972

Additional Parameters: -ignoreDirPrereq -enableTty

Privilege Delegation Settings: /bin/su - %SUDOAS% -c \"%COMMAND%\"

SSH Service Timeout: 10 minute(s)

5. Click on the **Browse** icon in the Network field. Specify if the network should be allocated by the system or selected manually.
6. If you want to install the Management Agent on an assembly tier (collection of Guest VMs) or selected guest virtual machines, check the **Enable EM Agent Push Configuration** check box. You can then select the tiers and / or the guest virtual machines on which the Management Agent is to be installed using the Agent Push approach in the Assembly Node Selection for the EM Agent Push Selection table.
  - **Node Name:** The name of the assembly tier on which the Management Agent is to be deployed.
  - **Installation Type:** The type of installation which can be Push Agent Installation, No Agent Installation, and RPM Installation.
  - **Enable Agent Push:** This checkbox is enabled only if the Enable EM Agent Push Configuration checkbox has been selected. Select this checkbox to select the tiers on which the Management Agent is to be installed by way of the Agent Push approach.

In a selected tier, you can also choose the guest virtual machines on which the Management Agent is to be installed.

---



---

**Note:**

- If you select Push Agent Installation, you can specify the guest virtual machines that need to be enabled or disabled in the Miscellaneous tab.
  - For tiers that have the Management Agent as a product (as part of the binary package), the **Enable Agent Push** checkbox is disabled and only RPM Installation is available. In this case, you can specify the Management Agent details in the Product Configuration tab of the selected tier.
- 
-

7. After you have specified the options to configure the Management Agent, enter the following details:
  - **Installation Base Directory:** This is the parent directory on the remote host where the Oracle home of the Management Agent is created.
  - **Platform:** The operating system on which the Management Agent is to be installed.
  - **User:** The name of the Oracle Install user is displayed here.
  - **Port:** The port used by the Management Agent for communication.
  - **Password:** This field is mandatory for PS5 assemblies if the **Enable Agent Push Configuration** checkbox has been selected.
  - **Privilege Delegation Settings:** Run as root.
  - **SSH Service Timeout:** The timeout settings for SSH.
  - **Additional Parameters:** Specify additional parameters to be used during the installation of the Management Agent. This field is pre-populated with the entry `-ignoreDirPrereq -enablePty` to ignore the prerequisite checks during installation.
8. Click **Next** to go to the Deploy Assembly: Deployment Configuration page.

**Figure 8–10 Assembly Deployment: Deployment Configuration**

Instance Details Network and Agent **Deployment Configuration** Schedule Review

Assembly Deployment : Deployment Configuration Review Save As Plan... Back Step 3 of 5 Next Cancel

Select each Component in the table to View and Customize the default configuration properties.

| Component                                    | Guest VM(s) |         |         |         | Auto Scalable            | VM Size | Deployment Option | Guest VM Name Prefix       |
|----------------------------------------------|-------------|---------|---------|---------|--------------------------|---------|-------------------|----------------------------|
|                                              | Default     | Minimum | Maximum | Initial |                          |         |                   |                            |
| example_assembly                             |             |         |         |         |                          |         |                   |                            |
| assembly_ohs_1119/FMW_Ohs:example_assembly   | 1           | 1       | 1       | 1       | <input type="checkbox"/> | Custom  | Not Defined       |                            |
| assembly_ohs_1119/FMW_Wls/AdminServer:exam   | 1           | 1       | 1       | 1       | <input type="checkbox"/> | Custom  | Not Defined       |                            |
| assembly_ohs_1119/FMW_Wls/CacheCluster:exam  | 2           | 1       | 2       | 2       | <input type="checkbox"/> | Custom  | Not Defined       | assembly_ohs_1119/FMW_Wls/ |
| assembly_ohs_1119/FMW_Wls/Cluster1:example_a | 2           | 1       | 2       | 2       | <input type="checkbox"/> | Custom  | Not Defined       | assembly_ohs_1119/FMW_Wls/ |

Selected Components

No Tier/Assembly Instance selected in the table above.

The following details are displayed:

- **Component:** The assembly component that is being deployed.
- **Guest VMs:** The default number of guest virtual machines as specified in the deployment profile appears. Specify the minimum and maximum number of guest virtual machines that are to be created.
- **Auto Scalable:** Select this check box to indicate that the tier instances can be auto scaled based on their policies.
- **VM Size:** The amount of memory to be allocated to the guest virtual machine. The virtual server should have sufficient memory to run the guest virtual machine. Otherwise, the guest virtual machine will be created, but will not be started. The size of the guest virtual machine (memory and CPU values).
- **Deployment Option:** The deployment for this assembly instance.

- **Guest VM Name Prefix:** This string will be used as a prefix in the guest virtual machine name. When a guest virtual machine is created using the Enterprise Manager, guest virtual machine names will be prefixed by the virtual machine name prefix you provide here and a number.
9. Select the assembly component to view and customize the default configuration properties.

**Figure 8–11 Assembly Deployment: Deployment Configuration: Product Configuration**

Instance Details Network and Agent **Deployment Configuration** Schedule Review

Assembly Deployment : Deployment Configuration Review Save As Plan... Back Step 3 of 5 Next Cancel

Select each Component in the table to View and Customize the default configuration properties.

| Component                                      | Guest VM(s) |         |         |         | Auto Scalable            | VM Size | Deployment Option | Guest VM Name Prefix       |
|------------------------------------------------|-------------|---------|---------|---------|--------------------------|---------|-------------------|----------------------------|
|                                                | Default     | Minimum | Maximum | Initial |                          |         |                   |                            |
| example_assembly                               |             |         |         |         |                          |         |                   |                            |
| assembly_ohs_1119/FMW_Ohs:example_assembly     | 1           | 1       | 1       | 1       | <input type="checkbox"/> | Custom  | Not Defined       |                            |
| assembly_ohs_1119/FMW_Wls/AdminServer:example  | 1           | 1       | 1       | 1       | <input type="checkbox"/> | Custom  | Not Defined       |                            |
| assembly_ohs_1119/FMW_Wls/CacheCluster:example | 2           | 1       | 2       | 2       | <input type="checkbox"/> | Custom  | Not Defined       | assembly_ohs_1119/FMW_Wls/ |
| assembly_ohs_1119/FMW_Wls/Cluster1:example_a   | 2           | 1       | 2       | 2       | <input type="checkbox"/> | Custom  | Not Defined       | assembly_ohs_1119/FMW_Wls/ |

Selected Component: example\_assembly

**Details example\_assembly**

Product Configuration Miscellaneous

Show All Properties ☒ Show only unassigned properties Search %

| Name                             | Default Value                            | Value        | Category            | Description |
|----------------------------------|------------------------------------------|--------------|---------------------|-------------|
| version                          | 1.0.0                                    | 1.0.0        |                     |             |
| uuid                             | c2a588f257fb57b52f-c2a588f257fb57b52f    |              |                     |             |
| sysprop admin-username           | weblogic                                 | weblogic     | Assembly Properties |             |
| sysprop domain-name              | WLS_OHS_COH                              | WLS_OHS_COH  | Assembly Properties |             |
| component                        | coherenceweb                             | coherenceweb | Assembly Properties |             |
| CoherenceCluster1.multicast-port | 7777                                     | 7777         | Assembly Properties |             |
| sysprop admin-vr-jmx-protocol    | iiop                                     | iiop         | Assembly Properties |             |
| sysprop domain-root              | /u01/apps/oracle/con/u01/apps/oracle/con |              | Assembly Properties |             |

- **Product Configuration:** This tab shows the properties of the products in the assembly. You can choose to view All Properties or Required Properties. For each property, the default value and the current value is displayed. Select a property to modify the value if applicable.
  - **Miscellaneous:** You can create one or more anti-affinity groups. An anti-affinity group is a group of virtual machines cannot run on the same server. Click **Add** and enter the Group Name and select the guest virtual machines that are to be part of the anti-affinity group.
10. Click a row in the table to select the tier you want to configure. There are 3 tabs that can be configured for each tier in the assembly: **Guest VM Configuration**, **Product Configuration**, and **Miscellaneous**.
  11. Let us look at the first tab: **Guest VM Configuration**. This tab allows you to modify the resource configuration parameters for all guest virtual machine instances of the selected assembly tier.

**Figure 8–12 Tier: Deployment Configuration: GuestVM Configuration**

**Assembly Deployment: Deployment Configuration**

Select each Component in the table to View and Customize the default configuration properties.

| Component                                                | Def | Minimum | Maximum | Initial | Aut Sca | VM Size | Deployment Option | Guest VM Name Prefix          |
|----------------------------------------------------------|-----|---------|---------|---------|---------|---------|-------------------|-------------------------------|
| example_assembly                                         |     |         |         |         |         |         |                   |                               |
| assembly_ohs_1102/FMW_Ohs:example_assembly               | 1   | 1       | 1       | 1       | 1       | Custom  | Not Defined       |                               |
| assembly_ohs_1102/FMW_Wlts/AdminServer:example_assembly  | 1   | 1       | 1       | 1       | 1       | Custom  | Not Defined       |                               |
| assembly_ohs_1102/FMW_Wlts/CacheCluster:example_assembly | 2   | 1       | 2       | 2       | 2       | Custom  | Not Defined       | assembly_ohs_1102/FMW_Wlts/Ca |
| assembly_ohs_1102/FMW_Wlts/Cluster1:example_assembly     | 2   | 1       | 2       | 2       | 2       | Custom  | Not Defined       | assembly_ohs_1102/FMW_Wlts/Ch |

Selected Component: assembly\_ohs\_1102/FMW\_Ohs:example\_

**Details: assembly\_ohs\_1102/FMW\_Ohs:example\_assembly**

**GuestVM Configuration**

**General**

☐ Enable High Availability

☒ Start VM After Creation

CPU Scheduling Priority: Medium

CPU Scheduling Cap (%): High

\* Root Password:

\* Confirm Root Password:

Keymap: en-us (English, United States)

Tags:

**VM Size: Custom**

Name: assembly\_ohs\_1102/FMW\_Ohs:example\_assembly

\* Max. Memory (in MB): 512

\* Memory (in MB): 512

\* Max. Number Of CPU Cores: 1

\* Number Of CPU Cores: 1

**Storage Options**

Storage Type: ☒ Default ☐ Override

**Disk**

| Disk Name            | Size (MB) | Type  | Disk Target Storage |
|----------------------|-----------|-------|---------------------|
| KjYJSolSqlaYW_System | 5122      | Local | Default             |
| AB                   | 0         | Local | Default             |

The following parameters can be configured:

■ **Guest VM Configuration:**

- **Enable High Availability:** If you want to enable high availability for this guest virtual machine, select this check box. If this option is enabled, the virtual server is shut down or restarted, the guest virtual machine is migrated to or restarted on another virtual server.

**Note:** High availability must be enabled both on the virtual server pool and on the guest virtual machine. If it is not enabled on both, high availability is disabled.

- **CPU Scheduling Priority:** Specify the priority of a guest virtual machine to access the physical CPUs, while competing with other guest virtual machines on the same virtual server. The higher the priority, the higher the stake of the guest virtual machine in claiming CPU cycles.
- **CPU Scheduling Cap (%):** Restricts the amount of physical CPU that a guest virtual machine can use. Use this to constrain guest virtual machine resource allocation.
- **Root Password:** Specify the operating system root password to log in to the guest virtual machine once it is up and running.
- **Keymap:** The keyboard mapping that will be used by the guest virtual machine being cloned.
- **Tags:** Specify one or more tags for the virtual machine. These tags are used to logical group the virtual machines and can be used to search for one or more virtual machines that meet a certain criteria.

- **Memory (in MB):** The amount of memory to be allocated to the guest virtual machine. The virtual server should have sufficient memory to run the guest virtual machine. Otherwise, the guest virtual machine will be created, but will not be started.
- **CPUs:** The number of CPU cores to assign to this guest virtual machine. This will show up as virtual CPU (VCPU) configuration parameter for the guest virtual machine.
- **Disk:** Specify the storage type. You can use the **Default** type as specified by the server pool or select **Override**. If you select **Override**, you must specify the Storage QoS and click the **Search** icon in the Storage field to specify the Storage Type. In the Storage Type window, specify the following:

**Clone Target Type:** This can be Repository or Storage Array.

**Clone Target:** If you select Repository, select the storage location where the disk is to be created. If you select Storage Array, specify the type, which can be SAN or iSCSI storage server.

**Clone Type:** This can be Sparse or Non-Sparse Copy. A sparse copy is a disk image file of a physical disk, taking up only the amount of space actually in use; not the full specified disk size. A non-sparse copy is a disk image file of a physical disk, taking up the space equivalent to the full specified disk size, including empty blocks.

- **Disk:** Every guest virtual machine is created with at least one disk. You can specify additional disks and their sizes. These disks will be added to the guest virtual machine. After the guest virtual machine is created, log in to the guest virtual machine and reformat and mount the additional disks. Click **Add** to add a storage device (disk).

Specify the name of the hard disk, its size, and QoS Priority Class. You can enable the disk I/O priority by setting an appropriate QoS Priority Class. The priority class ranges from 0 to 7. The priority class 0 has the highest priority, and 7 the lowest.

Select the Shared Disk check box to share this disk between multiple guest virtual machines. You can enable the disk I/O priority by setting an appropriate QoS Priority Class. The priority class ranges from 0 to 7. The priority class 0 has the highest priority, and 7 the lowest.

- **NIC:** Specify the network interfaces for this guest virtual machine. The minimum value is 1 and maximum is 8. Click **Add** to add a network interface and specify the following details:

**Name:** Enter a name for the network resource.

**Network:** Assembly Network / Backend network.

**Network Type:** Specify if the network type is Internet Routable, Non Internet Routable, or a predefined set of IP addresses.

**IP Assignment:** Select the protocol to be used to assign the IP address. This can be DHCP, Network Profile, or None.

The configuration parameters you specify here will be applied to all the guest virtual machine instances of the selected node. If you want to specify a different configuration for a guest virtual machine instance, click the Guest VM Instances option. In the Deployment Configuration: GuestVM



Instance page, you can select a guest virtual machine instance from the View list and specify configuration parameters for that instance.

12. The next tab is Product Configuration. This tab shows the properties of the products in the tier. You can choose to view **All Properties** or **Required Properties**. For each property, the default value and the current value is displayed. Select a property to modify the value if applicable.

**Figure 8–13 Tier Deployment Configuration: Product Configuration**

Assembly Deployment : Deployment Configurati

Select each Component in the table to View and Customize the default configuration properties.

| Component                                   | Guest VM(s) |         |         |         | Auto Scalable            | VM Size | Deployment Option | Guest VM Name Prefix                |
|---------------------------------------------|-------------|---------|---------|---------|--------------------------|---------|-------------------|-------------------------------------|
|                                             | Default     | Minimum | Maximum | Initial |                          |         |                   |                                     |
| example_assembly                            |             |         |         |         |                          |         |                   |                                     |
| assembly_ohs_1119/FMW_Ohs:example_assem     | 1           | 1       | 1       | 1       | <input type="checkbox"/> | Custom  | Not Defined       |                                     |
| assembly_ohs_1119/FMW_Wfs/AdminServer:exar  | 1           | 1       | 1       | 1       | <input type="checkbox"/> | Custom  | Not Defined       |                                     |
| assembly_ohs_1119/FMW_Wfs/CacheCluster:exar | 2           | 1       | 2       | 2       | <input type="checkbox"/> | Custom  | Not Defined       | assembly_ohs_1119/FMW_Wfs/CacheClu  |
| assembly_ohs_1119/FMW_Wfs/Cluster1:example  | 2           | 1       | 2       | 2       | <input type="checkbox"/> | Custom  | Not Defined       | assembly_ohs_1119/FMW_Wfs/Cluster1_ |

Selected Component assembly\_ohs\_1119/FMW\_Ohs:example\_assembly

Details assembly\_ohs\_1119/FMW\_Ohs:example\_assembly

Guest VM Instances

GuestVM Configuration Product Configuration Miscellaneous

Show All Properties ☐ Show only unassigned properties Search %

| Name                             | Default Value       | Value               | Category             | Description |
|----------------------------------|---------------------|---------------------|----------------------|-------------|
| ocm.runConfiguration             | false               | false               | Appliance Properties |             |
| sysprop oracleInstance           | /u01/apps/oracle/co | /u01/apps/oracle/co | Appliance Properties |             |
| system-fileset ORACLE_INSTANCE   | /u01/apps/oracle/co | /u01/apps/oracle/co | File set properties  |             |
| sysprop ocm.ccrDirPath           | /u01/apps/oracle/pi | /u01/apps/oracle/pi | Appliance Properties |             |
| scaling abs-max                  | 2147483647          | 2147483647          | Appliance Properties |             |
| ocm.metalink.CsrRegistration.CSI |                     |                     | Appliance Properties |             |
| sysprop fmwHome                  | /u01/apps/oracle/pi | /u01/apps/oracle/pi | Appliance Properties |             |
| template-target-platform         | OVM                 | OVM                 | Template Properties  |             |
| system-fileset fmwHome root      | /u01/apps/oracle/pi | /u01/apps/oracle/pi | File set properties  |             |

13. The next tab is the Miscellaneous tab.

**Figure 8–14 Tier Deployment Configuration: Miscellaneous**

Assembly Deployment : Deployment Configurati

Select each Component in the table to View and Customize the default configuration properties.

| Component                                   | Guest VM(s) |         |         |         | Auto Scalable            | VM Size | Deployment Option | Guest VM Name Prefix                |
|---------------------------------------------|-------------|---------|---------|---------|--------------------------|---------|-------------------|-------------------------------------|
|                                             | Default     | Minimum | Maximum | Initial |                          |         |                   |                                     |
| example_assembly                            |             |         |         |         |                          |         |                   |                                     |
| assembly_ohs_1119/FMW_Ohs:example_assem     | 1           | 1       | 1       | 1       | <input type="checkbox"/> | Custom  | Not Defined       |                                     |
| assembly_ohs_1119/FMW_Wfs/AdminServer:exar  | 1           | 1       | 1       | 1       | <input type="checkbox"/> | Custom  | Not Defined       |                                     |
| assembly_ohs_1119/FMW_Wfs/CacheCluster:exar | 2           | 1       | 2       | 2       | <input type="checkbox"/> | Custom  | Not Defined       | assembly_ohs_1119/FMW_Wfs/CacheClu  |
| assembly_ohs_1119/FMW_Wfs/Cluster1:example  | 2           | 1       | 2       | 2       | <input type="checkbox"/> | Custom  | Not Defined       | assembly_ohs_1119/FMW_Wfs/Cluster1_ |

Selected Component assembly\_ohs\_1119/FMW\_Ohs:example\_assembly

Details assembly\_ohs\_1119/FMW\_Ohs:example\_assembly

Guest VM Instances

GuestVM Configuration Product Configuration Miscellaneous

Additional Properties

Log Location /assemblybuilder/logs

☒ TIP Copying of logs may consume space in database where job logs are stored. Be careful while specifying location of logs to copy to prevent unnecessary files from being copied.

Timeout Settings

Product Configuration Timeout 60 minute(s)

Network Configuration Timeout 60 minute(s)

- **Miscellaneous:** Specify the following details:

- **Log File Location:** Specify the directory in which the log files are to be stored.
  - **Timeout Settings:** Specify the timeout settings for the Product Configuration and Network Configuration. The default value is 60 minutes and this value can be modified for each Oracle VM Manager target. The timeout value specifies how long Enterprise Manager should wait for either Product Configuration or Network Configuration to be completed before proceeding with the assembly deployment.
14. If a tier has more than one Guest VM Instance, you can configure each Guest VM Instance separately. Select the tier for each Guest VM Instance has to be separately configured. Click the **Guest VM Instances** link on the top right corner of the lower region. The Deployment Configuration: Configure Guest VM Instance page appears. Click **Edit VM Name** to modify the name of the Guest VM Instance, specify a new name and click **OK**. You can modify the configuration of the Guest VM and click **Continue** to return to the Assembly Deployment: Deployment Configuration page.
  15. Click **Next**. In the Assembly Deployment: Schedule page, specify the Start and End Date for the assembly deployment. The Start Date is the date on which the request is being submitted and the End Date is the date on which the assembly deployment instances are retired. Click **Next**.
  16. Click **Next**. The Assembly Deployment: Review page appears.

**Figure 8–15 Assembly Deployment: Review**

| Component               | Default | Guest VM(s) |         | Initial | Auto Scalable | VM Size | Deployment Option | Guest VM Name                        |
|-------------------------|---------|-------------|---------|---------|---------------|---------|-------------------|--------------------------------------|
|                         |         | Minimum     | Maximum |         |               |         |                   |                                      |
| example_assembly        |         |             |         |         |               |         |                   |                                      |
| assembly_ohs_1119/FMW_1 | 1       | 1           | 1       | 1       | ✗             | Custom  |                   |                                      |
| assembly_ohs_1119/FMW_1 | 1       | 1           | 1       | 1       | ✗             | Custom  |                   |                                      |
| assembly_ohs_1119/FMW_1 | 2       | 1           | 2       | 2       | ✗             | Custom  |                   | assembly_ohs_1119/FMW_Wls/CacheClus  |
| assembly_ohs_1119/FMW_1 |         |             |         |         |               |         |                   |                                      |
| assembly_ohs_1119/FMW_1 | 2       | 1           | 2       | 2       | ✗             | Custom  |                   | assembly_ohs_1119/FMW_Wls/Cluster1_v |
| assembly_ohs_1119/FMW_1 |         |             |         |         |               |         |                   |                                      |

17. Click **Next**. Review the details entered so far. You can save as a plan that can be used later. Click **Save as Plan**, specify a name, and description for the plan. Select the folder in which the deployment plan is to be saved and click **OK**. Note that the deployment plan is stored as a software library component but can be created only through this procedure or by a Cloud API that generates a default deployment plan for an assembly.
18. Click **Submit** to submit the job. From the **Enterprise** menu, select **Job**, then select **Activity** to view the Job Activity page. This page shows the details of the submitted job. You can click on the Name link to view more details.

To view details of the deployment request, from the Enterprise menu, select Cloud, then select Infrastructure Request Dashboard. See [Section 8.9, "Viewing the Infrastructure Request Dashboard"](#) for details.

---

**Note:** To ensure that Enterprise Manager has been configured correctly to push a 12c Management Agent automatically on the guest virtual machines when the assembly deployed, you must change the agentpush entry in the <OMS\_ORACLE\_HOME>/sysman/prov/agentpush/agentpush.properties file as follows:

From: oracle.sysman.prov.agentpush.pdpShellOutEnabled=true

To: oracle.sysman.prov.agentpush.pdpShellOutEnabled=false

After this change has been made, you must restart the Oracle Management Service.

---

## 8.8.5 Provisioning a Guest Virtual Machine Using Oracle VM Templates

You can create one or more virtual machines by deploying a template. The deployment of the template is based on the constraints defined while the template was created. Typically, a template can contain an operating system, basic configuration (number of CPUs, memory, disk size), and preinstalled applications. The initial placement of virtual machines will be based on the storage, network and computational resources. Once a template is deployed, additional instances can be added (scale up) or removed (scale down) based on application demand.

---

**Note:** Enterprise Manager supports deployment of legacy Oracle VM 2.2 style Templates. This templates are available in the .tgz file format. Guest virtual machines are created when these templates are deployed, but the you need to manually login to the VNC console of the newly created virtual machines to complete the network and product configuration. Also, unlike for Oracle Virtual Assemblies (.ova file format), Enterprise Manager does not automatically push the Management Agent for Oracle VM 2.2 Templates (.tgz file format).

---

### Prerequisites

- Enterprise Manager Cloud Control is set up to manage the cloud environment.
- The template is available to the zone and you have sufficient privileges to deploy it.
- There are sufficient resources available in the zone for successful deployment.

Follow this process to deploy a template:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Right click on a virtual server pool target and select **Deploy Template**.

**Figure 8–16 Template Deployment - Select Target and Source**

CDC\_ZONE

Select Target and Source Instance Details Schedule Review

Template Deployment : Select Target and Source Back Step 1 of 4 Next Save As Plan... Review Submit Cancel

**Destination**  
Select the destination target on which the Guest VMs are to be deployed.

Target CDC\_ZONE

Target Type oracle\_vm\_zone

**Source**  
Plans allows you to record inputs and use them later while performing deployments using standard source images and values. Selecting a plan at this stage allows you to pre-populate the interview with previously saved values of deployment and configuration parameters. Select an existing plan below, or skip the plan selection to continue with your deployment.

☐ Use Deployment Plan

Deployment Plan

\* Template EL4U5\_X86\_PVM

Description EL4U5\_X86\_PVM

**Request Details**

\* Request Name TEMPLATE\_DEPLOYMENT\_2013-01-14\_13-18-52-709

3. On the Select Target and Source page, select the following:
  - Select the target machine on which the guest virtual machine is to be created.
  - In the Source section, you can select either of the following options:
    - **Deployment Plan:** If you select the **Use Deployment Plan** check box, you can select a deployment plan (if available) with predefined configurations. Selecting a plan at this stage allows you to prepopulate the interview with previously saved values of deployment and configuration parameters.
    - **Template:** The template component to be used for this deployment. This can either be a component present in the Software Library, or at an external location. If it is stored in an external location, specify the NFS or HTTP location where the component is stored. For example, `http://myvm_server/template` or `file://myvm_server/template`.
4. Click **Next**. The Template Deployment: Deployment Configuration page appears.

**Figure 8–17 Template Deployment - Deployment Configuration**

The screenshot shows the 'CDC\_ZONE' deployment configuration page, specifically the 'Instance Details' tab. The page is divided into several sections:

- Guest VM:** Includes a 'Guest VM(s)' field set to 1, a 'Configure Each Guest VM Instance' link, 'Guest VM Naming' options (Generate Guest VM Names using prefix or Upload Guest VM Names from file), and a 'Guest VM Name' field containing 'EL4US\_X86\_PVM\_vm'.
- General:** Includes checkboxes for 'Enable High Availability' and 'Start VM After Creation', 'CPU Scheduling Priority' (Medium), 'CPU Scheduling Cap (%)' (High), 'Keymap' (en-us), and a 'Tags' field.
- VM Size:** Includes a 'VM Size' dropdown set to 'Custom', 'Max. Memory (in MB)' (1024), 'Memory (in MB)' (1024), 'Max. Number Of CPU Cores', and 'Number Of CPU Cores'.
- NIC:** A table with columns 'Name', 'Network Type', 'IP Assignment', and 'MAC Address'. It currently shows 'No Data To Display'.
- Disk:** Includes a 'Storage Type' dropdown (Default/Override), an 'Add' button, and a table with columns 'Disk Name', 'Size', 'Type', and 'Disk QoS'. The 'Total Size' is 0 GB.

At the bottom, there is a tip: 'If storage array is selected for the disk, system will select an existing physical'.

This page contains the following sections:

#### ■ Guest VM

- **Guest VM:** Number of guest virtual machines to be deployed. Click **Configure Each Guest VM Instance** if you want to configure each guest virtual machine instance individually.
- **Guest VM Naming:** You can choose either of the following options:
  - Generate Guest VM Names Using Prefix:** This string will be used as a prefix in the guest virtual machine name. When a guest virtual machine is created using the Enterprise Manager, guest virtual machine names will be prefixed by the virtual machine name prefix you provide here and a number.
  - Upload Guest VM Names from File:** If you select this option, you can Browse and upload a file with user defined guest virtual machine names.

#### ■ General Configuration

- **Enable High Availability:** If you want to enable high availability for this guest virtual machine, select this check box. If this option is enabled, the virtual server is shut down or restarted, the guest virtual machine is migrated to or restarted on another virtual server.

---

**Note:** High availability must be enabled both on the virtual server pool and on the guest virtual machine. If it is not enabled on both, high availability is disabled.

---

- **Start VM After Creation:** Specify whether the guest virtual machine should be started automatically after it is created.

- **CPU Scheduling Priority:** Specify the priority of a guest virtual machine to access the physical CPUs, while competing with other guest virtual machines on the same virtual server. Higher the priority, higher is the stake of the guest virtual machine in claiming CPU cycles.
  - **CPU Scheduling Cap (%):** Restricts the amount of physical CPU that a guest virtual machine can use. Use this to constrain guest virtual machine resource allocation.
  - **Keymap:** The keyboard mapping that will be used by the guest virtual machine being cloned.
  - **Tags:** Specify one or more tags for the virtual machine. These tags are used to logical group the virtual machines and can be used to search for one or more virtual machines that meet a certain criteria.
  - **Memory (in MB):** The amount of memory to be allocated to the guest virtual machine. The virtual server should have sufficient memory to run the guest virtual machine. Otherwise, the guest virtual machine will be created, but will not be started.
  - **CPU Cores:** The number of CPU cores to assign to this guest virtual machine. This will show up as virtual CPU (VCPU) configuration parameter for the guest virtual machine.
  - **VM Size:** The amount of memory to be allocated to the guest virtual machine. The virtual server should have sufficient memory to run the guest virtual machine. Otherwise, the guest virtual machine will be created, but will not be started. The size of the guest virtual machine (memory and CPU values). Select the VM Size which can be Custom or Medium. If you select Custom, specify the Memory size and the number of CPUs.
  - **NIC:** Specify the network interfaces for this guest virtual machine. The minimum value is 1 and maximum is 8. Click **Add** to add a network interface and specify the following details:
    - **Name:** Enter a name for the network resource.
    - **Network Type:** Specify if the network type is Internet Routable, Non Internet Routable, or a predefined set of IP addresses.
    - **IP Assignment:** Select the protocol to be used to assign the IP address.
  - **Disk:**

The list of available disks appears in the Disk section. Click **Add** to add an additional storage. In the Add Storage Device window, enter the name of the disk and its size. Indicate if it is a shared disk and specify the priority of a guest virtual machine to access the physical CPUs, while competing with other guest virtual machines on the same virtual server. Higher the priority, higher is the stake of the guest virtual machine in claiming CPU cycles.
5. Click **Next** and **Review** the details entered so far.
  6. **Review** the details entered so far. You can save as a plan that can be used later. Click **Save as Plan**, specify a name, and description for the plan. Select the folder in which the deployment plan is to be saved and click **OK**. If you do not wish to save the deployment procedure, click **Submit** to submit the job. From the Enterprise menu, select **Job**, then select **Activity** to view the Job Activity page. This page shows the details of the submitted job. You can click on the Name link to view more details.

To view details of the deployment request, from the Enterprise menu, select Cloud, then select Infrastructure Request Dashboard. See [Section 8.9, "Viewing the Infrastructure Request Dashboard"](#) for details.

## 8.8.6 Provisioning a Guest Virtual Machine Using an ISO Image

ISO images are imported into Enterprise Manager from installation media. You can create one or more virtual machines by deploying an ISO image.

### Prerequisites

- Enterprise Manager Cloud Control is set up to manage the cloud environment.
- The ISO image is available to the zone and you have sufficient privileges to deploy the image.
- There are sufficient resources available in the zone for successful deployment.
- The Domain Type of the guest virtual machine being deployed must be set to HVM (hardware virtualized).

Follow this process to deploy an ISO image:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Right click on a virtual server pool target and select **Deploy ISO**.

**Figure 8–18 ISO Deployment: Select Target and Source**

The screenshot shows the 'Pool1' deployment configuration page. The page has a breadcrumb trail: 'Select Target and Source' > 'Deployment Configuration' > 'Schedule' > 'Review'. The main title is 'ISO Deployment : Select Target and Source'. There are navigation buttons: 'Back', 'Step 1 of 4', 'Next', 'Save As Plan...', 'Review', 'Submit', and 'Cancel'. The 'Destination' section is expanded, showing 'Select the destination target on which the Guest VMs are to be deployed.' Below this, it says 'Target: Pool1' and 'Target Type: Oracle VM Server Pool'. The 'Source' section is also expanded, showing 'Plans allow you to record inputs and use them later while performing deployments using standard source images and values. Selecting a plan at this stage allows you to pre-populate the interview with previously saved values of deployment and configuration parameters. Select an existing plan below, or skip the plan selection to continue with your deployment.' Below this text is a checkbox for 'Use Deployment Plan'. Underneath, there is a 'Deployment Plan' dropdown menu with a search icon. Below the dropdown, there is an entry for 'ISO' with a search icon. The 'Request Details' section at the bottom shows a 'Request Name' field with the value 'ISO\_DEPLOYMENT\_2011-08-25\_09-48-19-689'.

3. On the Select Target and Source page, the server pool target on which the ISO image is to be deployed appears. In the Destination section, you can select either of the following:
  - **Deployment Plan:** If you select the Use Deployment Plan check box, you can select a deployment plan (if available) with predefined configurations. Selecting a plan at this stage allows you to pre-populate the interview with previously saved values of deployment and configuration parameters.
  - **ISO:** The ISO image to be used for this deployment. This can either be a component present in the Software Library, or at an external location. If it is stored in an external location, specify the NFS or HTTP location where the component is stored. For example, `http://myvm_server/image` or `file://myvm_server/image`.

4. Enter the name of the request and click **Next** to continue. The ISO Deployment: Deployment Configuration page appears.
5. This page contains the following sections:
  - **Guest VM**
    - **Guest VM:** Number of guest virtual machines to be deployed. Click **Configure Each Guest VM Instance** if you want to configure each guest virtual machine instance individually.
    - **Guest VM Naming:** You can choose either of the following options:
      - Generate Guest VM Names Using Prefix:** This string will be used as a prefix in the guest virtual machine name. When a guest virtual machine is created using the Enterprise Manager, guest virtual machine names will be prefixed by the virtual machine name prefix you provide here and a number.
      - Upload Guest VM Names from File:** If you select this option, you can Browse and upload a file with user defined guest virtual machine names.
  - **General Configuration**
    - **Enable High Availability:** If you want to enable high availability for this guest virtual machine, select this check box. If this option is enabled, the virtual server is shut down or restarted, the guest virtual machine is migrated to or restarted on another virtual server.

---

**Note:** High availability must be enabled both on the virtual server pool and on the guest virtual machine. If it is not enabled on both, high availability is disabled.

---

- **Start VM After Creation:** Specify whether the guest virtual machine should be started automatically after it is created.
  - **CPU Scheduling Priority:** Specify the priority of a guest virtual machine to access the physical CPUs, while competing with other guest virtual machines on the same virtual server. Higher the priority, higher is the stake of the guest virtual machine in claiming CPU cycles.
  - **CPU Scheduling Cap (%):** Restricts the amount of physical CPU that a guest virtual machine can use. Use this to constrain guest virtual machine resource allocation.
- **VM Size:** The amount of memory to be allocated to the guest virtual machine. The virtual server should have sufficient memory to run the guest virtual machine. Otherwise, the guest virtual machine will be created, but will not be started. The size of the guest virtual machine (memory and CPU values). Select the VM Size which can be Custom or Medium. If you select Custom, specify the Memory size and the number of CPUs.
  - **NIC:** Specify the network interfaces for this guest virtual machine. The minimum value is 1 and maximum is 8. Click **Add** to add a network interface and specify the following details:
    - **Name:** Enter a name for the network resource.
    - **Network Type:** Specify if the network type is Internet Routable, Non Internet Routable, or a predefined set of IP addresses.



- **IP Assignment:** Select the protocol to be used to assign the IP address.
- **Disk:**

The list of available disks appears in the Disk section. Click **Add** to add an additional storage. In the Add Storage Device window, enter the name of the disk and its size. Indicate if it is a shared disk and specify the priority of a guest virtual machine to access the physical CPUs, while competing with other guest virtual machines on the same virtual server. Higher the priority, higher is the stake of the guest virtual machine in claiming CPU cycles.
- 6. Click **Next**. In the Schedule page, specify when the ISO image is to be deployed and click **Next**.
- 7. **Review** the details entered so far. You can save as a plan that can be used later. Click **Save as Plan**, specify a name, and description for the plan. Select the folder in which the deployment plan is to be saved and click **OK**. If you do not wish to save the deployment procedure, click **Submit** to submit the plan and complete the process.

### 8.8.7 Creating PXE Bootable Guest Virtual Machines

One of the key requirements of provisioning is the hardware server's ability to boot over the network instead of a diskette or CD-ROM. There are several ways computers can boot over a network, and Preboot Execution Environment (PXE) is one of them. PXE is an open industry standard supported by a number of hardware and software vendors. A detailed document on PXE specification can be found at <http://www.pix.net/software/pxeboot/archive/pxespec.pdf>.

You can create a network bootable (PXE boot) virtual machine that has the minimum configuration information, then start the virtual machine through Preboot Execution Environment (PXE) over a network later to install the guest operating system.

#### Prerequisites

- Enterprise Manager Cloud Control is set up to manage the cloud environment.
- There are sufficient resources available in the zone for successful deployment.

Follow this process to deploy an ISO image:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Select a Virtual Server Pool target. Select **Deploy PXE** from the Virtual Server Pool menu.
3. On the Select Target and Source page, select the following:
  - Select the target machine on which the guest virtual machines are to be deployed.
  - Select the deployment plan to be used while creating PXE bootable guest virtual machines. The Deployment Plan contains pre-configured values and configuration parameters. This is an optional field.
4. Click **Next**. In the General Configuration page, enter the following details:
  - Name of the guest virtual machine to be deployed. If you are deploying more than one guest virtual machine, you can use a prefix or upload a file with user defined guest virtual machine names.
  - Number of guest virtual machines to be deployed. Click **Configure Each Guest VM Instance** if you want to configure each guest virtual machine instance individually.

- Select the VM Size which can be Custom or Medium. If you select Custom, specify the Memory size and the number of CPUs.
  - Specify the priority of a guest virtual machine to access the physical CPUs, while competing with other guest virtual machines on the same virtual server. Higher the priority, higher is the stake of the guest virtual machine in claiming CPU cycles.
  - Specify the cap to restrict the amount of physical CPU that a guest virtual machine can use.
  - CPU Cores: The number of CPU cores to assign to this guest virtual machine. This will show up as virtual CPU (VCPU) configuration parameter for the guest virtual machine.
  - Keymap: The keyboard mapping to be used for the guest virtual machine.
  - Memory (in MB): The amount of memory to be allocated to the guest virtual machine. The virtual server should have sufficient memory to run the guest virtual machine. Otherwise, the guest virtual machine will be created, but will not be started.
  - CPU Cores: The number of CPU cores to assign to this guest virtual machine. This will show up as virtual CPU (VCPU) configuration parameter for the guest virtual machine.
  - In the NIC Configuration section, you can specify the number of virtual network interfaces for this guest virtual machine. The minimum value is 1 and the maximum value is 8. You can add one or more network interfaces or edit existing ones.
  - Enterprise Manager Cloud Control provides a VNC console to the newly created guest virtual machine. This is the VNC password to access the VNC console.
  - Select a Domain Type for the virtual machine:
    - **xen\_hvm** (hardware virtualization, or fully virtualized): Enables you to select an ISO file from which to create the virtual machine.
    - **xen\_pvm** (para-virtualized): Enables you to select a location for the mounted ISO file from which to create the virtual machine. Before you create the virtual machine using the paravirtualized method, mount the ISO file on an NFS share, or HTTP or FTP server. The ISO file must be available in the relevant repository.
    - **xen\_hvm\_pv\_drivers**: Identical to xen\_hvm but with additional paravirtualized drivers for improved performance of the virtual machine. This Domain Type is used to run Windows guest operating systems with an acceptable performance level.
  - Select the **Enable High Availability** option if you want the guest virtual machine to run automatically on a different virtual server if the earlier virtual server shuts down.
  - Select the **Start VM After Creation** check box to indicate whether the machine should automatically be started after it is created.
  - The list of available disks appears in the Disk section. Click **Add** to add an additional storage.
5. Click **Next**. In the Schedule page, specify when the PXE bootable guest virtual machines are to be created and click **Next**.

6. Review the details entered so far and click **Submit** to create the PXE bootable guest virtual machines.

## 8.9 Viewing the Infrastructure Request Dashboard

The SSA Administrator can use the Request Dashboard to track the status of requests and reservations. The SSA administrator can view the requests across all zones for all users. The requests the SSA user can view depends on his zone and his instance level privileges.

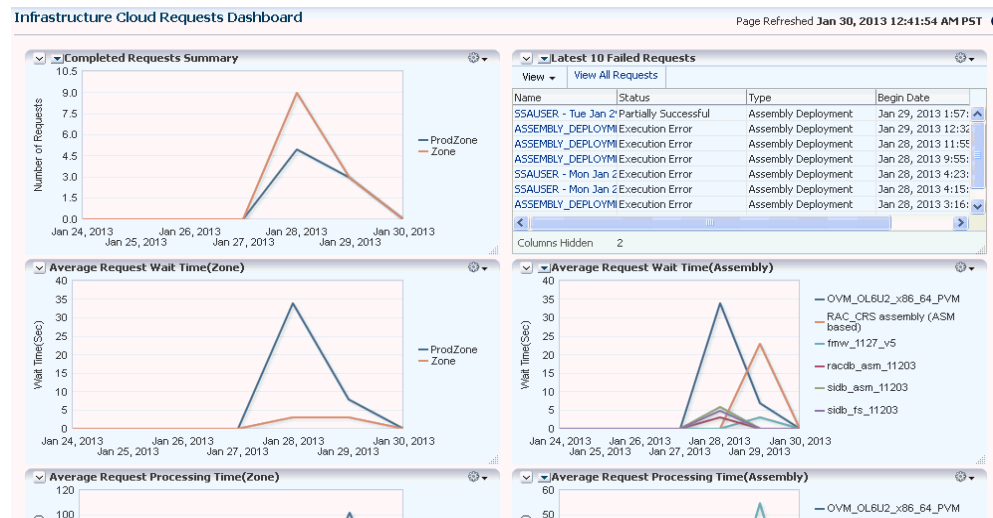
### Prerequisites

- Enterprise Manager Cloud Control is set up to manage the cloud environment.
- Appropriate self service users have been defined in the LDAP directory.
- Requests and reservations for resources have been submitted.

To view the Infrastructure Request Dashboard, follow these steps:

1. From the Enterprise menu, select **Cloud**, then select **Infrastructure Request Dashboard**. The Infrastructure Cloud Requests Dashboard appears.

**Figure 8–19 Infrastructure Requests Dashboard**



2. The following charts are displayed:
  - Latest 10 Failed and Pending Requests
  - Oldest 10 Running Requests
  - Request Summary - Failed, Successful, and Partially Successful Requests
  - Average Request Wait Time (Zone)
  - Average Request Wait Time (Assemblies and Templates)
  - Average Request Processing Time (Zone)
  - Average Request Processing Time (Assemblies and Templates)
3. Click on the **View All Requests** link to go to the All Cloud Requests page.
4. In this page, you can search for one of the following:

- Request Name: The name of the request.
  - Request Type: The type of request that you want to view. This can be Assembly Deployment, Template Deployment, ISO Deployment, PXE Deployment, Clone Guest VM, Update Guest VM.
  - Requesting Application: The application making the request. You can view only Self Service requests from this dashboard.
  - Status: The status of the request such as Successful, Canceled, Scheduled, Error Ending, and so on.
5. For each request type you select, you can view the Request Name, Request Type, Created By, Status, Begin Date, End Date, Memory, CPU, VMs, and so on. You can also select a request and click **Delete** to remove the request.
  6. Click on a **Request Name** link to drill down to the Request Details page.

### 8.9.1 Viewing Request Details

The Request Details page shows the details of the selected request. The name of the request, type, the OVM Manager associated with the request, status, date on which it was created, and so on are displayed.

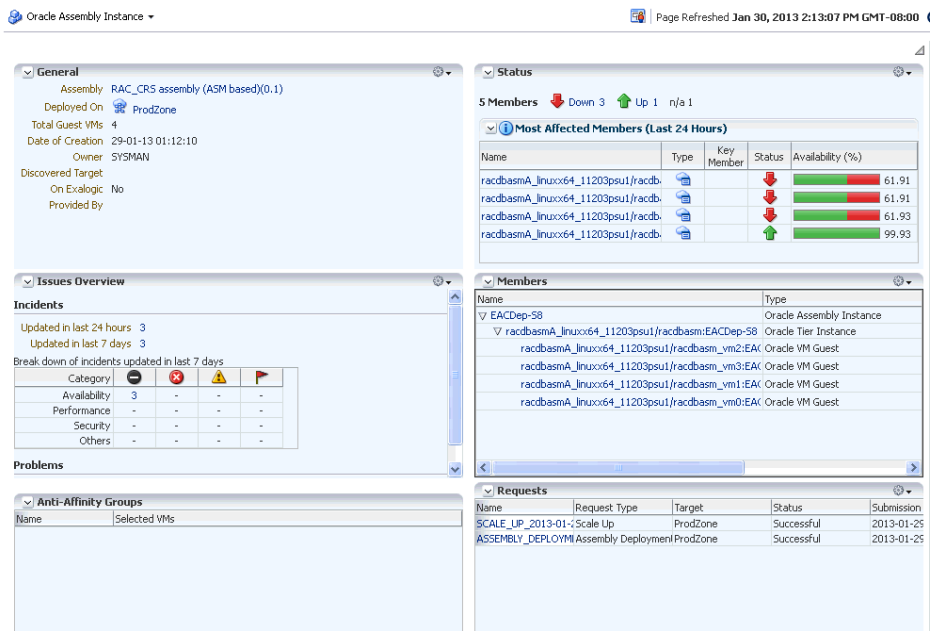
- **General:** This section displays the name of the request, request type, target name, and other details.
- **Deployment:** This section shows the name of the requests and the targets that have been deployed. Click on request link to drill down to the Job Details page. Click on a target link to drill down to the Home page of the target.

### 8.9.2 Viewing the Assembly / Tier Instance Page

To view the Home page for the assembly or tier instance, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Request Dashboard**.
2. The Infrastructure Request Dashboard page appears. Click the **All Requests** link in the Requests table.
3. In the All Requests page, click on an assembly deployment request name in the table. The Assembly Deployment page appears. Click on the **Deployment** tab. The assembly instance that has been deployed is displayed.
4. Click on the Assembly Instance link. The following page appears.

Figure 8–20 Assembly Instance



- **General:** This section shows the name of the assembly, the target on which it was deployed, the total number of Guest VMs, the date on which it was created, and the owner of this assembly instance.
- **Status:** The status of the most affected members over the last 24 hours is displayed. Click on a Name link to drill down to the target Home page.x
- **Issues Overview:** Shows the various incidents and policy violations that have occurred in the last 24 hours.
- **Affinity Relations:** The rules that have been defined for the assembly or tier are displayed here. For example, you can define rules to specify whether two guest VMs can exist together on the same server or server pool, rules to ensure that two guest VMs must exist on the same server, and so on.
- **Members:** The members in this instance, their status, and the installed software is displayed.
- **Requests:** The name of the request specified during deployment. Click on the link to navigate to the Deployment page.



# Part IV

---

## Setting Up and Using Platform as a Service

This section describes the steps required to enable Platform as a Service, which encompasses the Database as a Service (DBaaS) and Middleware as a Service (MWaaS) Cloud models.

It contains the following chapters:

- [Chapter 9, "Setting Up the PaaS Infrastructure"](#)
- [Chapter 10, "Monitoring the PaaS Infrastructure"](#)
- [Chapter 11, "Setting Up a DBaaS Cloud"](#)
- [Chapter 12, "Setting Up the DBaaS Self Service Portal"](#)
- [Chapter 13, "Using the DBaaS Self Service Portal"](#)
- [Chapter 14, "Monitoring a DBaaS Cloud"](#)
- [Chapter 15, "Setting Up an MWaaS Cloud"](#)
- [Chapter 16, "Setting Up the MWaaS Self Service Portal"](#)
- [Chapter 17, "Using the MWaaS Self Service Portal"](#)
- [Chapter 18, "Monitoring a MWaaS Cloud"](#)





## Setting Up the PaaS Infrastructure

The PaaS model allows users to create platforms onto which consumers can deploy their own applications. A platform resource is typically comprised of a host, operating system, and application server, and may be virtualized. The platform may also include a database and even multiple hosts in a cluster.

This chapter covers the common self service operations that need to be performed to enable Platform as a Service (PaaS), which encompasses Database as a Service (DBaaS) using Oracle Single Instance, and RAC or Middleware as a Service (MWaaS) using WebLogic Server. These common tasks must be completed before setting up either DBaaS or MWaaS.

The chapter contains the following sections:

- [Getting Started](#)
- [Adding Hosts](#)
- [Creating a PaaS Infrastructure Zone](#)
- [Creating a PaaS Pool](#)

### 9.1 Getting Started

This section helps you get started by providing an overview of the steps involved in setting up Platform as a Service (PaaS).

**Table 9–1** *Getting Started*

| Step | Task                                                                                                                                               | Role                |
|------|----------------------------------------------------------------------------------------------------------------------------------------------------|---------------------|
| 1    | Configure the storage location for the Software Library. See <a href="#">Section 3.1, "Setting Up the Software Library"</a>                        | Cloud Administrator |
| 2    | Deploy the necessary plug-ins. See <a href="#">Section 3.3, "Deploying the Required Plug-ins"</a>                                                  | Cloud Administrator |
| 3    | Define roles for administrators and self service users. See <a href="#">Section 3.4, "Defining Roles and Assigning Users"</a> .                    | Super Administrator |
| 4    | Install the Management Agent on unmanaged hosts so that they can be monitored by Enterprise Manager. <a href="#">Section 9.2, "Adding Hosts"</a> . | Cloud Administrator |
| 5    | Create PaaS Infrastructure zones. See <a href="#">Section 9.3, "Creating a PaaS Infrastructure Zone"</a> .                                         | Cloud Administrator |

## 9.2 Adding Hosts

Oracle Management Agent (Management Agent) is one of the core components of Enterprise Manager Cloud Control. It works in conjunction with the plug-ins to monitor the targets running on a managed host. You must install Oracle Management Agents on your unmanaged hosts to monitor them in Enterprise Manager.

---

**Note:**

- If the Management Agent does not have Fusion Middleware 12.1.0.3 GC FMW Plug-in for monitoring, the following options will not work:
    - Create and Delete Data Sources
    - Start and Stop Applications
    - View List of Data Sources
    - View Application URLs
  - If you are deploying the Management Agent on several servers at a time, the total time required will increase exponentially.
  - Verify if the Oracle Fusion Middleware Plug-in has been successfully deployed. From the **Setup** menu, select **Extensibility**, then select **Plug-ins**. Select **Oracle Fusion Middleware**, then select **Deploy On**, then select **Management Agent**. Add the required agent on the target hosts and click **Deploy**.
- 

To install the Management Agent on an unmanaged server, follow these steps:


1. Log in to Enterprise Manager as a user with the `EM_CLOUD_ADMINISTRATOR` role.
2. From the **Setup** menu, select **Add Target**, then select **Add Target Manually**.
3. Select the **Add Host Targets** option and click **Add Host**.
4. Add the host names and select the platform. Check the Self Update console, if the Agent software for the desired platform is listed as not available, click **Next**.
5. Provide the Agent Install Location, click in the Instance directory field to auto-generate the path, create a named credential for the agent user (provide access to root user either by using `sudo` or `pbrun`), and clear the value in the port field (this will enable automatic selection of the port), click **Next**.
6. Review all entered values, and click **Deploy Agent**.
7. Track the progress of agent deployment on the Add Host Status page. The agent deployment takes between 5-10 minutes.
8. Select **Targets**, then select **Hosts** to navigate to the Hosts page. Confirm if all the hosts are listed on the page.

## 9.3 Creating a PaaS Infrastructure Zone

A PaaS Infrastructure Zone can contain a group of resources such as OVM zones or hosts. Each zone can contain homogeneous resources of only one type (hosts or OVM zones). Each resource in a zone represents a location at which a service instance is to be deployed.

The first step in building a PaaS cloud is to create at least one PaaS Infrastructure Zone, which is done by selecting one more Enterprise Manager host targets and identifying the SSA enabled roles that users will need to in order to request resources in this zone.

**Figure 9–1 PaaS Infrastructure Zone**



### Prerequisites

Before you create a PaaS Infrastructure Zone, you must ensure that the following prerequisites are met:

- The EM\_CLOUD\_ADMINISTRATOR, EM\_SSA\_ADMINISTRATOR, and EM\_SSA\_USER roles must have been created.
- The Software Library must have been configured.

---

**Note:** If you are using an OMS Shared System location in the Software Library, the credentials owner must grant privileges to the `CLOUD_ENGINE_USER`. Since the `CLOUD_ENGINE_USER` is a hidden user account, the owner of the named credential cannot grant **View** privileges from the Enterprise Manager Console. To address this issue, (especially on a Windows host where OMS Agent Filesystem is the recommended approach for setting up Software Library) you need to run the following EMCLI commands:

```
emcli login -username=<username> -password =<password>

emcli grant_privs -name=CLOUD_ENGINE_USER -privileges="GET_CREDENTIAL; CRED_NAME=<namedcred>:CRED_OWNER=<loginusername>"
```

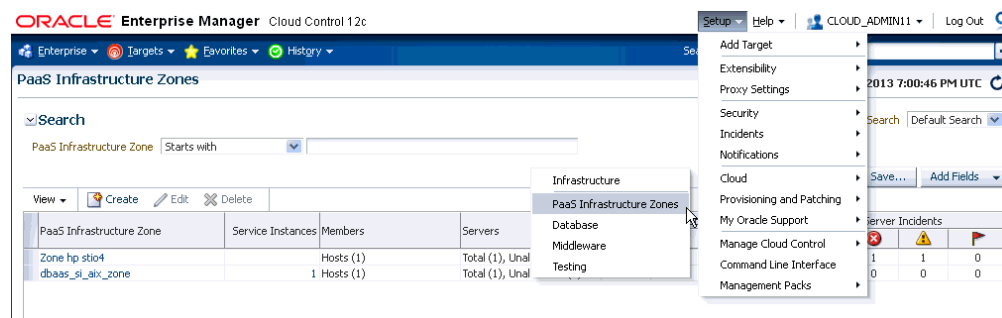
---

- You must log in as a user with `EM_CLOUD_ADMINISTRATOR` privileges to create a PaaS Infrastructure Zone.

To create a PaaS Infrastructure Zone, follow these steps:

1. From the **Setup** menu, select **Cloud**, then select **PaaS Infrastructure Zone**.

**Figure 9–2 Create PaaS Infrastructure Zone: Navigation**



2. In the PaaS Infrastructure Zones page, click **Create**. The Create PaaS Infrastructure Zone: General page appears.

---

**Note:** You can also launch the Create PaaS Infrastructure Zone wizard from the Enterprise menu. From the **Enterprise** menu, select **Cloud**, then select **Middleware and Database Home**. From the **Middleware and Database Cloud** menu, select **Create PaaS Infrastructure Zone**.

---

**Figure 9–3 Create PaaS Infrastructure Zone: General**

**PaaS Infrastructure Zones**

General Members Roles Review

Create PaaS Infrastructure Zone : General

Back Step 1 of 4 Next Cancel

**General**

\* Name

Description

Member Type ☒ Host ☐ OVM Zone

**Placement Policy Constraints Per Host**

Placement Policy Constraints allow the self service administrator to set maximum ceilings for any host in the PaaS Infrastructure zone. These constraints are used during deployment to filter out hosts that are overburdened. For example, in a production environment, the CPU utilization may be limited to 80 percent, whereas in a development environment, the CPU utilization can be up to 95 percent.

\* Maximum CPU Utilization (%)

\* Maximum Memory Allocation (%)

☒ **TIP** The service instance will be provisioned on the first host that meets the placement constraints.

3. Enter a name and description for the zone.
4. Select the Member Type which can be:
  - **Host:** If you select this type, only host targets can be part of this PaaS Infrastructure zone. You must specify appropriate values for the placement policy constraints based on the resources available to the members. The EM\_CLOUD\_ADMINISTRATOR can use the placement policy constraints to set maximum ceilings for any host in the zone. Placement policy values are used during deployment to filter out hosts that are overburdened. You can define the following placement policy constraints for each host:
    - **Maximum CPU Utilization:** The average CPU utilization percentage for the past 7 days.
    - **Maximum Memory Allocation:** The average memory allocation for the past 7 days.

Hosts exceeding these thresholds will not be considered for provisioning operations.
  - **OVM Zone:** An Oracle VM Zone is a collection of virtual server pools that share storage. If you select this type, only Oracle VM Zone targets can be part of this PaaS Infrastructure zone.
5. Click **Next**. The Create PaaS Infrastructure Zone page appears. If you selected the Member Type as **Host**, the following page appears:

**Figure 9–4 Create PaaS Infrastructure Zone: Targets**

**PaaS Infrastructure Zones**

General Targets Roles Review

Create PaaS Infrastructure Zone : Targets

Back Step 2 of 4 Next Cancel

**Credentials**

Credentials to be used for provisioning in this PaaS infrastructure zone. The credentials should be the same for all hosts. You can only use Named Credentials that you own.

\* Named Credential to use for this PaaS infrastructure zone

**Targets**

Add one or more targets to the PaaS infrastructure zone.

| Name | Description | Type | Status | VM Guest | Oracle VM Zone |
|------|-------------|------|--------|----------|----------------|
|      |             | Host | Up     | n/a      | n/a            |
|      |             | Host | Up     | n/a      | n/a            |

6. Specify the Named Credentials to be used for provisioning all the targets in this zone. Click the + icon to add new named credentials. The credentials you specify must be global credentials and valid on all host members present in the zone. Unless you explicitly click **Verify Credential**, the credentials are not verified for all hosts in zone.
7. In the Targets region, click **Add** to add one or more host targets to the zone. A target can be a host or an OVM Zone. For more information on OVM zones, see [Section 4.15, "Creating a Zone"](#) for details.

---

**Note:** The host that you select cannot be used in more than one PaaS Infrastructure Zone.

---

8. If you selected the Member Type as **OVM Zone**, click **Add** to add one or more Oracle VM Zone targets to the PaaS Infrastructure Zone.
9. Click **Next**. The Create PaaS Infrastructure Zone: Roles page appears.

**Figure 9–5 Create PaaS Infrastructure Zone: Roles**

PaaS Infrastructure Zones

General Targets Roles Review

Create PaaS Infrastructure Zone : Roles

Back Step 3 of 4 Next Cancel

**Roles**

A PaaS Infrastructure zone can be made available to a restricted set of users via the use of roles. Select the roles that can access this PaaS infrastructure zone.

+ Add... - Remove...

| Role Name  | Role Description  |
|------------|-------------------|
| SSACOROLE1 | Test plsql insert |
| SSACOROLE2 | Test plsql insert |

10. A PaaS Infrastructure Zone can be made available to a restricted set of self service users with the EM\_SSA\_USER role. Click **Add** to add and select the roles that can access this PaaS Infrastructure Zone.
11. Click **Next**. The Create PaaS Infrastructure Zone: Review page appears. Review the details of the zone being created and click **Submit** to create the zone. You will return to the Database and Cloud Home page.

### 9.3.1 Editing a PaaS Infrastructure Zone

To edit a PaaS Infrastructure Zone, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Middleware and Database Cloud Home**.
2. Click on the **PaaS Infrastructure Zones** link in the General region to drill down to the PaaS Infrastructure Zones page.
3. Select a PaaS Infrastructure Zone from the list and click **Edit**.

---

**Note:** If you do not have access to or own the Named Credentials that have been defined for the PaaS Infrastructure Zone, you must specify new credentials.

---

4. Follow the steps in the wizard to edit the necessary parameters and click **Submit** to accept the changes.

### 9.3.2 Deleting a PaaS Infrastructure Zone

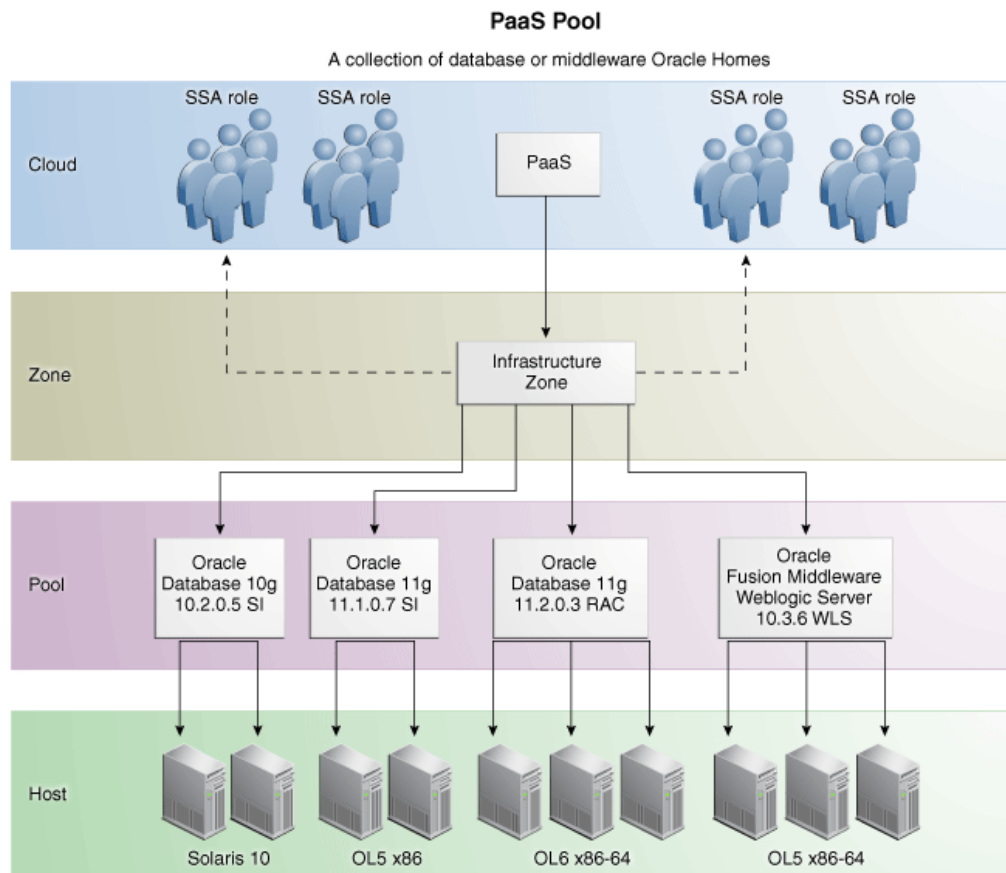
To delete a PaaS Infrastructure Zone, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Middleware and Database Cloud Home**.
2. Click on the PaaS Infrastructure Zone link to drill down to the PaaS Infrastructure Zones page.
3. Select a PaaS Infrastructure Zone from the list and click **Delete**.
4. A confirmation message is displayed. Click **Delete PaaS Infrastructure Zone** to delete the zone.

## 9.4 Creating a PaaS Pool

The second step in building a PaaS cloud is to create software pools (for example, database or middleware Oracle Homes that exist on hosts within the PaaS Infrastructure Zone, or database instances for schema as a service). The figure below shows an example of a PaaS pool.

**Figure 9–6 PaaS Pool**



A PaaS pool must be a homogeneous collection of Oracle Homes of the same platform, type, and version. For example, Solaris 10 Oracle Database single instance version

10.2.0.5 or Oracle Enterprise Linux 5 x86-64 WebLogic Server version 10.3.6. You can create a Database Pool or a Middleware Pool.

- To create a Database Pool, see [Section 12.5, "Creating a Database Pool"](#).
- To create a Middleware Pool, see [Section 16.2, "Creating a Middleware Pool"](#).



---

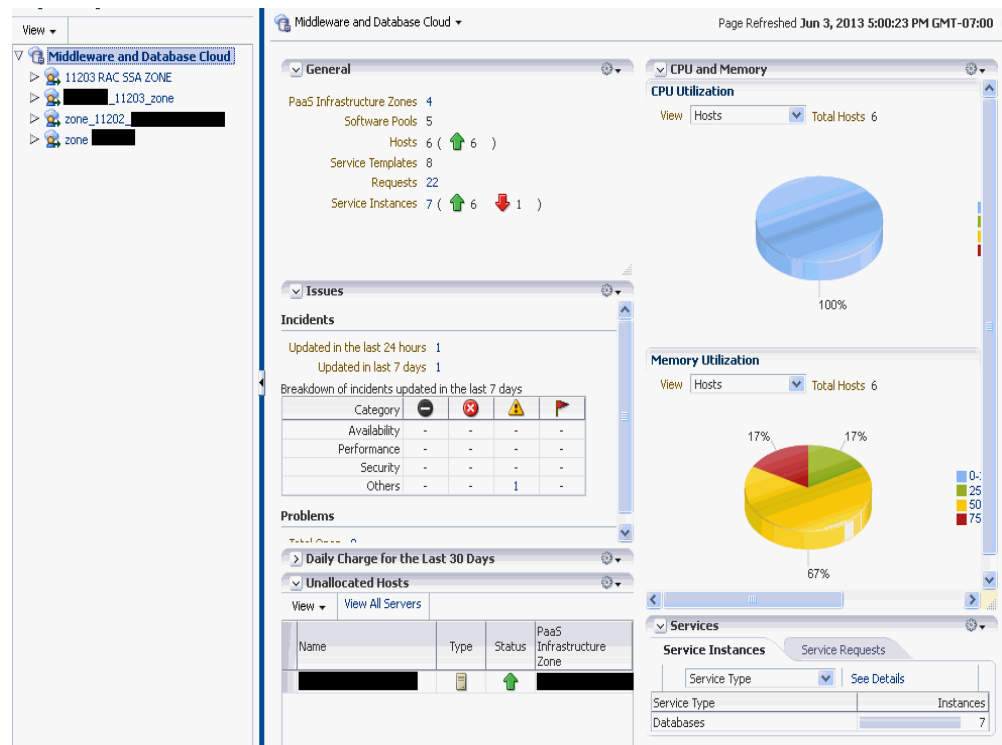
## Monitoring the PaaS Infrastructure

This chapter describes the monitoring operations that can be performed on DBaaS or MWaaS clouds. It contains the following sections:

- [Viewing the Middleware and Database Cloud Home Page](#)
- [Viewing the Middleware and Database Request Dashboard Page](#)
- [Viewing All Requests Page](#)
- [Viewing the Service Instances Page](#)
- [Viewing the PaaS Infrastructure Zones Page](#)
- [Viewing the PaaS Infrastructure Zone Home Page](#)

### 10.1 Viewing the Middleware and Database Cloud Home Page

After the `SSA_ADMINISTRATOR` has completed DBaaS and MWaaS setup and one or more SSA users have submitted requests, a unified view of the DBaaS and MWaaS service types can be seen on the Middleware and Database Cloud Home page. To view the Middleware and Database Cloud Home page, from the **Enterprise** menu, select **Cloud**, then select **Middleware and Database Cloud**.

**Figure 10–1 Middleware and Database Cloud Home**

It contains the following regions:

- **General:** This region contains the following:
  - **PaaS Infrastructure Zone:** The total number of PaaS Infrastructure zones that have been created is displayed. Click on the link to drill down to the PaaS Infrastructure Zones page.
  - **Software Pools:** The total number of software pools that have been defined across MWaaS and DBaaS. The software pools are applicable only to PaaS Infrastructure zones which contain the host member type. They are not applicable to OVM zone member types.
  - **Oracle VM Zones:** The Oracle VM Zones associated with the PaaS Infrastructure zones.
  - **Hosts:** The physical hosts associated with the software pools.
  - **Service Templates:** This total number of service templates that have been defined.
  - **Requests:** The total number of requests for MWaaS and DBaaS. Click on the link to drill down to the All Requests page.
  - **Service Instances:** The total number of members that are up or down. Click on the link to drill down to the Service Instances page.
- **Issues:** The availability, performance, security, and other incidents across all the zones.
- **Daily Charge for the Last 30 Days:** Shows rolled up chargeback data across all service types.

- **CPU and Memory:** The average CPU and memory titillations by all servers (virtual servers and physical hosts) across all zones.
- **Unallocated Hosts:** The physical hosts that are not associated with any software pool are listed here. You can see the name of the host, type, and the status. Click **View All Servers** to drill down to the Servers page.
- **Services:** This region displays the following:
  - **Service Instances:** The total number of instances by service type, zone, service template, status or owner. Select a option from the View drop down list to see the rollup inventory for the type. Click See Details to view the Service Instance Details page. Click on the bar graph next to a service type to drill down to the detailed view of the selected service.
  - **Service Requests:** The Service Requests table lists requests by request type, zone, service type, service template, status, or owner. You can choose to view all the requests or for a specific period by selecting the appropriate option from the Show drop down list. Click See Details to view the Request Details page. Click on the bar graph next to a specific request to drill down to the detailed view of the selected request.

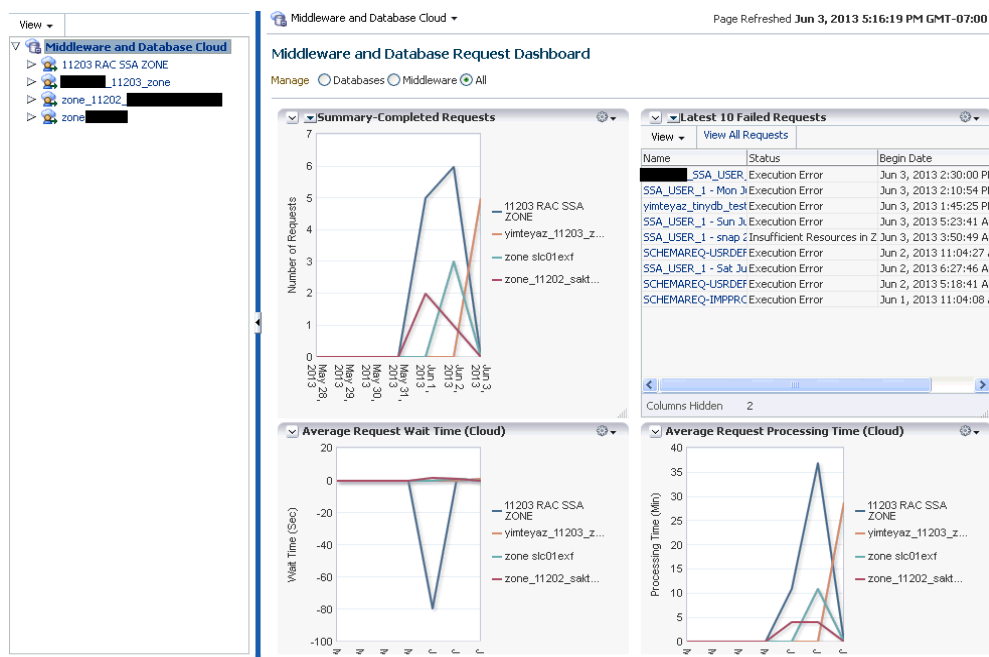
From the Middleware and Database Cloud menu, you can do the following:

- **Create PaaS Infrastructure Zone:** You can launch the wizard to create the zone. See [Section 9.3, "Creating a PaaS Infrastructure Zone"](#) for details.
- **Service Instances:** You can view detailed information of all the members. See [Section 10.4, "Viewing the Service Instances Page"](#) for details.
- **Request Dashboard:** You can view all the database and middleware requests. See [Section 10.2, "Viewing the Middleware and Database Request Dashboard Page"](#) for details.

## 10.2 Viewing the Middleware and Database Request Dashboard Page

The SSA Administrator can use the Request Dashboard to track the status of all database and middleware requests. The SSA administrator can view the requests across all zones for all users. The requests the SSA user can view depends on his zone and his instance level privileges. You can choose to view Middleware, Databases, or All requests by clicking the appropriate radio button.

To view the Middleware and Database Request Dashboard, from the **Enterprise** menu, select **Cloud**, then select **Middleware and Database Request Dashboard**.

**Figure 10–2 Middleware and Database Request Dashboard**

The following charts are displayed:

- **Summary:** You can view a summary of all Completed Requests, Failed Requests, Successful Requests, and Partially Successful Requests.
- **Latest:** You can view a list of the Latest 10 Failed Requests, Oldest 10 Running Requests, and Latest 10 Pending Requests. Click **View All Requests** to drill down to the All Requests page.
- **Average Request Wait Time (Zone):** The average request wait time across all zones for the past 7 days.
- **Average Request Processing Time (Zone):** The average request wait time across all zones for the past 7 days.

## 10.3 Viewing All Requests Page

The All Requests page shows a list of requests made so far. You can choose to view database requests, middleware requests, or both. To view the All Requests page, click the View All Requests link in the Latest 10 Failed Requests region in the Middleware and Database Request Dashboard page. For each request, the name of the request, date and time on which was created, the type of request, user who created the request, and the status of the request are displayed.

Click on the **Name** link to drill down to the Request Details page.

**Figure 10–3 Request Dashboard Details General**

Request Dashboard > All Requests > SCHEMAREQ-IMPPROFILE\_CREATE\_9\_56

**SCHEMAREQ-IMPPROFILE\_CREATE\_9\_56**

General Deployment

Request Name SCHEMAREQ-IMPPROFILE\_CREATE\_9\_56  
 Type Create Service Instance  
 Status Success  
 Description Create Database for "SCHEMAREQ-IMPPROFILE".  
 Service Type Databases  
 Deployed On 11203 RAC SSA ZONE  
 Submission Date Jun 1, 2013 4:09:57 AM PDT  
 Created By SSA\_USER\_1  
 Begin Date Jun 1, 2013 4:10:57 AM PDT  
 Last Modified By SYSTEM  
 Last Modification Date Jun 1, 2013 4:21:08 AM PDT

This page contains two tabs:

- **General:** This tab shows details of the request such as Request Name, Request Type, Start Date, End Date, Status, and other details. Click the Deployment tab.

**Figure 10–4 Request Dashboard Details Deployment**

Request Dashboard > All Requests > SCHEMAREQ-IMPPROFILE\_CREATE\_9\_56

**SCHEMAREQ-IMPPROFILE\_CREATE\_9\_56**

General Deployment

Procedure DBAAS-CREATE-#4

✓ Deployed Targets

| Name                   | Type |
|------------------------|------|
| Service_2683886D90615B |      |

- **Deployment:** This tab shows details of the service instance that was created. Click the **Procedure** or **Deployed Targets** link to navigate to the deployment or Job instance page that can be used to troubleshoot failed requests.

## 10.4 Viewing the Service Instances Page

The Service Instance page shows the list of all service instances. To view this page, from the Enterprise menu, select Cloud, then select Middleware and Database Cloud. From the Middleware and Database Cloud menu, select Service Instances.

**Figure 10–5 Service Instances**

Middleware and Database Cloud

Middleware and Database Cloud

Page Refreshed Jun 25, 2012 10:32:33 PM GMT-07:00

**Service Instances**

Search

Match: ☒ All ☐ Any

Service Instance Name: Starts with

Service Type: Equals

Search Reset Save... Add Fields

| Service Instance Name | Service Instance Type | Service Type | Status         | Creation Date                | Expiry Date                 | PaaS Infrastructure Zone | Software Pool            | Ser |
|-----------------------|-----------------------|--------------|----------------|------------------------------|-----------------------------|--------------------------|--------------------------|-----|
|                       |                       | Middleware   | Status Pending | Jun 19, 2012 9:38:46 AM PDT  | Jun 20, 2012 9:39:24 AM PDT | test paas zone           | test mw software poc Ser |     |
|                       |                       | Middleware   |                | Jun 18, 2012 12:11:40 PM PDT |                             | test paas zone           | test mw software poc Ser |     |

For each service instance, the following details are displayed:

- **Service Instance Name:** The name associated with the service instance. Click on the link to drill down to the target home page.
- **Service Instance Type:** The type of request (Middleware or Database).
- **Status:** The status of the target.
- **Creation Date:** The date on which the service was created.
- **Expiry Date:** The date on which the service will expire and no longer be available.
- **PaaS Infrastructure Zone:** The PaaS Infrastructure Zone in which this service instance was created. Click on the link to drill down to the Zone Home page.
- **Software Pool:** The software pool to which this service instance was created. Click on the link to drill down to the Software Pool Home page.
- **Service Template:** The service template that was used to create this service instance.
- **Incidents:** Shows any incidents on the virtual servers and physical hosts present in the zone.

## 10.5 Viewing the Service Template Details Page

To view the Service Template Details page, click on the Service Template Name link in the Database / Middleware Cloud Self Service Portal Setup page. The following details are displayed:

- Name and description of the service template.
- The configuration details of the database or WebLogic Server instance being provisioned.
- The PaaS Infrastructure Zone and the Software Pool in which the service template is being provisioned.
- The users (roles) that can access this service template.

## 10.6 Viewing the PaaS Infrastructure Zones Page

This page lists all the PaaS Infrastructure zones that have been created. To view this page, from the **Enterprise** menu, select **Cloud**, then select **Middleware and Database Cloud**. Click the **PaaS Infrastructure Zones** link in the General region to view this page. The name of the zone, the members, servers (total and unallocated), status, and incidents are shown for each zone.

---

---

**Note:** You can also reach this page from the **Setup** menu. Select **Cloud**, then select **PaaS Infrastructure Zones** from the menu.

---

---

**Figure 10–6 PaaS Infrastructure Zones**

Middleware and Database Cloud Page Refreshed Aug 18, 2012 12:43:13 AM GMT+00:00

**PaaS Infrastructure Zones**

Search Saved Search Default Search

PaaS Infrastructure Zone Starts with Search Reset Save... Add Fields

View Create Edit Delete

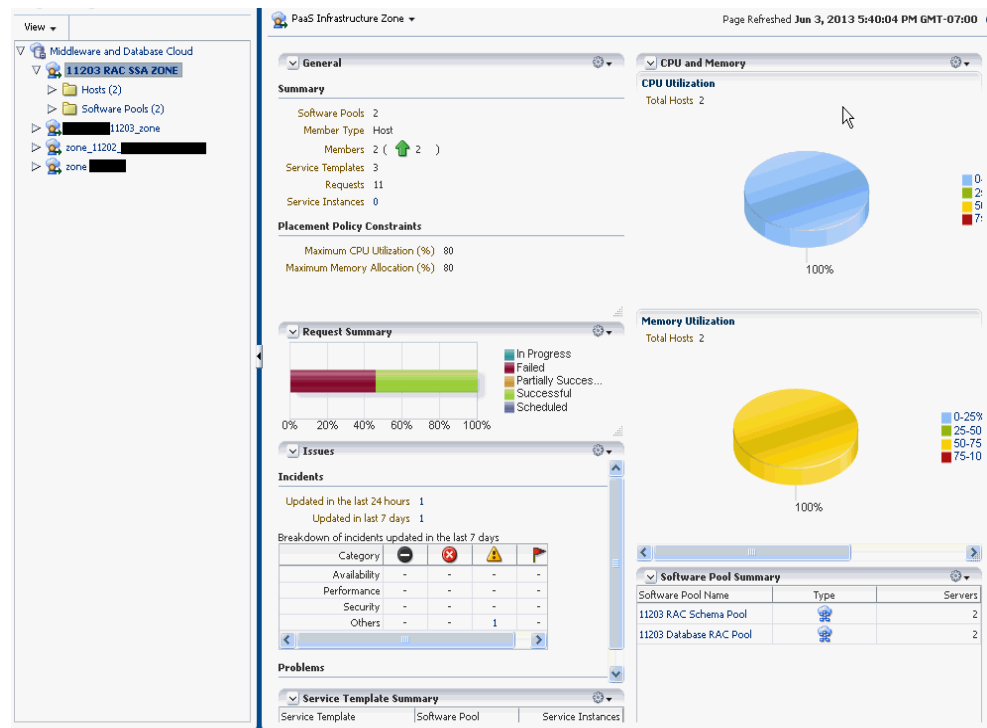
| PaaS Infrastructure Zone | Service Instances                                           | Members | Servers | Server Status Summary |   |   |   |   | Server Incidents |   |   |   |
|--------------------------|-------------------------------------------------------------|---------|---------|-----------------------|---|---|---|---|------------------|---|---|---|
| zone1                    | 2 Hosts (2), Oracle VM Zones (1 Total (2), Unallocated (0)) |         |         | 0                     | 2 | 0 | 0 | 0 | 0                | 0 | 0 | 0 |

You can do the following:

- **Create:** Click **Create** to launch the Create PaaS Infrastructure Zone wizard. See [Section 9.3, "Creating a PaaS Infrastructure Zone"](#) for details.
- **Edit:** Select a zone from the list and click **Edit** to modify the details.
- **Delete:** Select a zone from the list and click **Delete** to delete the details.
- **View:** Click on the zone name link to drill down to the PaaS Infrastructure Zone Home page.

## 10.7 Viewing the PaaS Infrastructure Zone Home Page

A zone is a logical entity that allows homogenous resources to be grouped together. These homogeneous resources consist of components (servers, storage, networks, operating systems, applications) that are required to create new services. To view this page, from the **Enterprise** menu, select **Cloud**, then select **Middleware and Database Cloud**. Click the **PaaS Infrastructure Zones** link in the General region to view this page. The PaaS Infrastructure Zones page appears. Click a PaaS Infrastructure Zone link to drill down to the PaaS Infrastructure Zone Home page.

**Figure 10–7 PaaS Infrastructure Zone Home**

The PaaS Infrastructure Zone Home page contains the following regions:

- **General:** This region lists the total number of Software Pools, Oracle VM Zones, Member Type (hosts or Oracle VM Zones), Members, Service Templates, Requests, and Service Instances in the zone. It also shows the Placement Policy Constraints that have been defined for the zone.
- **Request Summary:** Shows the percentage of requests that are in progress, failed, successful, partially successful, or scheduled.
- **CPU and Memory:** Shows bar charts indicating the CPU and memory utilization distributed across all servers (virtual servers and physical hosts) in the zone.
- **Issues:** Shows the outstanding incidents and problems for the target.
- **Software Pool Summary:** Shows the name and type of each software pool in the zone. Click on the Software Pool Name link to drill down to the Software Pool Home page. This field is applicable only for the Host Member Type (displayed in the General region).
- **Unallocated Servers:** The servers that are not associated with any software pool are listed here. You can see the name of the server, type (host or virtual server), and the status. This field is applicable only for the Host Member Type (displayed in the General region). Click **View All Servers** to drill down to the Servers page.
- **Members:** Shows the members of the zones and the member type (Oracle VM Zone or Host). Click on the Name link to drill down to the target Home page.
- **Service Template Summary:** Shows the service templates associated with the zone. The software pool associated with each service template and number of service instances in each resource pool is also displayed.



## Setting Up a DBaaS Cloud

This chapter covers the initial configuration required to enable a Database as a Service Cloud. It contains the following sections:

- [Getting Started](#)
- [Setting Up Credentials for Provisioning](#)
- [Provisioning Database Software](#)
- [Deploying the Database](#)
- [Configuring the Oracle Listener](#)
- [Registering and Managing Storage Servers](#)

### 11.1 Getting Started

This section helps you get started by providing an overview of the steps involved in setting up a Private Database Cloud. Before you set up the database cloud, you must download and deploy the required plug-ins. For more details, see [Section 3.3, "Deploying the Required Plug-ins"](#).

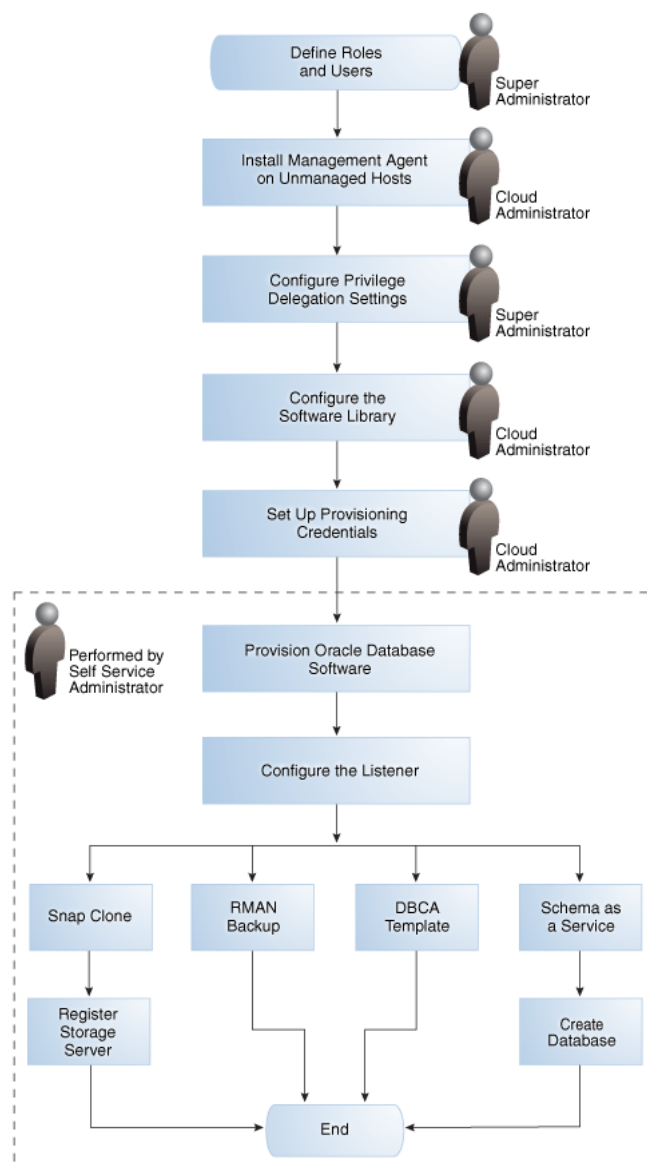
**Table 11–1** *Getting Started with DBaaS*

| Step | Task                                                                                                                                                                                                                                                                                                                                                        | Role                       |
|------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------|
| 1    | Define roles for administrators and self service users. See <a href="#">Section 3.4, "Defining Roles and Assigning Users"</a> .                                                                                                                                                                                                                             | Super Administrator        |
| 2    | Install the Management Agent on unmanaged hosts so that they can be monitored by Enterprise Manager. <a href="#">Section 9.2, "Adding Hosts"</a> .                                                                                                                                                                                                          | Cloud Administrator        |
| 3    | Configure Privilege Delegation Settings on your managed hosts. <a href="#">Section 3.6, "Configuring Privilege Delegation Settings"</a> .                                                                                                                                                                                                                   | Super Administrator        |
| 4    | Configure the storage location for the Software Library. See <a href="#">Section 3.1, "Setting Up the Software Library"</a>                                                                                                                                                                                                                                 | Cloud Administrator        |
| 6    | Set up provisioning credentials. See <a href="#">Section 11.2, "Setting Up Credentials for Provisioning"</a> .                                                                                                                                                                                                                                              | Self Service Administrator |
| 7    | If you are: <ul style="list-style-type: none"> <li>■ Setting up DBaaS, you must provision the database software or create Oracle Homes. See <a href="#">Section 11.3, "Provisioning Database Software"</a>.</li> <li>■ Setting up Schema as a Service, you must deploy the database. See <a href="#">Section 11.4, "Deploying the Database"</a>.</li> </ul> | Self Service Administrator |

**Table 11–1 (Cont.) Getting Started with DBaaS**

| Step | Task                                                                                                         | Role                       |
|------|--------------------------------------------------------------------------------------------------------------|----------------------------|
| 8    | Configure the Listener. See <a href="#">Section 11.5, "Configuring the Oracle Listener"</a> .                | Self Service Administrator |
| 9    | Register the storage servers. See <a href="#">Section 11.6, "Registering and Managing Storage Servers"</a> . | Self Service Administrator |

Figure 11–1 shows the sequence of steps involved in enabling Database or Schema as a Service.

**Figure 11–1 Enabling Database / Schema as a Service**

## 11.2 Setting Up Credentials for Provisioning

Before you perform any operations on the Managed Servers or databases, you must define the credentials that will be used by Enterprise Manager to connect to the targets.

You need to set up the following types of credentials:

- Normal credentials are the host operating system credentials used to provision the database software and create databases. For example, `oracle/<login password>`. These credentials are saved when the Database Pool is created and are used when the `EM_SSA_USER` requests a database or a schema.
- Privileged credentials are the host operating system credentials used to perform privileged actions like executing root scripts. These credentials are used deploying software (for running `root.sh` during deployment), for mounting and unmounting storage volumes (for databases created with snapshots) and so on. These credentials are saved along with the Database Pool if the pool is used for creating databases using snapshots.
- Database SYSDBA credentials are used and saved for schema as a service database pool. These credentials are required only for schema as a service.

---

**Note:** It is recommended that the same OS user who owns the Oracle Home on the host, creates the database.

---

To create named credentials, follow these steps:

1. Log in to Enterprise Manager as an administrator with the `EM_SSA_ADMINISTRATOR` role.
2. From the **Setup** menu, select **Security**, then select **Named Credentials**.
3. Click **Create** in the Named Credentials page.
4. Enter the Credential Name and Credential Description. Set the Authenticating Target Type field to Host and Scope field to Global. Enter the user name and password in the Credential Properties section. If you need to set privileged credentials, select Sudo or PowerBroker in the Run Privilege field and enter values in the Run As and Profile fields.
5. Click **Test** and **Save**.
6. Verify these credentials against a host target and click **OK**.

## 11.3 Provisioning Database Software

Before you can enable database as a service, the database software must already be provisioned on all hosts. Database software can be provisioned by an administrator with the `EM_SSA_ADMINISTRATOR` role in the following ways:

- **Provisioning Profile**
  - Capture a gold image of an existing database using a Provisioning Profile. See the Enterprise Manager Lifecycle Management Administrator's Guide for details.
  - Use the Provisioning Profile to provision the Clusterware/ASM or Grid Infrastructure (for Real Application Cluster databases), and Database Oracle Home. This method ensures that the necessary database plug-in (monitoring

part of the database plug-in) is deployed onto the Management Agent as part of the database provisioning Oracle Home installation.

To create a provisioning profile, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning** and select the database provisioning deployment procedure to be used. You can select either **Provision Oracle Database** or **Provisioning Oracle RAC Database** deployment procedure.

**Note:** Do not create a new database as part of this deployment procedure.

- **Using the Database Installer**

- From the **Setup** menu, select **Extensibility**, then select **Plug-ins**. and deploy the complete database plug-in on all the Management Agents in a PaaS Infrastructure Zone.
- Run the Clusterware/ASM or Grid Infrastructure installer to set up the cluster and ASM (for RAC databases).
- Run the Database Installer and ensure you select the **create database** option on all hosts.
- Discover the database. From the **Setup** menu, select **Add Target**, then **Add Targets Manually**, and then select **Add Non-Host Targets Using Guided Process** (Also Adds Related Targets).
- From the **Enterprise** menu, you can also select **Job**, then select **Library** and submit the **Discover Promote Oracle Home Target** job to add the Oracle Home.

For more details on provisioning the database software, see the Enterprise Manager Lifecycle Management Administrator's Guide.

## 11.4 Deploying the Database

For schema as a service, you must deploy a single instance or RAC database. To deploy a database, you must use the Provision Oracle Database deployment procedure. See the Enterprise Manager Lifecycle Management Administrator's Guide for details.

## 11.5 Configuring the Oracle Listener

You need to configure an Oracle Home and the Oracle Listener before you can add them as Enterprise Manager targets.

To set up the Oracle Listener (Listener) for the database hosts, follow these steps:

1. Log in as a user with the `EM_SSA_ADMINISTRATOR` role and perform mass deployment of database homes on the newly added hosts as described in [Section 9.2, "Adding Hosts"](#).
2. To configure a Listener running from the same Oracle Home on which the database instance is to be created, launch a Bash shell and enter the following commands:
  1. `<AGENT_BASE>/agent_inst/bin/emctl stop agent`
  2. `export TNS_ADMIN=<DB_HOME_LOCATION>/network/admin`
  3. `<AGENT_BASE>/agent_inst/bin/emctl start agent`
  4. `export ORACLE_HOME=<DB_HOME_LOCATION>`

5. Run `$ORACLE_HOME/bin/netca` and create the listener. Make sure you have the same Listener name and Listener port on all the hosts.
3. To configure a Listener running from the Single Instance High Availability (SIHA) Oracle Home, launch a Bash shell and enter the following commands:
  1. `export ORACLE_HOME=<SIHA_HOME_LOCATION>`
  2. Run `$ORACLE_HOME/bin/netca` and create the listener. Make sure you have the same listener name and listener port on all the hosts
4. Log in as the user with the `DBAAS_ADMIN_ROLE` and discover the newly added Listener target on all the hosts. From the **Setup** menu, select **Add Target**, then select **Add Target Manually**.
5. Select the **Add Non-Host Targets Using Guided Process** option and select Target Type as **Oracle Database, Listener, and Automatic Storage Management** and click **Add Guided Discovery** and follow the steps in the wizard. Before you add the new Listener target, ensure the `ORACLE_HOME` for the Listener is pointing to the correct `ORACLE_HOME` location. This process adds the Oracle Home target which is used when a database pool is created.

## 11.6 Registering and Managing Storage Servers

This chapter contains the following sections:

- [Before You Begin](#)
- [Overview of Registering Storage Servers](#)
- [Prerequisites for Registering Storage Servers](#)
- [Registering Storage Servers](#)
- [Administering the Storage Server](#)
- [Managing Storage Servers](#)

### 11.6.1 Before You Begin

Before you begin, note the following:

- Windows databases are not discovered as part of storage discovery. This is because the Windows storage NFS collection does not happen at all. NFS collection is also not supported on certain OS releases, and thus databases on those OS releases cannot be Snap Cloned. For further details please refer to the *My Oracle Support* note 465472.1. Also, NAS volumes cannot be used on Windows for supporting Oracle databases.
- Snap Clone is supported on Sun ZFS Storage 7120, 7320, 7410 and 7420 models.
- Snap Clone supports Sun ZFS storage on HP-UX hosts only if the OS version is B.11.31 or higher. If the OS version is lower than that, the Sun Storage may not function properly thereby Snap Clone gives unexpected results.
- By default, the maximum number of NFS file systems that Enterprise Manager discovers on a target host is 100. However, this threshold is configurable. You can also choose a list of file systems to be monitored if you do not want all the extra file systems to be monitored.

The configuration file `$agent_inst/sysman/emd/emagent_storage.config` for each host agent contains various storage monitoring related parameters.

To configure the threshold for the NFS file systems, you need to edit the following parameters:

```
Collection Size:START
Disks=1000
FileSystems=1000
Volumes=1000
Collection Size:END
```

If you choose to provide a list of file systems to be monitored, it can be provided between the following lines:

```
FileSystems:START
FileSystems:END
```

Restart the Management Agent and refresh the host configuration for the changes to this configuration file to be effective.

- Ensure that there is no extra space in any file path name or software library name.
- If the OMS Repository is running on RDBMS with 11.1.0.7.0 and AL32UTF8 character set, you need to apply patch 11893621.

## 11.6.2 Overview of Registering Storage Servers

Registering a storage server, such as NetApp storage server or Sun ZFS storage server, in Enterprise Manager enables you to provision databases using the snapshot and cloning features provided by the storage.

The registration process validates the storage, and discovers the Enterprise Manager managed database targets on this storage. Once the databases are discovered, you can enable them for *Snap Clone*. Snap Clone is the process of creating database clones using the Storage Snapshot technology.

## 11.6.3 Prerequisites for Registering Storage Servers

Before you register a storage server, follow the prerequisites outlined in the following sections:

- [Configuring Storage Servers](#)
- [Customizing Storage Agent Proxy](#)
- [Granting Roles and Privileges](#)

### 11.6.3.1 Configuring Storage Servers

Before you register a storage server, you require the following privileges and licenses to successfully use Snap Clone:

- [Obtaining NetApp Hardware Privileges and Licenses](#)
- [Obtaining Sun ZFS Hardware Privileges and Licenses](#)

---

---

**Note:** Only NetApp and Sun ZFS storage servers are supported in Enterprise Manager Cloud Control 12c.

---

---

### 11.6.3.1.1 Obtaining NetApp Hardware Privileges and Licenses NetApp Hardware Privileges

*Privileges* is a generic term. The term varies for different storage server vendors. For example, NetApp refers to privileges as *Capabilities*.

For NetApp storage server, to use Snap Clone, assign the following privileges or capabilities to the NetApp hardware user:

---

---

**Note:** You can assign these capabilities individually or by using wildcard notations. For example:

```
'api-volume-*', 'api-*', 'cli-*
```

---

---

- api-aggr-list-info
- api-aggr-options-list-info
- api-file-delete-file
- api-file-get-file-info
- api-file-read-file
- api-license-list-info
- api-nfs-exportfs-append-rules
- api-nfs-exportfs-delete-rules
- api-nfs-exportfs-list-rules
- api-nfs-exportfs-modify-rule
- api-snapshot-create
- api-snapshot-delete
- api-snapshot-list-info
- api-snapshot-reclaimable-info
- api-snapshot-restore-volume
- api-snapshot-set-reserve
- api-system-api-get-elements
- api-system-api-list
- api-snapshot-set-schedule
- api-system-cli
- api-system-get-info
- api-system-get-version
- api-useradmin-group-list
- api-useradmin-user-list
- api-volume-clone-create
- api-volume-clone-split-estimate
- api-volume-create
- api-volume-destroy

- api-volume-get-root-name
- api-volume-list-info
- api-volume-list-info-iter-end
- api-volume-list-info-iter-next
- api-volume-list-info-iter-start
- api-volume-offline
- api-volume-online
- api-volume-restrict
- api-volume-set-option
- api-volume-size
- cli-filestats
- login-http-admin

### Setup for NetApp Storage Server

---

---

**Note:** Snap Clone is supported only on NetApp Data ONTAP® 7.2.1.1P1D18 or higher, and ONTAP@ 8.x (7-mode).

---

---

To set up the NetApp storage server, follow these steps:

1. Create ROLE `em_smf_admin_role` with all the recommended capabilities, such as `api-aggr-list-info`, `api-file-delete-file`, and the like.
2. Create GROUP `em_smf_admin_group` with the ROLE `em_smf_admin_role`.
3. Create USER `em_smf_admin` with GROUP `em_smf_admin_group` and a secure password.

---

---

**Note:** The user `em_smf_admin` must be a dedicated user to be used by Oracle Enterprise Manager. Oracle does not recommend sharing this account for any other purposes.

---

---

### NetApp Hardware Licenses

Snap Clone on a NetApp storage server requires a valid license for the following services:

- flex\_clone
- nfs
- snaprestore

#### 11.6.3.1.2 Obtaining Sun ZFS Hardware Privileges and Licenses

*Privileges* is a generic term. The term varies for different storage server vendors. For example, Sun ZFS refers to privileges as *Permissions*.

For Sun ZFS storage server, to use Snap Clone, assign the following privileges or permissions to the Sun ZFS hardware user:



---

**Note:** All the permissions listed must be set to `true`. The scope must be `'nas'` and there must not be any further filters.

---

- `changeProtocolProps`
- `changeSpaceProps`
- `clone and createShare`
- `destroy`
- `rollback`
- `takeSnap`

### Setup for Sun ZFS Storage Server

To setup the Sun ZFS storage server, follow these steps:

1. Create ROLE `em_smf_admin_role`.
2. Create AUTHORIZATIONS for the ROLE `em_smf_admin_role`.
3. Set SCOPE as `nas`.
4. Set the recommended permissions, such as, `allow_changeProtocolProps`, `allow_changeSpaceProps`, and the like to `true`.
5. Create USER `em_smf_admin` and set its ROLE property as `em_smf_admin_role`.

---

**Note:** The user `em_smf_admin` must be a dedicated user to be used by Oracle Enterprise Manager. Oracle does not recommend sharing this account for any other purposes.

---

### Sun ZFS Hardware Licenses

Snap Clone on Sun ZFS Storage Appliance requires a license for the Clones feature. A restricted-use license for the same is included with the Enterprise Manager Snap Clone.

#### 11.6.3.2 Customizing Storage Agent Proxy

Before you register a NetApp storage server, meet the following prerequisites:

- [Acquiring Third Party Licenses](#)
- [Uploading Storage Vendor SDK](#)
- [Overriding the Default SDK](#)
- [Overriding Third Party Server Components](#)

---

**Note:** Storage Proxy Agent is supported only on *Linux Intel x64* platform.

---

**11.6.3.2.1 Acquiring Third Party Licenses** The Storage Management Framework is shipped by default for *Linux x86-64* bit platform, and is dependent on the following third party modules:

- Source CPAN - CPAN licensing apply

- IO::Tty (version 1.10)
- XML::Simple (version 2.20)
- Net::SSLeay (version 1.52)
- Open Source - Owner licensing apply
  - OpenSSL(version 1.0.1e)

**11.6.3.2.2 Uploading Storage Vendor SDK** Before you register a NetApp storage server, do the following:

1. Download the *NetApp Manageability SDK* version 5.1 for all the platforms from the following NetApp support site:  
<http://support.netapp.com/NOW/cgi-bin/software>
2. Unzip the 5.1 SDK and package the Perl NetApp Data OnTap Client SDK as a tar file. Generally, you will find the SDK in the `lib/perl/NetApp` folder. The tar file when extracted should look as follows:

```
NetApp.tar
- netapp
 - NaElement.pm
 - NaServer.pm
 - NaErrno.pm
```

For example, the Software Library entity `Storage Management Framework Third Party/Storage/NetApp/default` should have a single file entry that contains `NetApp.tar` with the above tar structure.

---

---

**Note:** Ensure that there is no extra space in any file path name or software library name.

---

---

3. Once the tar file is ready, create the following folder hierarchy in software library:  
`Storage Management Framework Third Party/Storage/NetApp`
4. Upload the tar file as a Generic Component named `default`.

---

---

**Note:** To upload the tar file, you must use the OMS Shared filesystem for software library.

---

---

The tar file should be uploaded to this default software library entity as a Main File.

**11.6.3.2.3 Overriding the Default SDK** The `default` SDK is used for all the NetApp storage servers. However, the storage server may work with only a certain SDK. In such a case, you can override the SDK per storage server, by uploading an SDK and using it only for this particular storage server.

To override the existing SDK for a storage server, upload the tar file to the Software Library entity. The tar file should have the structure as mentioned in Step 3 of the previous section.

The Software Library entity name should be the same as the storage server name.

For example, if the storage server name is `mynetapp.example.com`, then the Software Library entity must be as follows:

Storage Management Framework Third  
 Party/Storage/NetApp/mynetapp.example.com

---

**Note:** A storage specific SDK is given a higher preference than the default SDK,

---

**11.6.3.2.4 Overriding Third Party Server Components** By default, all the required third party components are shipped for *Linux Intel 64* bit platform. If you need to override it by any chance, package the tar file as follows:

---

**Note:** The tar file should contain a `thirdparty` folder whose structure should be as mentioned below:

---

```
thirdparty
|-- lib
| |-- engines
| | |-- lib4758cca.so
| | |-- libaep.so
| | |-- libatalla.so
| | |-- libcapi.so
| | |-- libchil.so
| | |-- libcsuif.so
| | |-- libgmp.so
| | |-- libgost.so
| | |-- libnuron.so
| | |-- libpadlock.so
| | |-- libsureware.so
| | `-- libubsec.so
| |-- libcrypto.a
| |-- libcrypto.so
| |-- libcrypto.so.1.0.0
| |-- libssl.a
| |-- libssl.so
| `-- libssl.so.1.0.0
`-- pm
 |-- CPAN
 | |-- IO
 | | |-- Pty.pm
 | | |-- Tty
 | | `-- Constant.pm
 | `-- Tty.pm
 |-- Net
 | |-- SSLeay
 | | `-- Handle.pm
 | `-- SSLeay.pm
 |-- XML
 | `-- Simple.pm
 `-- auto
 |-- IO
 | `-- Tty
 | |-- Tty.bs
 | `-- Tty.so
 `-- Net
 |-- SSLeay
 |-- SSLeay.bs
 `-- SSLeay.so
```

Ensure that the tar file is uploaded to the Software Library entity which is named after the platform name, `x86_64`. The Software Library entity must be under the following:

Storage Management Framework Third Party/Server

The `x86_64` entity, when uploaded is copied to all the storage proxy hosts irrespective of which storage server it would be processing. To use this entity on a specific storage proxy agent, name the entity after the host name.

For example, `Storage Management Framework/Third Party/Server/x86_64` will be copied to any storage proxy host which is on an `x86_64` platform. Similarly, `Storage Management Framework Third Party/Server/myhost.example.com` is copied only to `myhost.example.com`, if it is used as a storage proxy host.

The host name is given a higher preference than the platform preference.

### 11.6.3.3 Granting Roles and Privileges

To perform various storage server activities, you need to grant the following roles and privileges:

- [Accessing Security Class](#)
- [Granting General Privileges](#)
- [Granting Target Privileges](#)
- [Granting Roles](#)
- [Granting Privileges for Provisioning](#)

**11.6.3.3.1 Accessing Security Class** Table 11–2 displays the security class you need for accessing privileges for storage management.

**Table 11–2 Security Class**

| Security Class | Description                                      | Base Table           | Base Table Column                                            |
|----------------|--------------------------------------------------|----------------------|--------------------------------------------------------------|
| STORAGE        | Defines access privileges for storage management | SMF_STORAGE_HARDWARE | STORAGE_HARDWARE_GUID<br>STORAGE_HARDWARE_NAME<br>CREATED_BY |

**11.6.3.3.2 Granting General Privileges** Table 11–3 displays the general privileges you need to set before you register a storage server.

**Table 11–3 General Privileges for Storage Server Registration and Management**

| Privilege          | Description                                              | Scope  | Notion | Included Privilege             |
|--------------------|----------------------------------------------------------|--------|--------|--------------------------------|
| VIEW_ANY_STORAGE   | Ability to view any storage.                             | class  | VIEW   | nil                            |
| REGISTER_STORAGE   | Ability to register storage.                             | class  | CREATE | nil                            |
| VIEW_STORAGE       | Ability to view storage details.                         | object | VIEW   | nil                            |
| MANAGE_STORAGE     | Ability to synchronize storage.                          | object | EDIT   | VIEW_STORAGE<br>CREATE_JOB     |
| MANAGE_ANY_STORAGE | Ability to manage any of the registered storage servers. | object | EDIT   | VIEW_ANY_STORAGE<br>CREATE_JOB |
| FULL_STORAGE       | Ability to modify or remove storage.                     | object | FULL   | MANAGE_STORAGE                 |

**11.6.3.3.3 Granting Target Privileges** Table 11–4 displays the target privilege you need to set, for enabling or disabling Snap Clone for a target.

**Table 11–4 Target Privilege for Enabling or Disabling Snap Clone**

| Target Privilege  | Scope  | Notion | Include In Privilege | Included Privilege | Applicable Target Type          |
|-------------------|--------|--------|----------------------|--------------------|---------------------------------|
| SNAP_CLONE_TARGET | object | Manage | FULL_TARGET          | VIEW_TARGET        | oracle_database<br>rac_database |

**11.6.3.3.4 Granting Roles** Table 11–5 displays the roles you need to grant to be able to register a storage server and perform various activities on the registered storage server.

**Table 11–5 Roles for Registering and Managing the Storage Server**

| Role                     | Description                                                      | Security Class        | Privilege                  | Granted To Role      |
|--------------------------|------------------------------------------------------------------|-----------------------|----------------------------|----------------------|
| EM_STORAGE_ADMINISTRATOR | Role has privileges to register storage hardware for Snap Clone. | STORAGE               | MANAGE_ANY_STORAGE         | EM_ALL_ADMINISTRATOR |
|                          |                                                                  | STORAGE               | REGISTER_STORAGE           |                      |
|                          |                                                                  | TARGET                | VIEW_ANY_TARGET            |                      |
|                          |                                                                  | TARGET                | PERFORM_OPERATION_ANYWHERE |                      |
|                          |                                                                  | NAMED_CREDENTIALS     | CREATE_CREDENTIAL          |                      |
|                          |                                                                  | JOB                   | CREATE_JOB                 |                      |
|                          |                                                                  | SWLIB_ENTITY_MGMT     | SWLIB_CREATE_ANY_ENTITY    |                      |
|                          |                                                                  |                       | SWLIB_VIEW_ANY_ENTITY      |                      |
| EM_STORAGE_OPERATOR      | Role has privileges to manage storage hardware for Snap Clone.   | STORAGE               | MANAGE_ANY_STORAGE         | EM_ALL_OPERATOR      |
|                          |                                                                  | TARGET                | VIEW_ANY_TARGET            |                      |
|                          |                                                                  | TARGET                | PERFORM_OPERATION_ANYWHERE |                      |
|                          |                                                                  | JOB                   | CREATE_JOB                 |                      |
|                          |                                                                  | SWLIB_VIEW_ANY_ENTITY | SWLIB_VIEW_ANY_ENTITY      |                      |

**11.6.3.3.5 Granting Privileges for Provisioning** You need the following privileges to be able to use the storage server for provisioning:

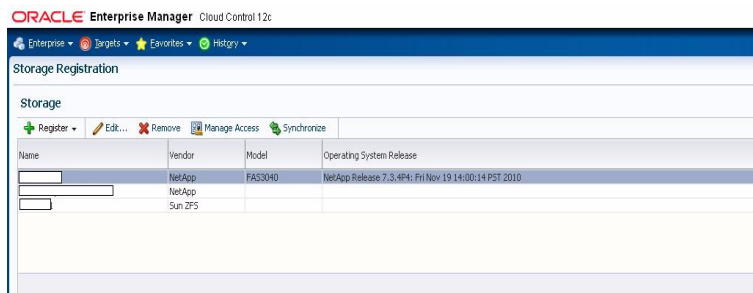
- VIEW\_STORAGE on the storage server
- GET\_CREDENTIAL on the storage server
- GET\_CREDENTIAL on all the Management Agent credentials of the storage server
- PERFORM\_OPERATION on the storage server Management Agent

## 11.6.4 Registering Storage Servers

To register the storage server, follow these steps:

1. From the **Setup** menu, click on **Provisioning and Patching**, and then select **Storage Registration**.

- On the Storage Registration page, in the Storage section, click on **Register**, and then select either **NetApp Storage** or **Sun ZFS Storage**, based on which storage server you want to register.



**Note:** Enterprise Manager Cloud Control 12c supports only NetApp and Sun ZFS storage servers.

You need the EM\_STORAGE\_ADMINISTRATOR role to complete the storage registration.

- On the NetApp or Sun ZFS Storage Registration page, in the Storage section, do the following:



- Enter the storage server name in the Name field. Ensure that the name is the valid host name and contains no spaces and invalid characters.
- Select the protocol.

**Note:** For NetApp storage, the connection is over http or https. For Sun ZFS storage, the connection is over ssh.

- Select the Storage Credentials, or click on the green plus sign to add.

**Note:** These credentials will be used by the Management Agent to execute storage (NetApp or Sun ZFS) APIs.

Only credentials owned by the user are displayed here.

In the display box that appears, enter the storage server name and password. Confirm the password and click **OK**.

- Enter storage name aliases (optional).

---

---

**Note:** A storage name alias is any name that may have been used when mounting a volume from the storage.

For example: IP address, FQDN, DNS alias, and the like.

Storage alias is necessary to identify the database targets on the storage. The database targets are identified by mapping the mount points to the files used by the database. For example, if the storage mystorage.com has an alias mystorage.net, and a database uses a data file mounted as mystorage.net:/u01, then mystorage.net must be added as an alias for the discovery to work.

---

---

4. In the Agent to Manage Storage section, do the following:
- Click **Add** to add a Management Agent host. A Storage Agent display box appears. Select a Management Agent from the Target Name column of the table. Then, click **Select**.

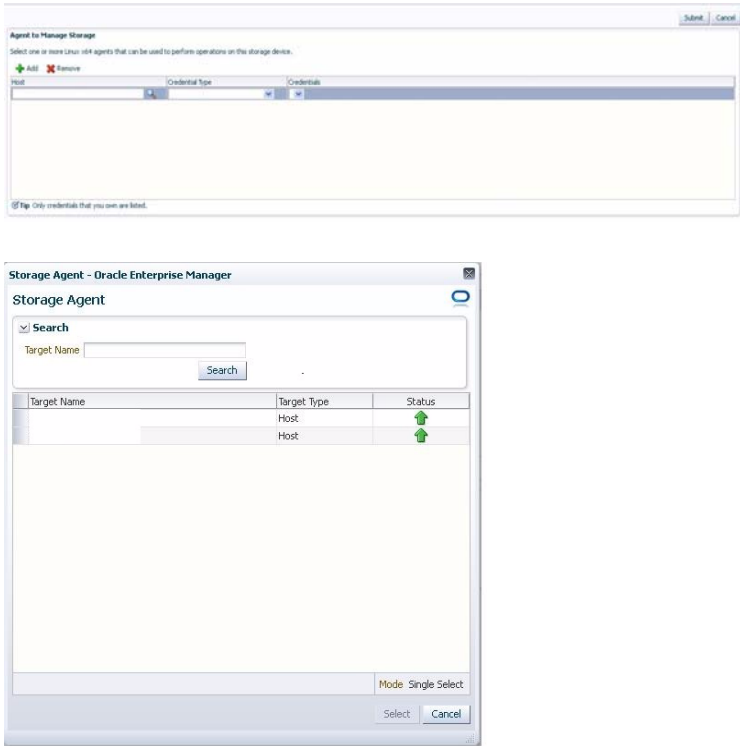
---

---

**Note:** The Management Agent list displays only *Linux X64* Management Agents.

---

---



The Management Agent selected is used for performing operations on the storage server.

- Once a Management Agent is selected, the Management Agent credentials are found and a named credential for the host is displayed.

---

---

**Note:** The Management Agent credentials are used to connect to the Management Agent from Oracle Management Service.

---

---


Multiple Management Agents can be configured to monitor the storage device. Click **Add** to choose a second Management Agent if required.

---

**Note:** Configuring multiple Management Agents to monitor the storage device provides you with a backup in the event that an host is down or the Management Agent is under blackout.

---

- Click **Submit** to register the storage server.



## 11.6.5 Administering the Storage Server

To administer the storage server, refer to the following sections:

- [Synchronizing Storage Servers](#)
- [Enabling or Disabling Snap Clone](#)
- [Deregistering Storage Servers](#)

### 11.6.5.1 Synchronizing Storage Servers

When you register a storage server for the first time, a synchronize job is run automatically. However, to discover new changes or creations, you should schedule a synchronize job to run at a scheduled time, preferably during a quiet period when Snap Clone actions are not in progress. To do this, follow these steps:

1. On the Storage Registration page, in the Storage section, click **Synchronize**.





---

**Note:** When you click on Synchronize, a deployment procedure is submitted which discovers all databases monitored by Enterprise Manager Cloud Control which can be used for Snap Clone.

You need EM\_STORAGE\_OPERATOR role along with GET\_CREDENTIAL privilege on the Storage Server and Storage Management Agent credentials to be able to synchronize the storage.

---

2. A confirmation box appears. Click **OK**.





This action now submits a one-time synchronization job.

---

---

**Note:** The synchronization job fetches latest storage information, and recomputes the mapping between storage volumes and databases.

---

---

3. On the Storage Registration page, in the Storage section, to view the procedure details of the Management Agent host, click on the value (for example, Scheduled) in the Status column.
4. On the Provisioning page, in the Procedure Steps section, click **View**, and then select **Expand All**. Keep clicking the Refresh button on the page to view the procedure activity as it progresses.

The synchronization status of the Management Agent on the Storage Registration page, changes to Succeeded once the synchronization process is complete.

5. To update a synchronize schedule of a registered storage server, select a storage server on the Storage Registration page and then click on **Edit**. On the Edit Storage page, in the Synchronize Storage section, edit the repetition time and frequency of the synchronize job.

---

---

**Note:** The frequency of a synchronization job is set at 3 hours by default.

---

---

Click **Submit**.

---

---

**Note:** The *Associating Storage Volumes With Targets* step relies on both database target metrics and host metrics. The database target (oracle\_database/rac\_database) should have up-to-date metrics for the *Controlfiles*, *Datafiles* and *Redologs*. The File Systems metric should be up to date for the hosts on which the database is running.

---

---

### 11.6.5.2 Enabling or Disabling Snap Clone

To enable or disable Snap Clone, do the following:

---

---

**Note:** You need SNAP\_CLONE\_TARGET or FULL\_TARGET privilege on the database to enable or disable Snap Clone.

---

---

1. On the Storage Registration page, in the Details section, select the Databases tab, and then select the database from the table which you want to use for Snap Clone, and then click **Enable Snap Clone** or **Disable Snap Clone**.

The Supported column in the table shows if the databases listed support Snap Clone or not.

---

---

**Note:** If a database listed shows the Not Supported icon in the Supported column, you can enable the Message column from the View menu. This column displays the reason. A possible reason is that the data files are missing on the storage server and may be located on either local disk or some other unregistered storage server. Another reason could be that the data files reside on another storage that is not synchronized.

---

---

2. Once Snap Clone has been enabled or disabled successfully for the database, a confirmation display box appears. Click **OK**.

---

---

**Note:** To view all the files located on the storage server, click **Show Files**.

To view all the projects (aggregates), volumes, and snapshots in the storage server, select the **Contents** tab.

---

---

### 11.6.5.3 Deregistering Storage Servers

To deregister a registered storage server, follow these steps:

---

---

**Note:** To deregister a storage server, you need `FULL_STORAGE` privilege on the storage along with `FULL_JOB` privilege on the Synchronization GUID of the storage server.

---

---

1. From the **Setup** menu, click on **Provisioning and Patching**, and then select **Storage Registration**.
2. On the Storage Registration page, in the Storage section, select a storage server from the list of registered storage servers.
3. Select **Remove**.

On the Remove Storage page, select the storage server that you want to deregister, and then click **Submit**.



The storage server is now deregistered.

---

---

**Note:** Once a storage is deregistered, the Snap Clone profiles and Service Templates on the storage will no longer be functional, and the relationship between these Profiles, Service Templates and Snap Cloned targets will be lost.

---

---

## 11.6.6 Managing Storage Servers

To manage the storage server, refer to the following sections:

- [Managing Storage Allocation](#)
- [Managing Storage Access Privileges](#)

- [Viewing Storage Registration Overview and Hierarchy](#)

### 11.6.6.1 Managing Storage Allocation

*Storage Ceiling* is the maximum amount of storage from a project or aggregate that Enterprise Manager is allowed to use. This ensures that Enterprise Manager creates clones in that project only till this limit is reached. When a storage project is discovered for the first time, the entire capacity of the project is set as the ceiling. In case of Sun ZFS, the quota set on the project is used.

---

**Note:** You must explicitly set quota property for the Sun ZFS storage project on the storage end. Also, the project should have a non zero quota set on the storage end. Else, Enterprise Manager will not be able to clone on it.

---

To edit the storage ceiling, do the following:

1. On the Storage Registration page, select the Contents tab and then click **Edit Storage Ceiling**.

---

**Note:** *Edit Storage Ceiling* option enables you modify the maximum amount of storage that Enterprise Manager can use. You can create clones or resize volumes only till this limit is reached.

---

2. In the Edit Storage Ceiling dialog box, enter the storage ceiling, and then, click **OK**.

### 11.6.6.2 Managing Storage Access Privileges

To manage storage access privileges for a registered storage server, follow these steps:

1. On the Storage Registration page, in the Storage section, select a storage server from the list of registered storage servers.

---

**Note:** the Storage Registration page displays only the databases which you have `VIEW_STORAGE` privilege on.

---

2. Click **Manage Access**.
3. On the Manage Access page, do the following:
  - Click **Change**, if you need to change the Owner of the registered storage server.

---

**Note:** The Owner of a registered storage server can perform all actions on the storage server, and grant privileges and roles to other Administrators.

---

- Click **Add Grant** to grant privileges to an Administrator, Role or both.
- On the Add Grant page, enter an Administrator name or select the type, and then click **Go**.

- Select an Administrator from the list of Administrators or Roles, and then click **Select**.
  - 4. On the Manage Access page, you can change privileges of an Administrator or Role by selecting the Administrator or Role from the Grantee column, and then clicking **Change Privilege**.
  - 5. In the Change Privilege display box, you can select one of the three following privileges:
    - View Storage (ability to view the storage)
    - Manage Storage (ability to edit the storage)
    - Full Storage (ability to edit or remove the storage)
- Click **OK**.
- 6. You can also revoke a grant to an Administrator by selecting the Administrator from the Grantee column, and then clicking **Revoke Grant**.
  - 7. When you are done with granting, revoking, or changing privileges to Administrators or Roles, click **Submit**.

---

**Note:** To be able to use the storage server, you also need to specifically grant privileges to the storage server and storage Management Agent credentials to the user.

---

### 11.6.6.3 Viewing Storage Registration Overview and Hierarchy

To view the storage registration overview, on the Storage Registration page, in the Details section, select the **Overview** tab. The Overview section provides a summary of storage usage information. It also displays a Snap Clone Storage Savings graph that shows the total space savings by creating the databases as a Snap Clone versus without Snap Clone.

---

**Note:** If you have NetApp volumes with no space guarantee, you may see negative allocated space in the Overview tab. Set guarantee to 'volume' to prevent this.


---



To view the storage registration hierarchy, on the Storage Registration page, in the Details sections, select the **Hierarchy** tab. This displays the storage relationships between the following:

- Test Master Database
- Database Profile
- Snap Clone Database
- Snap Clone Database Snapshots

You can select a row to display the corresponding Volume or Snapshot Details.



If a database profile or Snap Clone database creation was not successful, and it is not possible to delete the entity from its respective user interface, click on the **Remove** button to access the Manage Storage page. From this page, you can submit a procedure to dismount volumes and delete the snapshots or volumes created from an incomplete database profile or snap clone database.

---

---

**Note:** The Manage Storage page only handles cleanup of storage entities and does not remove any database profile or target information from the repository.

---

---

You can also select the **Procedure Activity** tab on the right panel, to see any storage related procedures run against that storage entity.



---

## Setting Up the DBaaS Self Service Portal

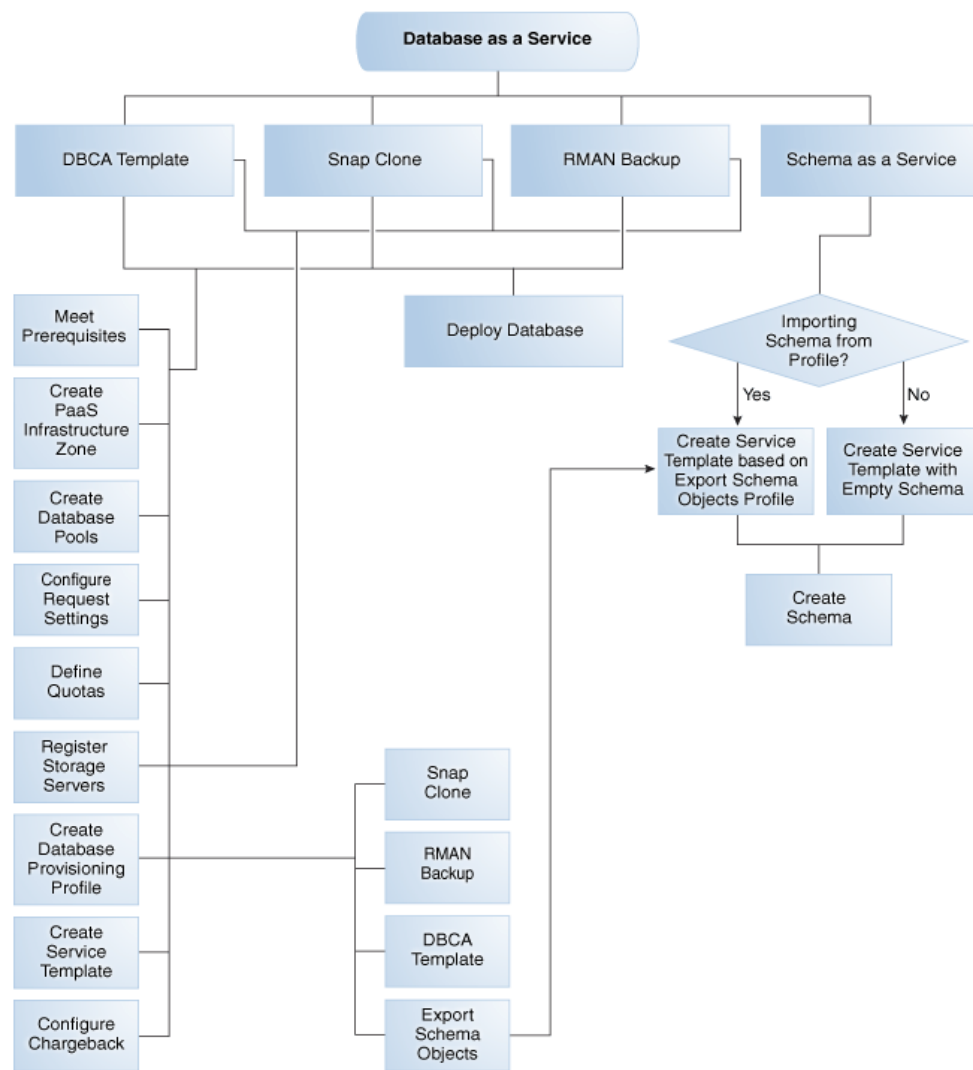
This chapter describes the procedure to set up and use the Database Self Service Portal. It contains the following sections:

- [Setting Up the Database Cloud Self Service Portal](#)
- [DBaaS Using Snap Clone Based Database Provisioning Profile](#)
- [DBaaS Using RMAN Backup Based Database Provisioning Profile](#)
- [DBaaS Using DBCA Template Based Database Provisioning Profile](#)
- [Creating a Database Pool](#)
- [Configuring Request Settings](#)
- [Setting Up Quotas](#)
- [Setting Up Profiles and Service Templates](#)
- [Creating a Database Provisioning Profile](#)
- [Creating a Database Service Template](#)
- [Configuring Chargeback](#)
- [Using Schema as a Service to Create Schemas](#)
- [Pre and Post Request Creation / Deletion Scripts](#)

### 12.1 Setting Up the Database Cloud Self Service Portal

The Database Cloud Self Service Portal is a Web page where Cloud users can request, monitor, and manage their own databases using Enterprise Manager. An administrator with the EM\_SSA\_ADMINISTRATOR role can configure and set up the database self service portal by defining database pools, service templates, and so on.

[Figure 12-1](#) lists the steps involved in setting up database or schema as a service.

**Figure 12–1 Setting up the Database Self Service Portal**

## 12.2 DBaaS Using Snap Clone Based Database Provisioning Profile

You can create a snap clone of a large database which will allow to quickly provision the database. The snap clone is stored in the Software Library and can be used to create service templates and provision databases. The snap clone feature is useful if your data size is very large and the database files are present on supported storage systems like ZFS Storage Appliance and NetApp Storage System. See [Section 11.6, "Registering and Managing Storage Servers"](#) for details.

To provision a database using snapshots, follow these steps:

1. Follow the steps in the Getting Started section to enable DBaaS. See [Section 11.1, "Getting Started"](#).
2. Set up one or more PaaS Infrastructure Zones. See [Section 9.3, "Creating a PaaS Infrastructure Zone"](#).
3. Create a Database Pool. See [Section 12.5, "Creating a Database Pool"](#).
4. Configure the request settings. See [Section 12.6, "Configuring Request Settings"](#).



5. Define quotas for each self service user. See [Section 12.7, "Setting Up Quotas"](#).
6. Register the storage server. See [Section 11.6, "Registering and Managing Storage Servers"](#).
7. Create a database provisioning profile using a snap clone. Snap clones are available only if the storage system such as ZFS Storage Appliance or Netapp Storage System has been registered and the database is running on it. See [Section 12.9.1, "Creating a Database Provisioning Profile Using Snapshots"](#).
8. Create a service template that is based on the profile you have created. See [Section 12.10, "Creating a Database Service Template"](#).
9. Configure the Chargeback Service. See [Section 22, "Chargeback Administration"](#).
10. While deploying a database, select the service template (associated with the snapshot based profile) you have created. See [Section 13.2, "Requesting Databases and Schemas"](#).

## 12.3 DBaaS Using RMAN Backup Based Database Provisioning Profile

Database backups created by RMAN (Recovery Manager) are stored as image copies or backup sets. To provision a database using RMAN backup, follow these steps:

1. Follow the steps in the Getting Started section to enable DBaaS. See [Section 11.1, "Getting Started"](#).
2. Set up one or more PaaS Infrastructure Zones. See [Section 9.3, "Creating a PaaS Infrastructure Zone"](#).
3. Create a Database Pool. See [Section 12.5, "Creating a Database Pool"](#).
4. Configure the request settings. See [Section 12.6, "Configuring Request Settings"](#).
5. Define quotas for each self service user. See [Section 12.7, "Setting Up Quotas"](#).
6. Register the storage server. See [Section 11.6, "Registering and Managing Storage Servers"](#).
7. Create a database provisioning profile using RMAN Backup. See [Section 12.9.2, "Creating a Database Provisioning Profile Using RMAN Backup"](#).
8. Create a service template that is based on the profile you have created. See [Section 12.10, "Creating a Database Service Template"](#).
9. Configure the Chargeback Service. See [Section 22, "Chargeback Administration"](#).
10. While deploying a database, select the service template (associated with the RMAN Backup based profile) you have created. See [Section 13.2, "Requesting Databases and Schemas"](#).

## 12.4 DBaaS Using DBCA Template Based Database Provisioning Profile

You can create database (DBCA) templates that can be used to provision databases. To provision a database using DBCA templates, follow these steps:

1. Follow the steps in the Getting Started section to enable DBaaS. See [Section 11.1, "Getting Started"](#).
2. Set up one or more PaaS Infrastructure Zones. See [Section 9.3, "Creating a PaaS Infrastructure Zone"](#).
3. Create a Database Pool. See [Section 12.5, "Creating a Database Pool"](#).

4. Configure the request settings. See [Section 12.6, "Configuring Request Settings"](#).
5. Define quotas for each self service user. See [Section 12.7, "Setting Up Quotas"](#).
6. Register the storage server. See [Section 11.6, "Registering and Managing Storage Servers"](#).
7. Create a database provisioning profile using a DBCA Template. See [Section 12.9.4, "Creating a Database Provisioning Profile Using DBCA Template"](#).
8. Create a service template that is based on the profile you have created. See [Section 12.10, "Creating a Database Service Template"](#).
9. Configure the Chargeback Service. See [Section 22, "Chargeback Administration"](#).
10. While deploying a database, select the service template (associated with the DBCA template based profile) you have created. See [Section 13.2, "Requesting Databases and Schemas"](#).

## 12.5 Creating a Database Pool

A Database Pool contains a set of resources that can be used to provision a database instance within a PaaS Infrastructure Zone. A database pool is a collection of homogenous targets such as servers or clusters with database software installed.

To create a database pool for database as a service, follow these steps:

1. Log into Enterprise Manager as a user with the EM\_SSA\_ADMINISTRATOR. role.
2. From the **Setup** menu, select **Cloud**, then select **Database**. The Database Cloud Self Service Portal Setup page appears.
3. From the **Create** menu, select **For Database**. The Create New Pool page appears.

**Figure 12–2 Create Database Pool**

**Database Pools**

**Create New Pool : Setup** [Back] Step 1 of 1 [Submit] [Cancel]

**Pool Details**

\* Name: new\_pool  
Description: [Text Area]

**Credentials**

Specify the host credentials that will be used for performing database creation operations. Root credentials are optional. They are needed if you plan to use this pool for provisioning database requests.

\* Host: SYSMAN\_ORACLE (SYSMAN) [Add]  
\* Root: SYSMAN\_ROOT (SYSMAN) [Add]

**Oracle Homes**

Add one or more Oracle Homes to the pool. All Oracle Homes must reside in the same PaaS infrastructure zone. Homogeneity is controlled by the target filters that cannot be modified once the pool is created.

\* PaaS Infrastructure Zone: [Dropdown] \* RAC Zone: [Dropdown] \* Platform: Linux x86-64 [Dropdown]  
\* Database Configuration: Database Instance [Dropdown] \* Version: 11.2.0.3.0 [Dropdown]

[Add] [Remove] [Test All Credentials]

| Name                                  | Location | On Host |
|---------------------------------------|----------|---------|
| Select targets for the pool creation. |          |         |

☒ **TIP:** A target can only belong to one pool.

Where is my Oracle Home?

**Placement Constraints**

Placement policy constraints allow the self service administrator to set maximum ceilings for resource utilization. This provides protection for the members of the software pool in terms of resource consumption. For example, a production software pool might enforce more conservative limits, whereas a development software pool might enforce more liberal limits.

\* Maximum Number of Database Instances per Host: 10 [Dropdown]

4. In the Credentials region, specify the following credentials:
  - **Host:** Required for creating and starting the database.
  - **Root:** Required for databases created using storage technologies. See [Section 12.2, "DBaaS Using Snap Clone Based Database Provisioning Profile"](#).

5. Enter a name and description for the database pool.
6. Specify the following details:
  - **PaaS Infrastructure Zone:** Select the PaaS Infrastructure Zone on which the database instance is to be provisioned.
  - **Database Configuration:** Specify the configuration of the database being provisioned. This can either be Single Instance or Cluster Database (RAC) but cannot be a combination of both.
  - **Platform and Version:** Specify the platform and version of the database being deployed.
7. Click **Add** and select one or more Oracle Homes to be added to the database pool.
8. In the Maximum Number of Database Instances (per host), enter the maximum number of database instances that can be running on each host in the pool.

---

**Note:** For RAC databases, the value specified in this field is used to ensure that the nodes on which the database is to be created meets this criteria.

---

9. Click **Submit** to create a database pool. The newly created pool will appear in the Database Cloud Self Service Portal Setup: Database Pools page. The database pool can now be added to the Service Template.

## 12.6 Configuring Request Settings

You can configure the database request settings by specifying when a request can be made, its duration, and so on. To configure the request settings, follow these steps:

1. Log into Enterprise Manager as a user with the `EM_SSA_ADMINISTRATOR` role.
2. From the **Setup** menu, select **Cloud**, then select **Database**. The Database Cloud Self Service Portal Setup page appears.
3. Click the **Request Settings** tab. Specify the following:
  - **Future Reservation Length:** How far in advance a request can be scheduled. You can select:
    - **No Restriction:** A new database request can be scheduled on any date later than the current date.
    - **Restricted Reservation:** A new database request can be scheduled in the future within the period specified here. Any future request scheduled beyond the Restricted Reservation date will be denied.

For example, if the Restricted Reservation is 1 year, and the current date is June 20, 2013, you can schedule a future request on or before June 20, 2014.
  - **Request Archive Retention:** Period after which the requests are automatically archived. You can select:
    - **No Restriction:** If this option is selected, requests are not archived.
    - **Restricted Retention:** Requests that older than the specified duration will be archived. The date on which a request is archived is based on the last modified date of the request.

For example, if a request has been created on June 10, 2013 and the **Restricted Retention** period specified is 10 days. Suppose there are some associated child requests as follows:

- Reschedule Request: Submitted on June 12
- Start Instance: Submitted on June 14
- Stop Instance: Submitted on June 17
- Start Instance: Submitted on June 18
- Delete Instance: Submitted on June 20

Since the date on which the request was last modified is June 20, the request will be purged on June 30th (Restricted Retention = 10 days).

- **Default Retirement Period:** The maximum period for which the service instance can be retained. You can select:
  - **No Restriction:** Requests do not have a end date and can be extended as required.
  - **Restricted Duration:** The period after which the service instance is to be released.

For example, if the Restricted Duration is 15 days, and the request begin date is June 5, 2013, the service instance will be retired or released on June 25, 2013.

4. Click **Apply**.

## 12.7 Setting Up Quotas

Quota is the aggregate amount of resources that can be granted to each self service user belonging to a certain role. This quota applies only to the databases provisioned through the Database Cloud Self Service Portal. To set up quotas, follow these steps:

1. Log into Enterprise Manager as a user with the EM\_SSA\_ADMINISTRATOR role.
2. From the **Setup** menu, select **Cloud**, then select **Database**. The Database Cloud Self Service Portal Setup page appears.
3. Click the **Quotas** tab and then click **Create**.

---

---

**Note:** Quota validation ensures that a request is executed only if the user's current usage is within the limits specified for the roles to which the user belongs. Only the direct roles are considered for the quota determination, and not nested ones.

---

---

4. In the Create New Quota window, select the Role Name for which the quota is to be granted. Only roles that have the EM\_SSA\_USER privileges are displayed in the Select Roles dialog box. After you have selected the role, enter the following details:
  - **Memory:** Total allocated memory that can be assigned for all the databases owned by a specific user.
  - **Storage:** Total allocated storage that can be assigned for all the databases owned by a user.

- Number of Database Requests: The number of databases that can be requested by a user at any point in time.
- Number of Schema Service Requests: The number of database services that can be requested by the user.

For example, if a user has 3 roles and the quota defined for each role is as follows:

| Role Name | Attribute       | Value         |
|-----------|-----------------|---------------|
| Dev Role  | Memory          | 2 GB          |
|           | Storage         | 15 GB         |
|           | DB Requests     | 10            |
|           | Schema Requests | 2             |
| Test Role | Memory          | 5 GB          |
|           | Storage         | 10 GB         |
|           | DB Requests     | 12            |
|           | Schema Requests | 3             |
| PM Role   | Memory          | Quota not set |
|           | Storage         | Quota not set |
|           | DB Requests     | Quota not set |
|           | Schema Requests | Quota not set |

The quota limits for each of the attributes would be determined based on all the roles to which a user directly belongs. For example, if a user has 3 roles with quota setup for each role as shown in the table below:

The quota is calculated as the absolute maximum across the roles for each attribute value. For the user who belongs to the Dev, Test, and PM Role, the Quota is calculated as:

Value for Memory = MAX (Dev, Test, PM) = MAX (2, 5, NA) = 5 GB

Value for Storage = MAX (Dev, Test, PM) = MAX (15, 10, NA) = 15 GB

Value for DB Requests = MAX (Dev, Test, PM) = MAX (10, 12, NA) = 12 Databases

Value for DB Requests = MAX (Dev, Test, PM) = MAX (2, 3, NA) = 5 Database Services

A sample quota validation log generated by the user request follows:

```
***** UserName - SSA_USER
***** Quota Limits:
*** Memory: 10
*** Storage: 10
*** Instances: 10
*** Services: 3
***** Current Usage:
*** Memory Used: 6.732421875
*** Storage Used: 5.28
*** Instances Used: 5
*** Services Used: 1
```

## 12.8 Setting Up Profiles and Service Templates

A database provisioning profile is an entity that captures source database information for provisioning. A profile can represent a complete database or a set of related schemas that form an application. A single profile can be used to create multiple service templates. See [Section 12.9, "Creating a Database Provisioning Profile"](#).

A service template is a standardized service definition that is offered to self service users to create databases or schemas. Self service users can create one or databases or schemas based on the service template definition. You can create a service template for:

- Database as a Service. See [Section 12.10, "Creating a Database Service Template"](#)
- Schema as a Service. See [Section 12.12.2, "Creating a Service Template for Schema as a Service"](#)

The following figure shows a list of database profiles and service templates that have been created.

**Figure 12–3 Profiles and Service Templates**

Database Cloud Self Service Portal Setup Page Refreshed Jan 4, 2013 9:51:02 PM PST

---

Database Pools

Request Settings

Quotas

**Profiles and Service Templates**

**Profiles**  
Profiles are entities that capture source database information for provisioning. A profile can represent a complete database or a set of related schemas that form an application.

View ▼ Create... Delete

| Name                                           | Status     | Description                                         |
|------------------------------------------------|------------|-----------------------------------------------------|
| Snapshot profile t112 database - SIDB          | Ready      | Database Reference Profile 16-12-2012 11:07 PM from |
| DBCA Template Profile for t112 database - SIDB | Ready      | Database Reference Profile 17-12-2012 12:11 AM from |
| HR Schema Profile from                         | Ready      | Database Reference Profile 17-12-2012 04:03 AM from |
|                                                | Incomplete | Database Reference Profile 18-12-2012 01:09 AM from |
|                                                | Incomplete | Database Reference Profile 18-12-2012 01:53 AM from |
|                                                | Incomplete | Database Reference Profile 18-12-2012 01:51 AM from |
|                                                | Ready      | Database Reference Profile 18-12-2012 02:13 AM from |
|                                                | Incomplete | Profile created for backup for                      |
|                                                | Ready      | Database Reference Profile 17-12-2012 02:37 AM from |
|                                                | Incomplete | Database Reference Profile 17-12-2012 06:31 AM from |

Snapshot Profile bugdemo Database 03

**Service Templates**  
A service template is a standardized service definition that is offered to self service users to create databases or schemas.

View ▼ Create Edit Delete

| Service Template Name                           | Created by  | Zones | Roles | Description                                         |
|-------------------------------------------------|-------------|-------|-------|-----------------------------------------------------|
| Create Empty schema                             | DBAAS_ADMIN | 1     | 2     | Create Empty schema for test applications           |
| Create HR schema                                | DBAAS_ADMIN | 1     | 2     | Create HR schema                                    |
| Create RAC Database Using RMAN Technology       | DBAAS_ADMIN | 1     | 1     | Create RAC Database Using RMAN Technology RDB       |
| Create RAC Database using Snapshot Technology 2 | DBAAS_ADMIN | 1     | 2     | Service template which uses a snapshot based profil |
| Create SIDB Using DBCA Profile Technology       | DBAAS_ADMIN | 1     | 1     | Create SIDB Using DBCA Profile Technology RDBMS     |
| Create SIDB Using Snapshot NetApp Technology    | DBAAS_ADMIN | 1     | 1     | Create SIDB Using Snapshot NetApp Technology RDI    |
| Create SIDB using RMAN Technology               | DBAAS_ADMIN | 1     | 1     | Create SIDB using RMAN Technology. RDBMS versio     |
| sample snapshot rac template                    | DBAAS_ADMIN | 1     | 2     |                                                     |

## 12.9 Creating a Database Provisioning Profile

A Database Provisioning Profile is an entity which contains software bits and configuration.

When a provisioning profile is created from an existing installation, it provides the flexibility to clone Grid Infrastructure (with software or configuration) and Oracle Database (with software or configuration).

An administrator can create a database provisioning profile as a one-time activity, which can be used by operators for mass deployments.

Using provisioning profile enables standardization in deployments and reduces need for rescheduling deployments by avoiding errors while configuring deployment procedures.

You can create a database provisioning profile, by using any of the following methods:

- [Creating a Database Provisioning Profile Using Snapshots](#)

- [Creating a Database Provisioning Profile Using RMAN Backup](#)
- [Creating a Database Provisioning Profile Using Existing RMAN Backup](#)
- [Creating a Database Provisioning Profile Using DBCA Template](#)
- [Creating a Database Provisioning Profile Using Export Schema Objects](#)

---

**Note:** Before you create a provisioning profile, follow these prerequisites:

- Ensure that the host has a reference database installed on it, for creating a database provisioning profile.
  - Ensure that the host is connected to a Management Agent installed and monitored in Enterprise Manager Cloud Control for communication.
- 

## 12.9.1 Creating a Database Provisioning Profile Using Snapshots

### Prerequisites for Creating a Database Provisioning Profile Using Snapshots

Before you create a database provisioning profile, follow these prerequisites:

- Ensure that the storage server you want to register for storage is available on the network.

---

**Note:** Only NetApp and SunZFS storage servers are supported in Enterprise Manager Cloud Control 12c.

---

- Ensure that the storage server is connected to a Management Agent installed and monitored in Enterprise Manager Cloud Control for communication.
- Ensure that the storage server is registered, and at least one database should be present which is enabled for Snap Clone.

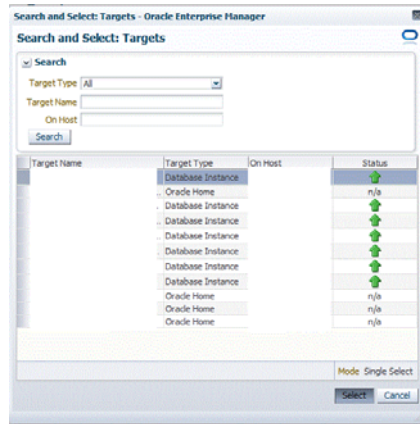
### Creating a Database Provisioning Profile Using Snapshots

To create a database provisioning profile, follow these steps:

1. You can access the Database Provisioning page using either of the following ways:
  - From the Setup menu, select **Cloud**, and then select **Database**. On the Database Cloud Self Service Portal Setup, select **Profiles and Service Templates**. In the Profiles section, click **Create**.
  - From the Enterprise menu, select **Provisioning and Patching**, and then select **Database Provisioning**. On the Database Provisioning page, in the Profiles section, click **Create**.
2. On the Reference Target page, click the search icon to select a target.



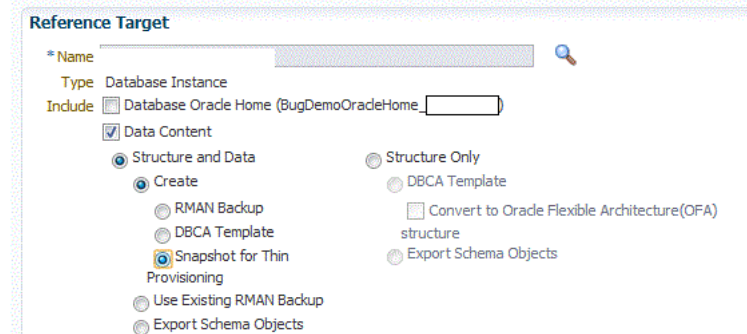
3. In the Search and Select:Targets display box, select the database you want to create a profile for, and then, click **Select**.




**Note:** Ensure that the database you select is enabled for Snap Clone. For information refer to *Enabling or Disabling Snap Clone*.

4. On the Reference Target page, select the following:
  - **Data Content**
  - **Structure and Data** to include physical and structural files from the database
  - **Create**
  - **Snapshots**

#### Create Database Provisioning Profile : Reference Target



5. In the Credentials section, select the database. In the Credentials column, you can select **Named Credentials** and then select a credential name from the Credential Name column (or) select **Preferred Credentials**. Click **Next**.



6. On the Profile page, do the following:
  - In the Profile Information section, enter a unique profile name of your choice.  
For example:  
Snap Clone Profile for HR database



You can also enter a description of the profile, and notes for your reference.

- In the Schedule section, you can choose to start the profile creation immediately, or you can schedule it for a later time.
- Click **Next**.

7. On the Review page, ensure that the selections you have made in the previous pages are correctly displayed and click **Submit**. Otherwise, click **Back** repeatedly till you reach the page where you want to make changes. Click **Cancel** to abort the provisioning profile creation.

| Volume                   | Mount Point          | Type        | Size (G) |
|--------------------------|----------------------|-------------|----------|
| central_emu_bugdemo_redo | /mnt/bugdb_demo_redo | Full Volume | 200.00   |
| central_emu_bugdemo_data | /mnt/bugdb_demo_data | Full Volume | 1100.0   |


Once you have submitted the provisioning profile creation job, the provisioning profile appears in the Database Provisioning page. The Database Provisioning page displays all the procedure steps.

## 12.9.2 Creating a Database Provisioning Profile Using RMAN Backup

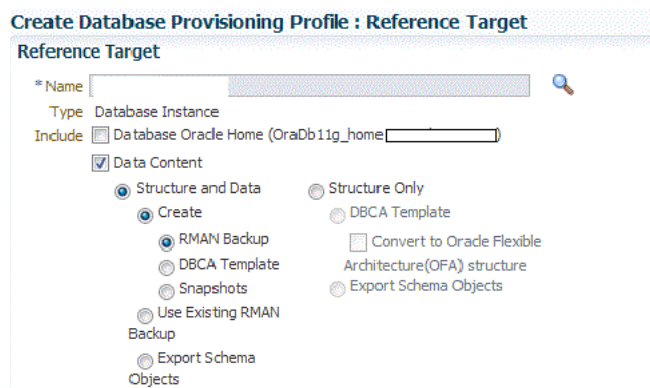
To create a database provisioning profile, follow these steps:

1. You can access the Database Provisioning page using either of the following ways:
  - From the Setup menu, select **Cloud**, and then select **Database**. On the Database Cloud Self Service Portal Setup, select **Profiles and Service Templates**. In the Profiles section, click **Create**.
  - From the Enterprise menu, select **Provisioning and Patching**, and then select **Database Provisioning**. On the Database Provisioning page, in the Profiles section, click **Create**.
2. On the Reference Target page, click the search icon to select a target.


3. In the Search and Select: Targets display box, select the database you want to create a profile for, and then, click **Select**.



4. On the Reference Target page, select the following:
  - **Data Content**
  - **Structure and Data** to include physical and structural files from the database.
  - **Create**
  - **RMAN Backup**



5. In the Credentials section, select the database. In the Credentials column, you can select **Named Credentials** and then select a credential name from the Credential Name column (or) select **Preferred Credentials**. Click **Next**.



6. On the Content Options page, do the following:
  - Select **Online Backup** or **Offline Backup**.
  - Enter the number of channels.
  - (optional) Select **Compress files** and then select **High**, **Medium**, or **Low** strength of compression.
  - (optional) Select **Use File Encryption** and then, enter the password and confirm the password.

- Select **Directory** for the Backup location, and enter the directory path (or) select **Fast Recovery Area**, which already contains a backup location.

---

**Note:** Ensure that the given location is accessible to the nodes in the database pool.

---

- Enter the backup file name format. This will be used to generate backup files.
- Enter the backup file tag. Backup files will be tagged with this tag.
- Enter the control file name. The control file backup will be generated with this name.
- Click **Next**.

7. On the Profile page, do the following:

- In the Profile Information section, enter a unique profile name of your choice.

For example:

RMAN Provisioning Profile for HR database

You can also enter a description of the profile, and notes for your reference.

- In the Schedule section, you can choose to start the profile creation immediately, or you can schedule it for a later time.
- Click **Next**.

8. On the Review page, ensure that the selections you have made in the previous pages are correctly displayed and click **Submit**. Otherwise, click **Back** repeatedly till you reach the page where you want to make changes. Click **Cancel** to abort the provisioning profile creation.

Once you have submitted the provisioning profile creation job, the provisioning profile appears in the Database Provisioning page. The Database Provisioning page displays all the procedure steps.

### 12.9.3 Creating a Database Provisioning Profile Using Existing RMAN Backup

To create a database provisioning profile, follow these steps:

1. You can access the Database Provisioning page using either of the following ways:
  - From the Setup menu, select **Cloud**, and then select **Database**. On the Database Cloud Self Service Portal Setup, select **Profiles and Service Templates**. In the Profiles section, click **Create**.
  - From the Enterprise menu, select **Provisioning and Patching**, and then select **Database Provisioning**. On the Database Provisioning page, in the Profiles section, click **Create**.
2. On the Reference Target page, click the search icon to select a target.

3. In the Search and Select: Targets display box, select the database you want to create a profile for, and then, click **Select**.

4. On the Reference Target page, select the following:
  - **Data Content**

- **Structure and Data** to include physical and structural files from the database or **Structure Only** to include only the structural files in the template.
- **Create**
- **Use Existing RMAN Backup**

5. In the Credentials section, select the database. In the Credentials column, you can select **Named Credentials** and then select a credential name from the Credential Name column (or) select **Preferred Credentials**. Click **Next**.

6. On the Content Options page, select an existing RMAN backup from the list provided.

---

**Note:** The list of existing RMAN backups are fetched from the database.

---

Click **Next**.

7. On the Profile page, do the following:
  - In the Profile Information section, enter a unique profile name of your choice.  
For example:  
Existing RMAN Provisioning Profile mydb.example.com
  - You can also enter a description of the profile, and notes for your reference.
  - In the Schedule section, you can choose to start the profile creation immediately, or you can schedule it for a later time.
  - Click **Next**.

Reference Target Content Options Profile Review

Create Database Provisioning Profile : Profile

Profile Information

- \* Profile Location: Database Provisioning Profiles/11.2.0.3.0/linux\_x64
- \* Profile Name: Profile for 750 GB Bug DB Database
- Description: Database Reference Profile 21-09-2012 02:13 AM from
- Profile Version: 11.2.0.3.0 Vendor: Oracle
- Notes: This is the profile created for bug database. The size of database is around 750 GB.

Schedule

Start: Immediately

Working Directory

\* Path: /tmp

Back Step 3 of 4 **Submit** Cancel

- On the Review page, ensure that the selections you have made in the previous pages are correctly displayed and click **Submit**. Otherwise, click **Back** repeatedly till you reach the page where you want to make changes. Click **Cancel** to abort the provisioning profile creation.

Reference Target Content Options Profile Review

Create Database Provisioning Profile : Review

Reference Target

Target: RMAN Backup

Host Credentials: Host Credentials

Content Details

Recovery Manager(RMAN) Backup Options

- Type: OPLDIE
- Channels: 2
- Compress Files: HGGH
- Use File Encryption: OFF
- Backup location: /scratch/backup
- Keep Local Backup: NONE
- Backup File Name Format: backup\_%.s
- Backup file Tag: data\_backup\_9\_28
- Control File Name: control\_%.s
- Control File Tag: data\_backup\_9\_28

Profile Details

- Profile Location: Database Provisioning Profiles/11.2.0.2.0/linux\_x64
- Profile Name: KSK\_TEST\_JMWAL\_PROFILE
- Description: Database Reference Profile 28-09-2012 02:17 AM from
- Profile Version: 11.2.0.2.0 Vendor: Oracle
- Notes: Host Name Database:

Schedule

Start: Immediately

Working Directory

Path: /tmp

Back Step 4 of 4 **Submit** Cancel

Once you have submitted the provisioning profile creation job, the provisioning profile appears in the Database Provisioning page. The Database Provisioning page displays all the procedure steps.

## 12.9.4 Creating a Database Provisioning Profile Using DBCA Template

To create a database provisioning profile, follow these steps:

- You can access the Database Provisioning page using either of the following ways:
  - From the Setup menu, select **Cloud**, and then select **Database**. On the Database Cloud Self Service Portal Setup, select **Profiles and Service Templates**. In the Profiles section, click **Create**.
  - From the Enterprise menu, select **Provisioning and Patching**, and then select **Database Provisioning**. On the Database Provisioning page, in the Profiles section, click **Create**.
- On the Reference Target page, click the search icon to select a target.

Reference Target Content Options Profile Review

Create Database Provisioning Profile : Reference Target

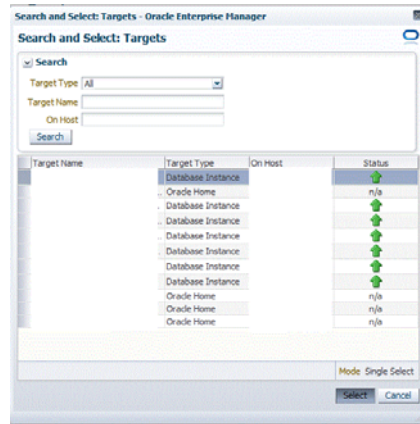
Reference Target

\* Name: [Search Icon]

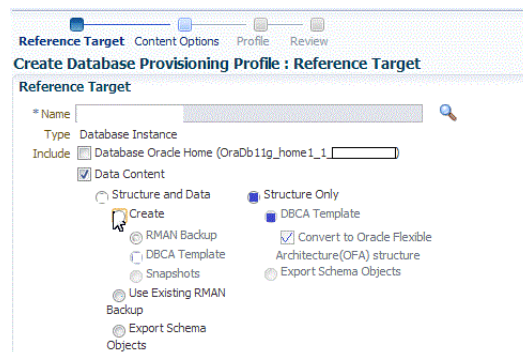
Back Step 1 of 4 Next Cancel

- In the Search and Select: Targets display box, select the database you want to create a profile for, and then, click **Select**.






4. On the Reference Target page, select the following:
  - **Data Content**
  - **Structure Only** to include only the structural files in the template.
  - **Create**
  - **DBCA Template**
  - **Convert to Oracle Flexible Architecture (OFA) structure** to ensure that discrepancy in directory structure in the reference host will not affect the database template creation.



5. In the Credentials section, select the database. In the Credentials column, you can select **Named Credentials** and then select a credential name from the Credential Name column (or) select **Preferred Credentials**. Click **Next**.



6. On the Profile page, do the following:
  - In the Profile Information section, enter a unique profile name of your choice.

For example:

DBCA Provisioning Profile for HR database

You can also enter a description of the profile, and notes for your reference.

- In the Schedule section, you can choose to start the profile creation immediately, or you can schedule it for a later time.
- Click **Next**.

7. On the Review page, ensure that the selections you have made in the previous pages are correctly displayed and click **Submit**. Otherwise, click **Back** repeatedly till you reach the page where you want to make changes. Click **Cancel** to abort the provisioning profile creation.

Once you have submitted the provisioning profile creation job, the provisioning profile appears in the Database Provisioning page. The Database Provisioning page displays all the procedure steps.

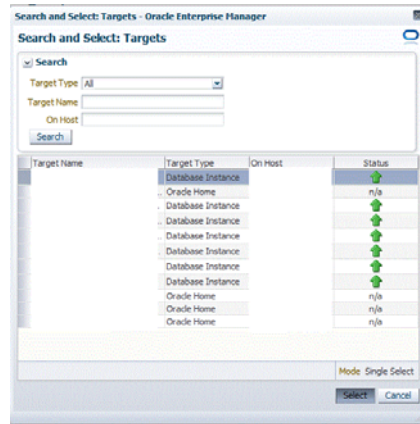
## 12.9.5 Creating a Database Provisioning Profile Using Export Schema Objects

To create a database provisioning profile, follow these steps:

1. You can access the Database Provisioning page using either of the following ways:
  - From the Setup menu, select **Cloud**, and then select **Database**. On the Database Cloud Self Service Portal Setup, select **Profiles and Service Templates**. In the Profiles section, click **Create**.
  - From the Enterprise menu, select **Provisioning and Patching**, and then select **Database Provisioning**. On the Database Provisioning page, in the Profiles section, click **Create**.
2. On the Reference Target page, click the search icon to select a target.

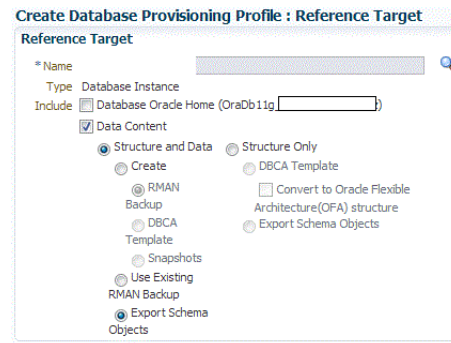
3. In the Search and Select: Targets display box, select the database you want to create a profile for, and then, click **Select**.






4. On the Reference Target page, select the following:

- **Data Content**
- **Structure and Data** to include physical and structural files from the database or **Structure Only** to include only the structural files in the template.
- **Create**
- **Export Schema Objects**



5. In the Credentials section, select the database. In the Credentials column, you can select **Named Credentials** and then select a credential name from the Credential Name column (or) select **Preferred Credentials**. Click **Next**.



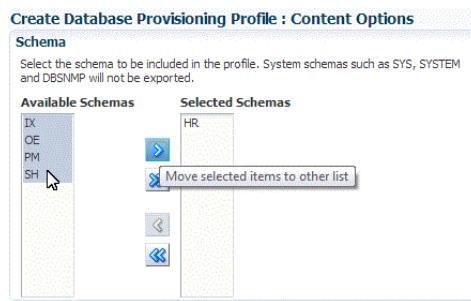
6. On the Content Options page, do the following:

- Select the schema you want to include in the profile by clicking on a schema in the Available Schema list and moving it to the Selected Schema list.

---

**Note:** When one or more schemas are selected, the end user has to ensure the dependent schemas are also selected for export. If the dependent schemas are not exported, the import operation will fail.

---




**Note:** The Available Schemas section displays only the user created schemas that has data content in it. All the empty schemas are filtered out by default.

**Note:** System schemas such as SYS, SYSTEM, and DBSNMP will not be exported.

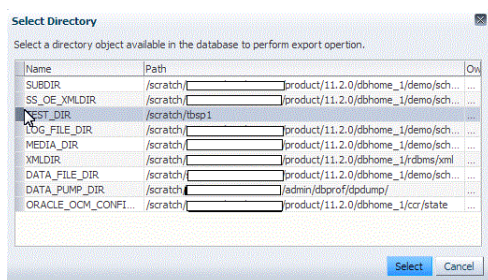
The user schemas which do not have their own tablespace are not shown in left side of the shuttle.

- In the Degree of Parallelism section, enter the number of threads required to perform the export operation.



**Note:** Increasing the number of threads helps reduce the time required to create the profile at the cost of using system resources.

- In the Directory Locations section, click on **Add**. In the Select Directory window, select a directory object listed to perform the export operation, and then click **Select**.



In the Log subsection, click on the search icon. In the Select Directory window, select the directory object you require to perform the export operation, and then click **Select**.

In the Log subsection, enter the file name, and then click on **Next**.

**Directory Locations**

**Dump**  
Specify the directory object, file name, and maximum size for the export files on the database server machine.

+ Add    X Delete...

| Directory | File name        | Maximum file size (MB) |
|-----------|------------------|------------------------|
| TEST_DIR  | export_dump1.dmp |                        |

**Log**

\* Directory: TEST\_DIR    File Name: log1.log

7. On the Profile page, do the following:

- In the Profile Information section, enter a unique profile name of your choice.

For example:

Schema Objects Provisioning Profile for HR database

You can also enter a description of the profile, and notes for your reference.

- In the Schedule section, you can choose to start the profile creation immediately, or you can schedule it for a later time.
- Click **Next**.

8. On the Review page, ensure that the selections you have made in the previous pages are correctly displayed and click **Submit**. Otherwise, click **Back** repeatedly till you reach the page where you want to make changes. Click **Cancel** to abort the provisioning profile creation.

Once you have submitted the provisioning profile creation job, the provisioning profile appears in the Database Provisioning page. The Database Provisioning page displays all the procedure steps.

## 12.10 Creating a Database Service Template

Service templates are standardized service definitions that allow self service users to create databases or schemas. Self service users can create one or more databases based on the service template definition. The service template can be based on:

- Snap Clone Based Profile
- RMAN Backup Based Profile
- DBCA Template Based Profile

### 12.10.1 Creating Service Template Using Snap Clone Profile

To create a snap clone based service template, follow these steps:

1. Log into Enterprise Manager as a user with the EM\_SSA\_ADMINISTRATOR. role.
2. From the **Setup** menu, select **Cloud**, then select **Database**. The Database Cloud Self Service Portal Setup page appears.
3. Click **Profiles and Service Templates**. From the **Create** menu, select **For Database**. The Create Service Template: General page appears.

**Figure 12–4 Create Service Template: General (Snap Clone)**

**Create Service Template**

General Database Initialization Parameters Scripts Zones Roles Review

Create Service Template : General

Back Step 1 of 7 Next Cancel

**General**

\* Name: Create SIDB Using Snapshot Netapp Technology

Description:

\* Profile: Snapshot profile t112 database v2 - SIDB

**Database Type**

Type: ☒ Single Instance ☐ Real Application Clusters (RAC)

Number of nodes: 1

**Identification**

Specify a prefix that should be used to generate a unique System Identifier(SID) at the time of database creation. The prefix helps to identify databases, which are created using this service template.

\* SID Prefix: \* Domain Name:

E.g., If the SID Prefix is 'DB' and Domain name is 'mycompany.com', the SID may be generated as 'DB23456.mycompany.com'

**Listener Port**

\* Port Number:

4. Enter a name and description for the service template. The description must be unique and provide information on the type of service template being created.
5. Click the **Search** icon next to the **Profile** field and select a snap clone profile from the list.
6. Select the Database Type which can be Single Instance or Real Application Cluster (RAC). If you select RAC, specify the Number of Nodes.
7. In the Identification region, specify the following:

- **SID Prefix:** Enter a prefix that is to be used to generate a unique System Identifier (SID) at the time of database creation. The prefix helps to identify databases created using this service template. The prefix can be a maximum of 6 characters.

The new database name generated will be based on the SID Prefix specified here. For example, if the prefix is specified as MYDB, the SID for the new database is generated as MYDB0000, MYDB0001, MYDB0002, and so on.

For existing databases (running databases and in progress requests), a unique SID name is generated for single instance databases, or a unique name is generated for real application cluster databases. For example, if the existing SIDs running on a host are MYDB0000, MYDB0001, MYDB0002 And In progress are MYDB0003, MYDB0004, the new SID or database name generated for the new request is MYDB0005.

- **Domain Name:** Enter a Domain Name to be used for the new database being created. You must ensure that the domain name does not contain a preceding "." (dot).
8. In the Listener Port field, specify listener port number that is to be associated with the new database. If you have selected a 11.2 or later RAC database, you must specify the scan port number. Click **Next**. The Create Service Template: Database page appears.

**Figure 12–5 Create Service Template: Database (Snap Clone)**

**Create Service Template**

General **Database** Initialization Parameters Scripts Zones Roles Review

**Create Service Template : Database** Back Step 2 of 7 Next Cancel

**Storage size**  
Configure the volumes that will be used by the database

Edit...

| Volume | Mount Point Prefix | Writable Space (GB) |
|--------|--------------------|---------------------|
|        | /fsk               | 2.00                |
|        | /fsk               | 2.00                |

**Snapshot Policy**  
Allow the user to take snapshots for the new database  
Maximum number of snapshots 5

**Working Directory**  
\* Path /tmp

**Administrator Credentials**  
Specify passwords for the administrative users (SYS, SYSTEM and DBSNMP) in the new database. These users are used by the SSA Administrator to manage the database. The SSA User has no access to these users.

☒ Use the same password  
Password \*\*\*\*\* Confirm Password \*\*\*\*\*

☐ Use different passwords

| User Name | Password | Confirm Password |
|-----------|----------|------------------|
| SYS       | *****    | *****            |
| SYSTEM    | *****    | *****            |
| DBSNMP    | *****    | *****            |

**Non Administrator Credentials**  
Select users that should not be accessible to the SSA User.

| Users      | SCOTT |
|------------|-------|
| OUTLN      |       |
| ORACLE_OCM |       |
| SH         |       |
| HR         |       |
| IX         |       |
| OE         |       |
| APPSYS     |       |
| WMSYS      |       |
| EXFSYS     |       |
| CTXSYS     |       |
| YDR        |       |

\*\*\*\*\* Confirm \*\*\*\*\*

9. Specify the following details:

- **Storage Size:** The volumes defined in the profile you have selected are displayed. Select a volume and click **Edit** to configure the volumes that will be used by the database:
  - **Mount Point Prefix:** The prefix for the new mount point for the database being provisioned. The volume will be mounted on `/<mount-prefix>_<unique-suffix-generated-by-EM>/<source-path>`.
  - **Writable Space:** The amount of writable space available on the database. The value specified here cannot exceed the Storage Ceiling specified for the server.
- **Snapshot Policy:** Snapshots are sets of historical data for specific time periods. By default, Oracle Database automatically generates snapshots of the performance data once every hour and retains the statistics in the workload repository for 8 days. Select the **Allow the user to take snapshots** for the new database checkbox and specify the **Maximum Number of Snapshots** that can be taken.
- **Working Directory:** Specify the directory in which the files are temporarily stored.
- **Administrator Credentials:** Specify passwords for the system schemas of the new database. These schemas will not be available to the `EM_SSA_USERS`. You can choose to use the same password for all the schemas or different passwords for each schema. If these values are not specified, default values will be used and the `EM_ADMINISTRATOR` can change them.
- **Non Administrator Credentials:** Apart from the system schemas, if you want to restrict access to other schemas, you can select them in the Non-Administrator Credentials region and specify the password.

10. Click **Next**. The Create Service Template: Initialization Parameters page appears.

**Figure 12–6 Create Service Template: Initialization Parameters (Snap Clone)**

**Create Service Template**

General Database **Initialization Parameters** Scripts Zones Roles Review

**Create Service Template : Initialization Parameters** Back Step 3 of 7 Next Cancel

View ▾ Detach Set

| Name                      | Value      | Editable | Modified | Description                                                                     |
|---------------------------|------------|----------|----------|---------------------------------------------------------------------------------|
| ✓ Audit                   |            |          |          |                                                                                 |
| audit_trail               | db         | Yes      | False    | Enable system auditing. Valid Values are none,os,db,(db,extended),xml,(xml,...) |
| ✓ Common                  |            |          |          |                                                                                 |
| db_block_size             | 8192       | No       | False    | Size of database block in bytes. Multiple of 512. Example 8192.                 |
| db_cache_size             | 312MB      | Yes      | False    | Size of DEFAULT buffer pool for standard block size buffers. Examples 0.5GB     |
| processes                 | 150        | Yes      | False    | User processes. Example 200.                                                    |
| compatible                | 11.2.0.0.0 | No       | False    | Database will be completely compatible with this software version.              |
| open_cursors              | 300        | Yes      | False    | Maximum # cursors per session. Example 500.                                     |
| remote_login_passwordfile | EXCLUSIVE  | Yes      | False    | Password file usage parameter. Valid Values are EXCLUSIVE,NONE,INTERNAL         |
| > Memory                  |            |          |          |                                                                                 |
| > Pool Size               |            |          |          |                                                                                 |
| > Recovery                |            |          |          |                                                                                 |
| > Undo Management         |            |          |          |                                                                                 |

11. In this page, you can configure the values of various initialization parameters that affect the operation of the database instance. Select the parameter and click the **Set** icon to modify the value of the parameter.
12. Click **Next**. The Create Service Template: Scripts page appears. Specify the custom scripts that need to be executed before and after the database is created. See [Section 12.13, "Pre and Post Request Creation / Deletion Scripts"](#) for details.
13. Click **Next**. The Create Service Template: Zones page appears.

**Figure 12–7 Create Service Template: Zones (Snap Clone)**

**Create Service Template : Zones** Back Step 5 of 7 Next Cancel

**Zones**

A service template can be configured to provision databases in one or more pools. Select the zone and associated pools that this service template can provision databases into.

+ Add - Remove Assign Pool

| Name                                                                     | Resource Pool |
|--------------------------------------------------------------------------|---------------|
| Salt Lake City Zone - SLC - For SNAPSHOT SIDB Pool for Linux 64 Machines |               |

14. Click **Add** to select a PaaS Infrastructure Zone into which the database instance is to be provisioned. Click **Assign Pool**. Select a pool from the list and assign this pool to the PaaS Infrastructure Zone. The database will be provisioned into this pool.
15. Click **Next**. The Create Service Template: Roles page appears.

**Figure 12–8 Create Service Template: Roles (Snap Clone)**

**Create Service Template : Roles** Back Step 6 of 7 Next Cancel

**Roles**

A service template can be configured for one or more roles in the Enterprise Manager.

+ Add - Remove

| Role Name     | Role Description |
|---------------|------------------|
| SSA_USER_ROLE |                  |

16. Click **Add** to select the SSA user roles to which this service template will be available. All users belonging to the selected role can use this service template. Click **Next**.
17. The Create Service Template: Review page appears. Click **Submit**. The newly created service template will appear in the Database Cloud Self Service Portal Setup: Service Templates page. You can click on the Service Template Name link to view additional details.

## 12.10.2 Create Service Template Using RMAN Backup Profile

1. Log into Enterprise Manager as a user with the EM\_SSA\_ADMINISTRATOR role.
2. From the **Setup** menu, select **Cloud**, then select **Database**. The Database Cloud Self Service Portal Setup page appears.
3. Click **Profiles and Service Templates**. From the **Create** menu, select **For Database**. The Create Service Template: General page appears.

**Figure 12–9 Create Service Template: General (RMAN Backup)**

4. Enter a name and description for the service template. The description must be unique and provide information on the type of service template being created.
5. Click the **Search** icon next to the **Profile** field and select a snap clone profile from the list.
6. Select the Database Type which can be Single Instance or Real Application Cluster (RAC). If you select RAC, specify the Number of Nodes.
7. In the Identification region, specify the following:
  - **SID Prefix:** Enter a prefix that is to be used to generate a unique System Identifier (SID) at the time of database creation. The prefix helps to identify databases created using this service template. The prefix can be a maximum of 6 characters.

The new database name generated will be based on the SID Prefix specified here. For example, if the prefix is specified as MYDB, the SID for the new database is generated as MYDB0000, MYDB0001, MYDB0002, and so on.

For existing databases (running databases and in progress requests), a unique SID name is generated for single instance databases, or a unique name is generated for real application cluster databases. For example, if the existing SIDs running on a host are MYDB0000, MYDB0001, MYDB0002 And In progress are MYDB0003, MYDB0004, the new SID or database name generated for the new request is MYDB0005.



- **Domain Name:** Enter a Domain Name to be used for the new database being created. You must ensure that the domain name does not contain a preceding "." (dot).
8. In the Listener Port field, specify listener port number that is to be associated with the new database. If you have selected a 11.2 or later RAC database, you must specify the scan port number. Click **Next**. The Create Service Template: Database page appears.

**Figure 12–10 Create Service Template: Database (RMAN Backup)**

**Create Service Template**

General **Database** Initialization Parameters Scripts Zones Roles Review

**Create Service Template : Database**

**Storage type**

☐ Automatic Storage Management

☒ File System

Location: /scratch/s1me/orabase11202\_2/ora

**Fast Recovery**

☒ Specify Fast Recovery Area

☒ Automatic Storage Management

Fast Recovery Area: /OracleEM/ThinProv/slc0Dec/slot01

☐ File System

Location:

Fast Recovery Size(MB):

☐ Enable Archiving

**Working Directory**

\* Path: /tmp

**Administrator Credentials**

Specify passwords for the administrative users (SYS, SYSTEM and DBSNMP) in the new database. These users are used by the SSA Administrator to manage the database. The SSA User has no access to these users.

☒ Use the same password

Password: Password Confirm Password

☐ Use different passwords

| User Name | Password | Confirm Password |
|-----------|----------|------------------|
| SYS       | Password | Confirm Password |
| SYSTEM    | Password | Confirm Password |
| DBSNMP    | Password | Confirm Password |

**Non Administrator Credentials**

Select users that should not be accessible to the SSA User.

| Users      | SCOTT |
|------------|-------|
| OUTLN      |       |
| ORACLE_OCM |       |
| SH         |       |
| HR         |       |
| IX         |       |
| OE         |       |
| APPQOSSYS  |       |
| WM SYS     |       |
| EXFSYS     |       |
| CTXSYS     |       |
| VNR        |       |

Password: Password Confirm Password

9. Specify the following:
- **Storage Type:** This can be:
    - **Automatic Storage Management:** The Oracle Automatic Storage Management (ASM) is a volume manager and a file system for database files that supports single-instance and RAC configurations. ASM groups the disks in your storage system into one or more disk groups. If you select ASM, specify the Disk Group here.
    - **File System:** The Oracle Database File System creates a standard file system interface on top of files and directories that are stored in database tables. If you select this option, you must specify the Location of the File System.
  - **Fast Recovery (Optional):** To simplify the management of backup and recovery files, you can create a fast recovery area for your database. The fast recovery area can be a ASM disk group or a file system that provides a centralized disk location for backup and recovery file. Specify the location of the Fast Recovery Area and the Fast Recovery Size. The amount of disk space to allocate for the fast recovery area depends on the size and activity levels of your database.
 

Select the **Enable Archiving** check box if the backups are to be archived.
10. Specify the **Working Directory** in which the files are temporarily stored before the database is provisioned.



11. Specify the **Administrator Credentials**. Specify passwords for the system schemas of the new database. These schemas will not be available to the EM\_SSA\_USERS. You can choose to use the same password for all the schemas or different passwords for each schema. If these values are not specified, default values will be used and the EM\_ADMINISTRATOR can change them.
12. Apart from the system schemas, if you want to restrict access to other schemas, you can select them in the Non-Administrator Credentials region and specify the password. These schemas will be locked and the EM\_SSA\_USERS cannot access them. Click **Next**.
13. Follow steps 10 to 17 in the [Section 12.10.1, "Creating Service Template Using Snap Clone Profile"](#) section to create the RMAN Backup based service template.

### 12.10.3 Create Service Template Using DBCA Template Profile

1. Log into Enterprise Manager as a user with the EM\_SSA\_ADMINISTRATOR role.
2. From the **Setup** menu, select **Cloud**, then select **Database**. The Database Cloud Self Service Portal Setup page appears.
3. Click **Profiles and Service Templates**. From the **Create** menu, select **For Database**. The Create Service Template: General page appears.

**Figure 12–11 Create Service Template: General (DBCA Template)**

The screenshot displays the 'Create Service Template: General' page. The 'General' tab is active, showing a form with the following fields: 'Name' (dbca\_template), 'Description', and 'Profile' (DBCA Template Profile for t112 database - SIDB). The 'Database Type' section has 'Single Instance' selected. The 'Identification' section has 'SID Prefix' set to 'DB' and 'Domain Name' set to 'us.oracle.com'. The 'Listener Port' section has 'Port Number' set to '1521'. Navigation buttons 'Back', 'Step 1 of 7', 'Next', and 'Cancel' are at the top right.

4. Enter a name and description for the service template. The description must be unique and provide information on the type of service template being created.
5. Click the **Search** icon next to the **Profile** field and select a snap clone profile from the list.
6. Select the Database Type which can be Single Instance or Real Application Cluster (RAC). If you select RAC, specify the Number of Nodes.
7. In the Identification region, specify the following:
  - **SID Prefix:** Enter a prefix that is to be used to generate a unique System Identifier (SID) at the time of database creation. The prefix helps to identify databases created using this service template. The prefix can be a maximum of 6 characters. The new database name generated will be based on the SID Prefix specified here. For example, if the prefix is specified as MYDB, the SID for the new database is generated as MYDB0000, MYDB0001, MYDB0002, and so on.

For existing databases (running databases and in progress requests), a unique SID name is generated for single instance databases, or a unique name is generated for real application cluster databases. For example, if the existing SIDs running on a host are MYDB0000, MYDB0001, MYDB0002 And In progress are MYDB0003, MYDB0004, the new SID or database name generated for the new request is MYDB0005.

- **Domain Name:** Enter a Domain Name to be used for the new database being created. You must ensure that the domain name does not contain a preceding "." (dot).
8. In the Listener Port field, specify listener port number that is to be associated with the new database. If you have selected a 11.2 or later RAC database, you must specify the scan port number. Click **Next**. The Create Service Template: Database page appears.

**Figure 12–12 Create Service Template: Database (DBCA Template)**

| User Name | Password | Confirm Password |
|-----------|----------|------------------|
| SYS       | *****    | *****            |
| SYSTEM    | *****    | *****            |
| DBSNMP    | *****    | *****            |

9. Specify the following:
- **Storage Type:** This can be:
    - **Automatic Storage Management:** The Oracle Automatic Storage Management (ASM) is a volume manager and a file system for database files that supports single-instance and RAC configurations. ASM groups the disks in your storage system into one or more disk groups. If you select ASM, specify the Disk Group here.
    - **File System:** The Oracle Database File System creates a standard file system interface on top of files and directories that are stored in database tables. If you select this option, you must specify the Location of the File System.
  - **Fast Recovery (Optional):** To simplify the management of backup and recovery files, you can create a fast recovery area for your database. The fast recovery area can be a ASM disk group or a file system that provides a centralized disk location for backup and recovery file. Specify the location of the Fast Recovery Area and the Fast Recovery Size. The amount of disk space to allocate for the fast recovery area depends on the size and activity levels of your database.

Select the **Enable Archiving** check box if the backups are to be archived.

10. Specify the **Working Directory** in which the files are temporarily stored before the database is provisioned.
11. Specify the **Administrator Credentials**. Specify passwords for the system schemas of the new database. These schemas will not be available to the EM\_SSA\_USERS. You can choose to use the same password for all the schemas or different passwords for each schema. If these values are not specified, default values will be used and the EM\_ADMINISTRATOR can change them.
12. Apart from the system schemas, if you want to restrict access to other schemas, you can select them in the Non-Administrator Credentials region and specify the password. These schemas will be locked and the EM\_SSA\_USERS cannot access them. Click **Next**.
13. Follow steps 10 to 17 in the [Section 12.10.2, "Create Service Template Using RMAN Backup Profile"](#) section to create the RMAN Backup based service template.

## 12.11 Configuring Chargeback

Click the **Chargeback** tab to add the database pools to the Chargeback Application. On this page, you can determine chargeback zones, services, items and prices. After the Chargeback is setup, you can view charges and plan details. For more details on setting up the Chargeback Service, see [Section 22, "Chargeback Administration"](#).

## 12.12 Using Schema as a Service to Create Schemas

Enterprise Manager Schema as a Service feature allows you to create one or more schemas on a running database. Follow these steps to create a schema on a database:

1. Create a database pool. See [Section 12.5, "Creating a Database Pool"](#).
2. Configure the request settings. See [Section 12.6, "Configuring Request Settings"](#).
3. Define quotas for each self service user. See [Section 12.7, "Setting Up Quotas"](#).
4. Create a database provisioning profile. This step is optional and is required if you are importing schemas from a profile.
5. Create a service template. A service template can contain:
  - User Defined Schema: In this case, the service template is created with an empty schema.
  - Schema from a Profile: In this case, you can import schemas from a database provisioning profile. You can select this option for applications with data such as eBusiness applications. See [Section 12.9.5, "Creating a Database Provisioning Profile Using Export Schema Objects"](#).

See [Section 12.12.2, "Creating a Service Template for Schema as a Service"](#).

6. Configure the Chargeback service. See [Section 22, "Chargeback Administration"](#).
7. Create the schema based on a specified service template. See [Section 13.2.1, "Requesting a Schema"](#).

### 12.12.1 Creating a Database Pool for Schema as a Service

You can create a database service with one or more schemas with or without seed data and database objects. To do so, follow these steps:

1. Log into Enterprise Manager as a user with the EM\_SSA\_ADMINISTRATOR role.

2. From the **Setup** menu, select **Cloud**, then select **Database**. The Database Cloud Self Service Portal Setup page appears.
3. From the **Create** menu, select **For Schema**.
4. Enter a name and description for the database pool.
5. In the Credentials region, specify the credentials that will be used for performing database operations. Specify the following credentials:

- **Host:** Required for starting the database.
- **Database:** Required for creating schemas on the database.
- **Grid Infrastructure:** Required for cluster databases.

After you have specified the credentials, click **Test All Credentials** to verify them.

6. Specify the following details:
  - **PaaS Infrastructure Zone:** Select the PaaS Infrastructure Zone in which the databases are present. The schemas will be created on the databases selected for the pool.
  - **Target Type:** Specify the configuration of the target database on which the schema is being deployed. This can be Single Instance or Cluster Database (RAC).
  - **Platform and Version:** Specify the platform and version of the database being deployed.
7. Click **Add** and select one or more databases to be added to the pool. All databases you select must be present in the same PaaS Infrastructure Zone.

---

**Note:** You can search for database groups that meet your requirements by specifying a search string in the Groups field.

---

8. Click **Next**. The Create New Pool: Policies page appears. Specify the placement policy constraints to set maximum ceilings for resource utilization on each database. Placement policy constraints for each database can be defined by services or workloads as follows:
  - **Services:** If you select this check box, you must specify the Maximum Number of Database Services that can be running on each database.
  - **Workloads Associated with Service Requests:** The amount of resources (such as CPU and memory) to manage the workload can be specified here. Specify the total workload for each database.
    - **Maximum CPU Allocation:** Specify the maximum CPU that can be allocated to each database.
    - **Maximum Memory Allocation:** Specify the maximum memory that can be allocated for each database.
9. If you have chosen to set the placement constraints by **Workloads Associated with Service Requests**, you can choose to enable the resource manager by selecting the **Enable Resource Manager for CPU** check box. You can use the resource manager to manage the CPU resource and ensure that `EM_SSA_USER` will not exceed the maximum allocated value.
10. Click **Submit** to create a database pool. The newly created database pool will appear in the Database Cloud Self Service Portal Setup: Database Pools page.

## 12.12.2 Creating a Service Template for Schema as a Service

One or more schemas can be created based on the service template definition. To create a service template for schema as a service, follow these steps:

1. Log into Enterprise Manager as a user with the `EM_SSA_ADMINISTRATOR` role.
2. From the **Setup** menu, select **Cloud**, then select **Database**. The Database Cloud Self Service Portal Setup page appears.
3. Click **Profiles and Service Templates**. In the Service Templates region, select **For Schema** from the **Create** menu. The Create Service Template: Name page appears.

**Figure 12-13 Create Service Template: Name (Schema)**

4. Enter a name and description for the service template. The description must be unique and provide information on the type of service template being created.
5. In the Schemas region, you can choose to:
  - **Create User Defined Schemas:** Select this option to create empty schemas which can be configured by the user. If you select this option, you can specify the Maximum Number of Schemas that can be created.
  - **Import Schemas from Profile:** Select a database provisioning profile from which the schema is to be imported. Click the **Search** icon next to the Profile field and select an export schema objects based profile from the list. If a profile does not exist, you must create one. See [Section 12.9.5, "Creating a Database Provisioning Profile Using Export Schema Objects"](#).
    - **Master Account:** If you select the **Import Schemas from Profile** option, you must select the Master Account that has privileges on all the schemas. If the Master Account is not selected, the `EM_SSA_USER` will select it when requesting a schema.
6. Click **Next**. The Create New Service Template: Schema Details page appears.

**Figure 12–14 Create Service Template: Schema Details (Schema)**

**Service Templates**

Name Schema Details Scripts Zones Roles Review

Create New Service Template : Schema Details Back Step 2 of 6 Next Cancel

**Workload Size**  
Workload represents the expected CPU, memory and storage requirements for each service.

View Create... Edit Delete

| Name                    | Description | CPU (cores) | Memory (GB) | Storage (GB) |
|-------------------------|-------------|-------------|-------------|--------------|
| No workloads specified. |             |             |             |              |

**Role**  
The database role to which all the privileges will be granted.

\* Name

\* Privilege for all accounts

☒ **TIP** Master account has complete access to all other schemas created as part of service request

**Tablespace**  
Specify the tablespace storage configurations.

Maximum Size ☒ Unlimited by the workload size selected at request time  
☐ Specified by the workload size selected at request time

☒ Auto Extend

Initial Size Workload size selected during request

Increment  GB

Specify the following details:

- **Workload Size:** Workload represents the total CPU, memory, and storage requirements for each service. You can define workload sizes such as small, medium, and large. Click **Create** and in the Create Workload window, enter the name, CPU, Memory, and Storage limit for each service.
- **Role:** Specify the role to which the database privileges such as `CREATE_SESSION`, `CREATE_DIMENSION`, and so on will be granted. Enter the name of the user and the privileges that will be granted to the user.

---

**Note:**

- If the role with the name specified already exists in the database, the database privileges cannot be modified.
  - The Master Account will have the `CREATE_SESSION`, `CREATE_DIMENSION`, `CREATE_INDEXTYPE`, `CREATE_OPERATOR`, `CREATE_PROCEDURE`, `CREATE_SEQUENCE`, `CREATE_TABLE`, `CREATE_TRIGGER`, `CREATE_TYPE`, `CREATE_VIEW`, and `CREATE_SYNONYM` privileges over all the schemas created as part of this service request.
- 
- **Tablespace:** Specify the tablespace storage configuration. Enter the following details:
    - **Maximum Size:** The maximum size of the tablespace can be **Unlimited** or **Specified by the workload size selected at request time**.
    - **Auto Extend:** Select this check box to automatically extend the new data file.
    - **Initial Size:** If you selected **Specified by the workload size selected at request time** in the Maximum Size field, enter the initial size of the tablespace.
    - **Increment:** If you have selected the **Auto Extend** checkbox, specify the size by which the data file is to be automatically incremented.

7. Click **Next**. The Create Service Template: Scripts page appears. Specify the custom scripts that need to be executed before and after the service instance has been created.

**Figure 12–15 Create Service Template: Scripts (Schema)**

8. Click **Next**. The Create Service Template: Zones page appears.

**Figure 12–16 Create Service Template: Zones (Schema)**

9. Click **Add** to select a PaaS Infrastructure Zone into which the service instance is to be provisioned. Click **Assign Pool** to add a Database Pool for the zone.
10. If you creating a service template based on an export schema objects profile, in the Dump File Location region, specify the Path or the shared location on all hosts in which the schema dumps are present. The Dump File Location must be accessible by all the databases.
11. Click **Next**. The Create Service Template: Roles page appears. Click **Add** to select the SSA user roles to which this service template will be available. All users belonging to the selected role can use this service template.
12. Click **Submit**. The newly created service template will appear in the Database Cloud Self Service Portal Setup: Service Templates page. You can click on the Service Template Name link to view additional details.

## 12.13 Pre and Post Request Creation / Deletion Scripts

You can run custom scripts before and after a service instance has been created. These requests can range from some additional host commands on the machine on which the

database was created or commands to perform additional book activities and other operations on the newly provisioned database.

Similarly, you can run scripts after a service instance has been deleted if you need to undo the changes made during service cleanup.

The custom script must follow certain guidelines and consists of four parts:

- **Response File:** The response file contains the service template and request specific information. Before the custom script is executed, the request procedure, the request procedure generates a response file (file with name-value pairs) containing the data from the service template as well as the request-specific data such as the SID that is auto computed, the hosts on which the database will be deployed, and so on.
- **Driver Script:** This is the key driver script that will be invoked by the request procedure. The driver script accepts only one argument which specifies the location of the response file.
- **Other Scripts:** Apart from the driver script, you can specify other perl or sql scripts. These scripts must be invoked from the driver script.
- **Additional Variables:** You can include a file containing additional variables that can be used by the custom script as applicable.

After the scripts have been created, they must be uploaded as a directive into the Software Library. The directive details are shown below:

**Figure 12–17 Pre or Post Database Script: Describe**

**View Entity: Pre or Post Database Script (0.1)**

Describe | Configure | Select Files

Name: Pre or Post Database Script  
Description: Pre or Post Database Script  
Type: Directive  
Subtype:   
Parent Directory: Samples  
Maturity: Production  
Status: Ready

**Other Attributes**

| Name            | Value | Description                               |
|-----------------|-------|-------------------------------------------|
| Product Version |       | Product version of the software component |
| Product         |       | Product the software component represents |
| Vendor          |       | Vendor of the software component          |

**Attachments**  
No attachment has been added yet.

**Notes**  
No note has been added yet.

In the Describe page, the name of the directive and the description is displayed. Click the **Configure** tab.

**Figure 12–18 Pre or Post Database Script: Configure**

**Software Library**

Software Library > View Entity: Pre or Post Database Script (0.1)

**View Entity: Pre or Post Database Script (0.1)**

Describe | Configure | Select Files

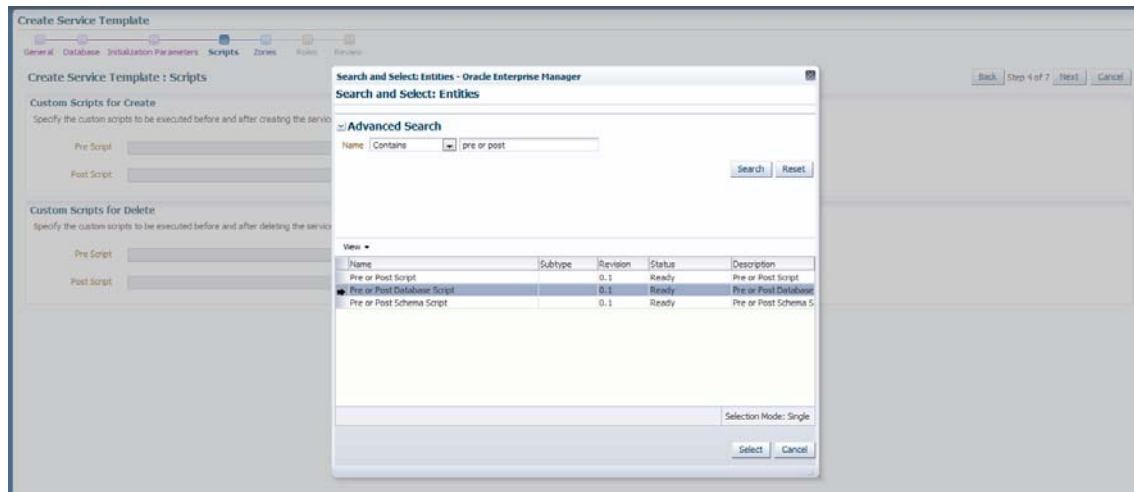
Configure Directive:

**Command Line Arguments**  
Command Line: %INPUT\_FILE% The directive can only take one argument and it MUST be named as INPUT\_FILE as shown here.

**Configuration Properties**  
Shell Type: Perl  
Run Privileged: ☐

The driver script accepts one command line argument which must be in the INPUT\_FILE format. This variable will be used at run-time to specify the location of the generated response file. Click the **Select Files** tab.



**Figure 12–19 Pre or Post Database Script: Select Files**

Specify any additional files that are required by the custom script. You can run the script from the same host on which the database instance or the database service was created or deleted.

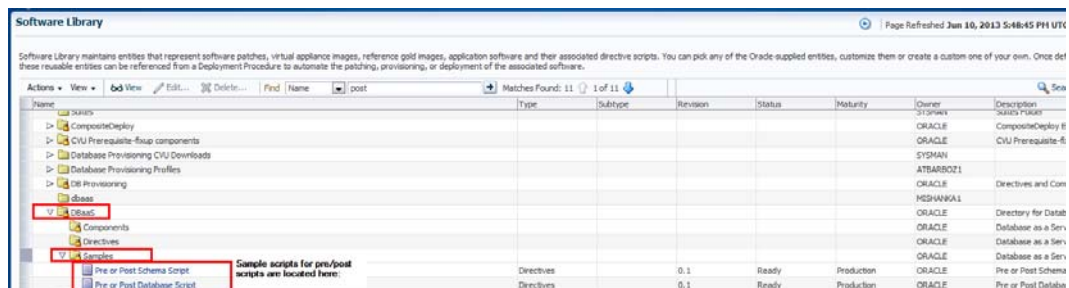
---

**Note:** if you need to change the content of the script, you must upload a newer version of the script to the Software Library. To use the latest version of the script, you must launch the Edit Service Template wizard and select the updated version of the script and save the template. This ensures that the latest version of the script will be used.

---

### 12.13.1 Sample Scripts

Sample scripts are available in the Software Library. To view the sample scripts, from the **Enterprise** menu, select **Provisioning and Patching**, then select **Software Library**. Select the **DBaaS** folder, then select **Directives**, and click the **Samples** folder to see the custom scripts.

**Figure 12–20 Sample Scripts**



## Using the DBaaS Self Service Portal

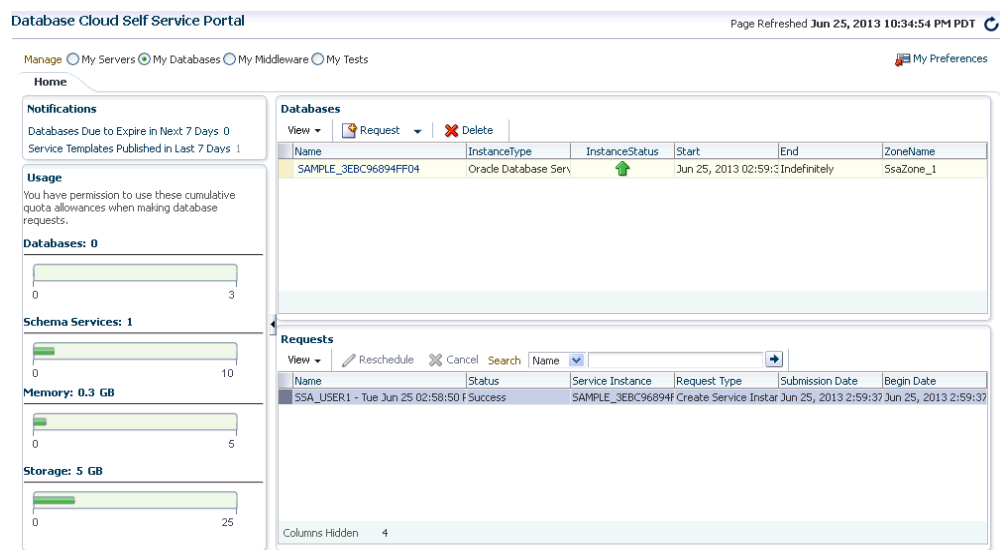
This chapter provides self service users with instructions on using the Database Cloud Self Service Portal to request, monitor, and manage database services. It contains the following sections:

- [Using the Database Cloud Self Service Portal](#)
- [Requesting Databases and Schemas](#)
- [Viewing the Database Service Home Page](#)
- [Viewing the Database Instance Home Page](#)
- [Viewing the Cluster Database Home Page](#)

### 13.1 Using the Database Cloud Self Service Portal

The Database Self Service Portal allows self service users to view, monitor and deploy databases on selected zones, as well as create schemas on running databases. To view the Database Cloud Self Service Portal, log in to Enterprise Manager as a user with EM\_SSA\_USER role. From the Enterprise menu, select **Cloud**, then select **Self Service Portal**, and click the **My Databases** radio button.

**Figure 13–1 Database Cloud Self Service Portal**



The following details are displayed:

- **Home:** This is the Home page for the Database Self Service Portal. It contains the following sections:
  - **Notifications:** This section shows the number of databases that are expiring in the next 7 days.
  - **Usage:** This region displays the quota that you currently have and how much has been used for the following:
    - \* **Databases:** The number of databases requested so far and the maximum number of databases that can be requested.
    - \* **Schemas:** The number of schemas requested so far and the maximum number of schemas that can be requested.
    - \* **Memory:** The total amount of memory allocated to the user and the memory used by all databases and schemas requested so far.
    - \* **Storage:** The total storage volumes allocated to the user and the storage used by all databases requested so far.
  - **Databases:** This table lists the services that you currently own. You can request more databases by clicking **Request Database**. You can also delete any databases you own to release it prior to expiration. Click on a **Service Name** link to drill down to the database details page.
  - **Requests:** This table tracks the requests that you have created. You can also delete requests that have been scheduled but not executed.
- **Chargeback:** This page displays all the chargeback break-up data for the databases and pricing for different zones. The charge incurred will be calculated from the database requested date to current date. This page contains the following sections:
  - **Charge Trend:** This graph show the usage and charge trends over time.
  - **Aggregate By:** You can aggregate the charge back data for the database by resources or metrics. You can export this data to an external application such as a spreadsheet.
  - **Charge Plans:** The charge plans that have been assigned to each zone are displayed. Select a zone to view the pricing. The rate for each metric in the database zone for a specific time period is displayed.
- **My Preferences:** Click on this link to set up your preferences such as Default Database Zone for New Requests, Default Email Address for Notifications, and so on.
- **My Preferences:** Click the **My Preferences** link to set up your preferences such as the Default Service Type, Default Email Address for Notifications, Default Locale, Default PaaS Infrastructure Zone, and so on. To set the Database Cloud Self Service Portal as the default page that appears when you log in to the Self Service Portal, select **Database** in the Default Service Type drop down list and click **Apply**. For more details on setting the preferences, see [Section 7.1.7, "Viewing My Preferences"](#).

## 13.2 Requesting Databases and Schemas

Using the Database Cloud Self Service Portal, you can:

- **Request a Schema:** You can create a database service with one or more schemas and populate the schema with the required data. See [Section 13.2.1, "Requesting a Schema"](#).

- **Request a Database:** You can deploy databases to a specific zone with a selected service template See [Section 13.2.2, "Requesting a Database"](#).

### 13.2.1 Requesting a Schema

You can create a database service with one or more schemas and populate the schema with the required data. To request a schema, follow these steps:

1. Log in to Enterprise Manager as a user with EM\_SSA\_USER role or any role that includes EM\_SSA\_USER role.
2. The Infrastructure Cloud Self Service Portal page appears. Click the **My Databases** radio button to navigate to the Database Cloud Self Service Portal.
3. From the **Request** menu in the Databases region, select **Schema**.
4. Choose a Schema Service Template from the list and click **Select**. The New Database Service Request page appears. The name and description of the service template you have selected is displayed. Enter the following details:
  - **Request Name:** Enter a name for the schema service request.
  - **Zone:** Select the zone in which the schema is to be created.
  - **Workload Size:** Specify the workload size for the service request.
  - **Schema Prefix:** Enter a prefix for the schema. For clustered databases, the service is registered with Grid Infrastructure credentials.
5. Click **Rename Schema** to enter a new name for the schema. If you wish to retain the source schema name, ensure that the Schema Prefix field is blank.
6. Specify the password for the schema. Select the **Same Password for all Schemas** checkbox to use the same password for all the schemas.
7. If you selected a service template based on an empty schema, you must select the schema with Master Account privileges. This schema will have access to all other schemas created as part of the service request.
8. Specify the schedule for the request.
9. Click **Submit** to deploy the database to the selected zone.

### 13.2.2 Requesting a Database

You can deploy databases to a specific zone with a selected service template. To request a database, follow these steps:

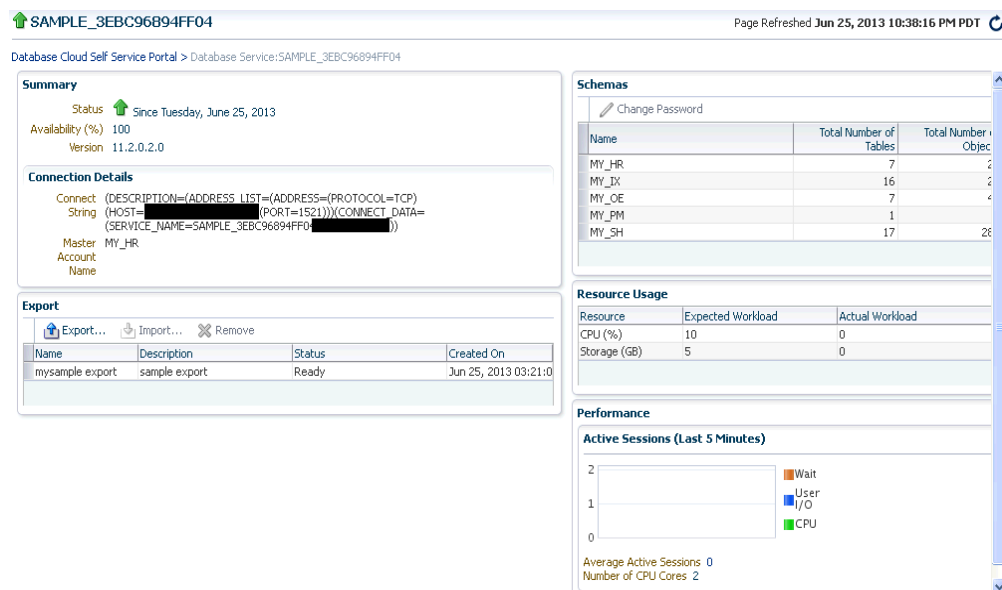
1. Login to Enterprise Manager as a user with EM\_SSA\_USER role or any role that includes EM\_SSA\_USER role.
2. The Infrastructure Cloud Self Service Portal page appears. Click the **My Databases** radio button to navigate to the Database Cloud Self Service Portal.
3. Click **Request Database** in the My Databases region.
4. Choose a service template from the list and click **Select**.
5. The New Database Service Request page appears. Specify the following:
  - **Request Name:** Enter the name of the request.
  - **Zone:** Select a PaaS Zone on which the database is to be deployed.
  - **Deployment Input:** Specify the user name and password for the database.

- Schedule Request: Specify the schedule for the request.
6. Click **Submit** to deploy the database to the selected zone.

## 13.3 Viewing the Database Service Home Page

Database services (services) are logical abstractions for managing workloads in Oracle Database. Services divide workloads into mutually disjointed groupings. Each service represents a workload with common attributes, service-level thresholds, and priorities. The Database Service Home page provides a detailed view of the database service.

**Figure 13–2 Database Service Home**



It contains the following regions:

- **Summary:** Displays the status of the database service, database version, and the connection details.
- **Export:** You can do the following:
  - Export: Take export dumps of all the schemas that are part of the service.
  - Import: Restore schemas from the selected export dump.
- **Connection Details:** Shows the database connection details including the Connect String and the Master Account Name.
- **Schemas:** Displays the name of the schema and tablespace associated with the database service.
- **Performance:** This region displays a graph showing the Average Number of Active Sessions over the last 5 minutes.
- **Resource Usage:** This region displays the workloads associated with the schema. This includes the CPU, Memory, and Storage allocated to the database service.

## 13.4 Viewing the Database Instance Home Page

This page contains the following regions:

- **Summary:** The region of the Database Instance Home page provides a quick overview of the database status and provides basic information about the database. It shows the current status of the database and includes details such as the time stamp of instance start time, database version, date of the last backup, and so on. It also displays the connect string used to connect to the database.
- **Backup and Restore:** If the fast recovery area is configured for the database, you can schedule a daily backup for this database. Click **Schedule Backup** to perform this task. Once backups are available, you can click **Restore** to restore the database to a specific point in time. If the fast recovery area is not configured, your Database Administrator can configure the fast recovery area or publish a new service template that allows you to request databases with fast recovery area configured at the time of request fulfillment.

---

**Note:** If the database instance has been created using Snap Clone, the Snapshot and Restore section is displayed here.

---

- **Jobs Running:** Shows the number of currently running jobs. If a value other than 0 appears, you can click the number to go to the Job Activity page where you can view information about all scheduled, currently running, and past jobs.
- **Performance:** Shows the performance monitoring information for the database. Click on the **Active Class** tab to view a bar chart showing the amount of time spent either waiting for a resource or consuming CPU.
- **Resources:** Shows graphs indicating the resource consumption and include Host CPU, Active Sessions, Memory, and Data Storage.
- **SQL Monitoring:** Shows the SQL Response Time chart that shows the current response of the tracked set of SQL versus the reference collection response.

#### Starting and Stopping the Database

Depending on the state of the database, you can use the Shutdown button to shut down the database if the database is open, the Startup button if the database is shut down, or the Shutdown button if the state of the database is unknown or in a Status Pending state.

## 13.5 Viewing the Cluster Database Home Page

Use the Cluster Database Home Page to:

- Determine the current status of the Oracle Real Application Clusters database by viewing a series of metrics.
- Start or stop the cluster database.
- Schedule daily backups and perform a point-in-time-restore of the database.

---

**Note:** The Home Page presents a high level view of the cluster database. Click the arrow key on the left margin to expand the master panel and view the Target Navigation pane. You can select a Database Instance from the left panel to view details for that instance. You can then perform operations such as Startup and Shutdown for the Database Instance.

---

#### Home Page Sections

The Cluster Database Home page displays information for the following sections:

- Summary
- Performance
- Configuration Changes
- Resources
- Jobs Running
- SQL Monitor
- Issues

### **Summary**

The Summary section displays the status and diagnostics information about the cluster database.

#### **Status**

This section provides a quick view of the status of the cluster database and provides basic information about the cluster database. The status can be Up, Down, Under Blackout, Unmonitored, or Unknown, and the status is mirrored in the arrow icon that appears to the left.

The number of database instances in the cluster, the Up Time, Version, Load, Total Sessions, Available Space and other details are displayed. Following are the fields displayed in this section:

- Instances: Displays the number of cluster database instances that are up.
- Up Time: Displays the total elapsed time since the cluster database was started.
- Availability: Displays the percentage of time that the cluster database was available.
- Version: Displays this version of the Oracle Database.
- Load: Displays the number of average active sessions since this cluster database was started.
- Total Sessions: Displays the total cluster database sessions running.
- Latest Backup: Displays when the last backup of the database was performed.
- Available Space: Displays space available on the cluster database.
- Total SGA: Displays total system global area.
- Problem Services: Displays the services that have problems

**Note:** QoS Status on Cluster Database Home page shows if the database is enabled for QoS Management. The possible states are:

- Enabled: Indicates that this database is enabled for QoS monitoring.
- Disabled: Indicates that this database is disabled for QoS monitoring.
- Active: Indicates that this database is enabled and monitored by QoS.

The QoS Status shown on the Cluster Database Home page is different from status shown on Quality of Service Dashboard. QoS Dashboard shows overall status of QoS itself in the Cluster.

### **Jobs Running**



This table displays a report of the job executions that shows the scheduled, running, suspended, and problem (stopped/failed) executions for all Enterprise Manager jobs on the cluster database, ordered by submission of the job to the cluster database or to any member instance. If a value other than 0 appears in a field, you can click the number to go to the Job Activity page where you can view information about all scheduled, currently running, and past jobs.

**Note:** The four job status categories displayed group together job executions of similar status. For a complete list of job status values, refer to About Job Status. Use the Jobs tab to view all jobs or search for jobs of a certain status.

### Performance

The Performance section shows active sessions during a one-hour window for activity classes and top services.

#### Activity Class

The Active Sessions chart shows potential problems inside the database. Categories, called wait classes, show how much of the database is waiting for a resource, such as CPU or disk I/O. The chart displays the load on the instance and identifies bottlenecks in performance.

To quickly identify problem areas, the chart displays a larger block of color to indicate more severe problems. Click the largest wait class on the highest peak, or alternatively click the corresponding wait class (indicated in yellow highlighting). Either action takes you to the Active Sessions Waiting page for the wait class selected, which shows top SQL, sessions, files, and objects for the wait class and also shows the associated wait events.

#### Services

The Active Sessions chart shows the top services waiting for the corresponding wait event during the time period shown. Only active services are shown. Click on one of the service legends to the right of the chart to go to the Service Activity page, where you can view real-time data showing the sessions load for all wait classes associated with the service.

### Resources

The Resources section displays a bar chart showing relative CPU utilization of the Oracle host. This instantaneous value is refreshed every minute from the host by the Agent. The 100% represents the total CPU that the host system provides. The Instances tab is a break down of instance-wise measure of the resources, and the Database tab presents the overall measure.

#### Host CPU

This section displays a bar chart showing the relative CPU utilization across all hosts in the cluster. This instantaneous value is refreshed every minute by the Agent. The 100% represents the total CPU across all hosts in the cluster. Two values appear in the bar chart. The bottom, darker color corresponds to the Database legend and represents how much of the CPU is consumed by all the instances of this database. Click the Database link to go to the Top Activity page to display all wait classes and related statistics. The upper, lighter color corresponds to the Other legend and represents all other processes. Click the Other link to go to the Host Performance page for a quick glimpse of the utilization statistics (CPU, Memory, Disk I/O, and Program Resource Utilization) for this host.

The Load number is unrelated to the Host CPU chart. The Load number relates to the sum of the current CPU load for all cluster hosts. Click the Load link to go to the Host

Performance page for a quick glimpse of the utilization statistics (CPU, Memory, Disk I/O, and Program Resource Utilization) for this host.

**Note:** The Host CPU section does not appear for pre-10g databases.

### Active Sessions

The bar chart shows the amount of time all instances consumed using I/O and CPU, and the amount of time it consumed in bottlenecks. The number shown beside the bar chart is a literal number representing the number of active sessions, rather than the total number of sessions. The chart shows the latest value instead of a historical value. The three session categories are always CPU, User I/O, and Wait.

The **Wait** category represents the value for all wait classes combined, excluding User I/O. All of the links go to the Cluster Database Performance page.

**User I/O** displays the average active sessions spent on waiting for user I/O. User I/O means that the workload originating from the user causes the database to read data from the disk or write data to the disk. Click the User I/O link to go to the Performance page to view potential problems inside and outside of the current database.

**CPU** displays the average active sessions spent on using CPU. Click the CPU link to go to the Top Activity page to display all wait classes and related statistics.

The **Maximum CPU** number is the total CPU count across all the cluster database hosts.

**Note:** The Active Sessions section does not appear for pre-10g databases.

### Memory

The chart shows the memory used by the database in GB.

**Shared Pool** displays various constructs that can be shared among users. For example:

- SQL statements that users implement are cached so that they can be reused if the same statement is used again.
- Information from the data dictionary is cached for quick access and reusability.
- Stored procedures, which are executable code that is stored in the database, can be cached for faster access.

**Buffer Cache** caches blocks of data retrieved from the database. If a piece of data can be found in the buffer cache, then there is no need to read it from disk. There are subcaches for multiple block sizes, if used.

**Large Pool** displays optional area used for buffering large I/O requests for various server processes.

**Java Pool** used for all session-specific Java code and data within the Java Virtual Machine (JVM).

**Other SGA** displays shared memory area that contains data and control information for the instance. Multiple users can share data within this memory area (controlled by Oracle), and information stored in the SGA can avoid repeated access from physical disk, a time consuming operation.

### Data Storage

The chart shows data storage details for tablespaces UNDOTBS, SYSAUX, and SYSTEM.

### SQL Monitor - Last Hour

This section lists the SQL monitoring details for the last hour. The real-time SQL monitoring feature of Oracle Database enables you monitor the performance of SQL statements while they are executing. By default, SQL monitoring automatically starts when a SQL statement runs parallel, or when it has consumed at least 5 seconds of CPU or I/O time in a single execution. It lists the status, duration, SQL ID, Session ID, Parallel, and Database Time.

**Instances**

The Instances table lists the instances for the cluster database, their availability, incidents, compliance score, ASM instance information, and ADDM findings. Click an instance name on the left hand panel to go to the home page for that instance. Click the links in the table to get more information about that particular incident.



## Monitoring a DBaaS Cloud

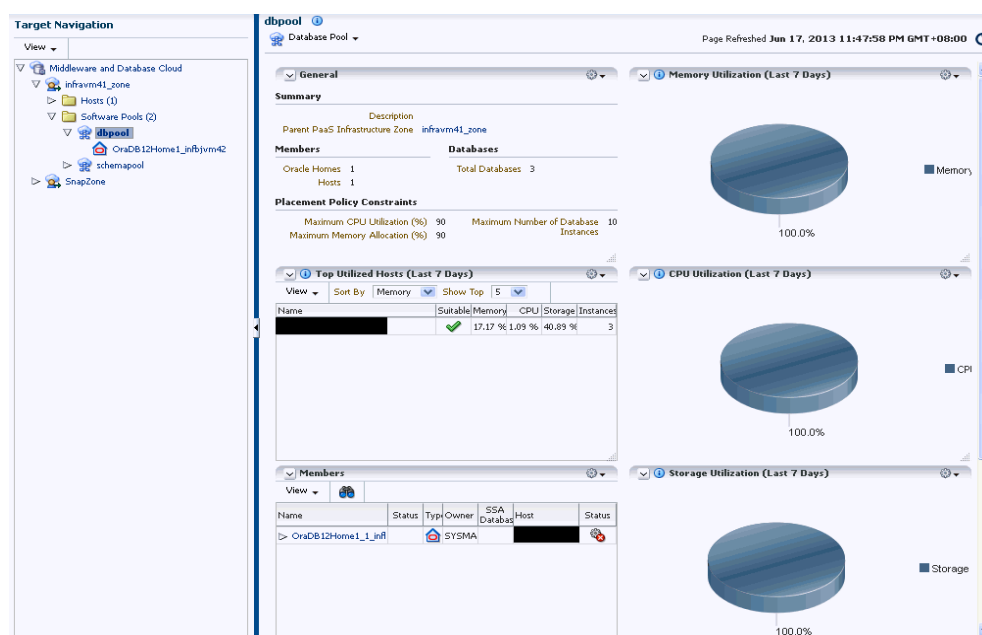
This chapter describes support for monitoring a DBaaS Cloud. It contains the following sections:

- [Viewing the Database Pool Home Page](#)
- [Viewing the Schema Pool Home Page](#)

### 14.1 Viewing the Database Pool Home Page

A database pool contains a set of resources that can be used to provision a database instance within a PaaS Infrastructure Zone.

**Figure 14–1 Database Pool Home**



This page contains the following regions:

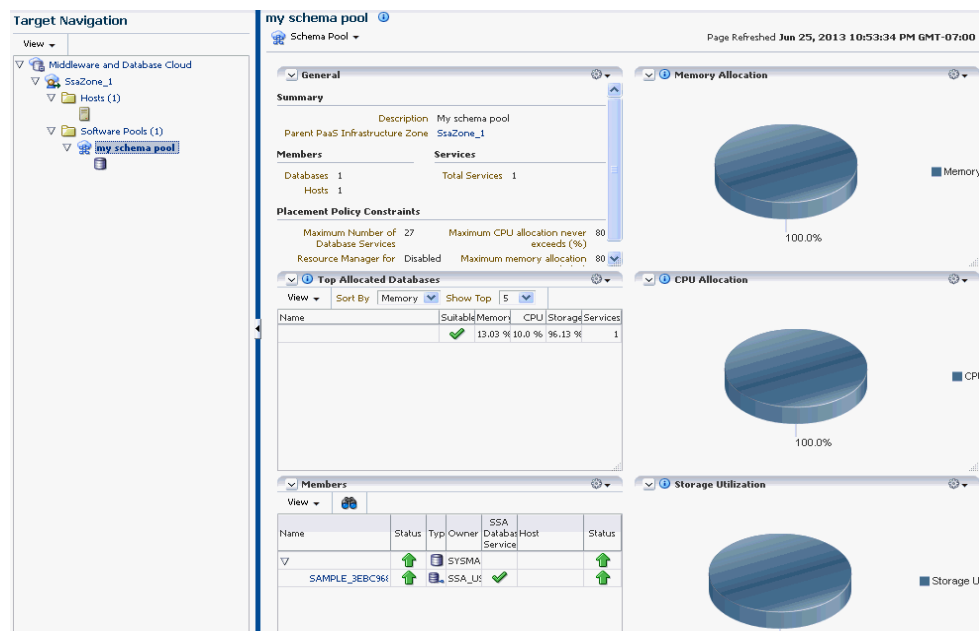
- **General:** This region displays the following:
  - **Summary:** Displays a description of the database pool and PaaS Infrastructure Zone to which it belongs. I

- **Members and Databases:** Shows the list of Members (Oracle Homes and Hosts) and Databases in the pool. The Databases region lists the total number of databases and the number of single instance or RAC databases in the pool.
- **Placement Policy Constraints:** The placement policy constraints that specify maximum ceilings for each host member of the pool are displayed.
- **Top Utilized Hosts:** This region shows the list of top utilized hosts over the last 7 days. Click on the link to drill down to the Host Home page.
- **Members:** Lists all the members in the database pool. Click on the links in the Members table to drill down to the Database or Host Home page.
- **Overview of Incidents and Problems:** This region lists any incidents or problems that have occurred.
- **Memory Utilization:** The chart shows the memory utilization by all the targets in the database pool over the last 7 days.
- **CPU Utilization:** The chart shows the CPU utilization by all the targets in the database pool over the last 7 days.
- **Storage Utilization:** The chart shows the storage utilization by all the targets in the database pool over the last 7 days.
- **Service Template and Instances:** Shows the service templates and service instances associated with the database pool.

## 14.2 Viewing the Schema Pool Home Page

You can define one or more schemas that can be deployed on a database.

**Figure 14–2** *Schema Pool Home*



- **General:** This region displays the following:
  - **Summary:** Displays a description of the schema pool and PaaS Infrastructure Zone to which it belongs. I

- **Members and Services:** Shows the list of Members (Databases and Hosts) and Services in the pool.
- **Placement Policy Constraints:** The placement policy constraints that specify maximum ceilings for each database member in the pool are displayed.
- **Top Allocated Databases:** This region shows the list of top utilized databases in the schema pool. Click on the link to drill down to the Database Instance page.
- **Members:** Lists all the members in the schema pool. Click on the links in the Members table to drill down to the Database Instance Home page.
- **Overview of Incidents and Problems:** This region lists any incidents or problems that have occurred.
- **Memory Utilization:** The chart shows the memory utilization by all the targets in the schema pool over the last 7 days.
- **CPU Utilization:** The chart shows the CPU utilization by all the targets in the schema pool over the last 7 days.
- **Storage Utilization:** The chart shows the storage utilization by all the targets in the schema pool over the last 7 days.
- **Service Template and Instances:** Shows the service templates and service instances associated with the schema pool.





## Setting Up an MWaaS Cloud

This chapter covers the initial configuration required to set up and enable a Middleware as a Service Cloud. It contains the following sections:

- [Getting Started with MWaaS for Physical Hosts](#)
- [Getting Started with MWaaS for Virtual Hosts](#)

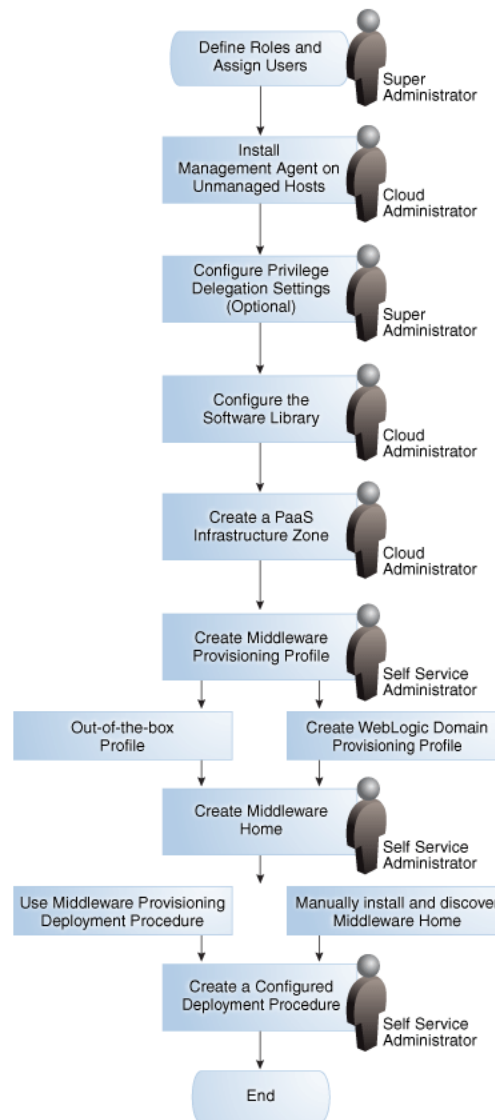
### 15.1 Getting Started with MWaaS for Physical Hosts

This section helps you get started by providing an overview of the steps involved in setting up the Middleware as a Service (MWaaS) to provision a service instance on one or more physical hosts. Before you set up MWaaS, you must download and deploy the required plug-ins. For more details, see [Section 3.3, "Deploying the Required Plug-ins"](#).

**Table 15–1** *Getting Started with MWaaS for Physical Hosts*

| Step | Task                                                                                                                                                                  | Role                       |
|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------|
| 1    | Define roles for administrators and self service users. See <a href="#">Section 3.4, "Defining Roles and Assigning Users"</a> .                                       | Super Administrator        |
| 2    | Install the Management Agent on unmanaged hosts so that they can be monitored by Enterprise Manager. <a href="#">Section 9.2, "Adding Hosts"</a> .                    | Cloud Administrator        |
| 3    | Optionally, you may need to configure Privilege Delegation Settings on your managed hosts. <a href="#">Section 3.6, "Configuring Privilege Delegation Settings"</a> . | Super Administrator        |
| 4    | Configure the storage location for the Software Library. See <a href="#">Section 3.1, "Setting Up the Software Library"</a>                                           | Cloud Administrator        |
| 5    | Load Middleware Provisioning Profiles into the Software Library. See <a href="#">Section 15.1.1, "Creating Provisioning Profiles"</a> .                               | Self Service Administrator |
| 6    | Create a Middleware Home. See <a href="#">Section 15.1.2, "Creating a Middleware Home"</a> .                                                                          | Self Service Administrator |
| 7    | Create a Configured Deployment Procedure. See <a href="#">Section 15.1.3, "Configuring and Saving the Deployment Procedure"</a> .                                     | Self Service Administrator |

[Figure 15–1](#) shows the sequence of steps required to enable MWaaS (physical hosts).

**Figure 15–1 Getting Started with MWaaS (Physical Hosts)**

### 15.1.1 Creating Provisioning Profiles

Provisioning profiles can be used to standardize deployments and help reduce errors when deployment procedures are configured. These provisioning profiles are used by operators for mass deployment. You can:

- Create a Middleware Provisioning Profile from a domain by using the Software Library. See [Section 15.1.1.1, "Creating a WebLogic Domain Provisioning Profile"](#) for details.
- Use out-of-the-box Middleware Provisioning Profiles. These profiles have been certified with WebLogic Server 10.3.5.0, and 10.3.6.0. See [Section 15.1.1.2, "Using the Out-of-the-box Provisioning Profiles"](#) for details.

#### 15.1.1.1 Creating a WebLogic Domain Provisioning Profile

A WebLogic Domain Provisioning Profile consists of the Middleware Home, binaries, and the domain configuration. You can create a profile, save it in the Software Library,

and then use the saved profile as the source for creating new WebLogic domains. This will ensure that future WebLogic installations follow a standard, consistent configuration.

### Prerequisites

- The Management Agent must be running on the Administration Server.
- You must have the host credentials for the Administration Server running on the source machine.
- The WebLogic Domain for which the profile is being created must be a monitored target in Enterprise Manager.
- The disk space required to create a profile is calculated as follows:  

$$\text{Disk Space} = \text{Middleware Home Size} + \text{WebLogic Domain Size} + \text{Space for Temporary Scripts}$$
- The JRF binaries are installed in the Middleware Home and configured in the domain.
- The WebLogic Domain that is used to create a profile must meet the following requirements:
  - The WebLogic Domain must be configured on a single host, with a single cluster and one Managed Server.
  - The machine associated with the WebLogic Server must not be of Unix type.
  - The Administration Server must not be associated with a machine.
- The maximum heap value of all the servers must have an appropriate -Xmx setting in multiples of 256. For example, Xms1024m -Xmx1536m -XX:PermSize=512m -XX:MaxPermSize=1536m and so on. You must define the -Xmx server setting to ensure that the MWaaS memory metric and quota checking feature works properly. The value of the memory metric for a MWaaS service is the sum of the maximum heap available for all servers in the service.

To create a profile, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Software Library**.
2. From the **Actions** menu, click **Create Folder** to create a folder in which the profile you are creating is to be stored. We recommend that you do not use an out-of-box folder to store the profile.
3. Select the folder that you have created, and from the **Actions** menu, select **Create Entity**, and then select **Component**.
4. The Create Component popup window appears. From the **Select Subtype** drop-down list, select the **WebLogic Domain Provisioning Profile** component and click **Continue**.
5. In the Create WebLogic Domain Provisioning Profile: Describe page, the Parent Directory under which the profile will be created and the Subtype are displayed. Enter a name (only alphanumeric characters and underscores are allowed) and description for the profile and enter the values for the Product Version, Product, and Vendor attributes. For example, enter the Product Version as 10.3.6, Product as WebLogic Domain, and Vendor as Oracle.
6. Click **Next**. The Create WebLogic Domain Provisioning Profile: Configure page appears.

**Figure 15–2 Create WebLogic Domain Provisioning Profile: Configure**

**Software Library**

Describe **Configure** Review

Create WebLogic Domain Provisioning Profile : Configure Back Step 2 of 3 Next Save Cancel

Parent Directory Components  
Subtype WebLogic Domain Provisioning Profile

In order to create a WebLogic Domain Provisioning Profile, you must first select an existing WebLogic Domain. By default, the domain selected and its Middleware Home would be archived and stored in the software library for future cloning operations. If you want only the domain configuration to be included in the profile, then modify the default settings by deselecting the checkbox below.

☐ Include the binaries from the Middleware Home in the profile to be created.

WebLogic Domain GCDomain

Host [REDACTED]

\* Working Directory /tmp/ps2mwhProvSrc

**Source Information**

| Component            | Location                                             |
|----------------------|------------------------------------------------------|
| GCDomain             | /scratch/aimc/gc_inst/user_projects/domains/GCDomain |
| Middleware Home      | /scratch/aimc/ps2mwh                                 |
| WebLogic Server 10.3 | /scratch/aimc/ps2mwh/wls_server_10.3                 |
| Oracle Home(s)       |                                                      |
| oracle_common        | /scratch/aimc/ps2mwh/oracle_common                   |

**Host Credentials**

☒ Preferred Credentials ☐ Named Credentials

Credential Set Normal Host Credentials

7. Click the **Search** icon next to the WebLogic Domain field and select a WebLogic Domain from which the profile is to be created. The profile will include both the software in the Middleware Home as well as the configuration in the WebLogic Domain home.
8. Deselect the **Include the binaries from the Middleware Home in the profile to be created** checkbox. Since the Middleware Home is already installed, deselecting this checkbox ensures that only the configuration is included.
9. In the Working Directory field, specify the directory on the host machine on which the files required for creating the profile are temporarily stored. If this directory is not present, it will be created. When the profile has been created, the contents of this directory will be deleted.
10. In the Host Credentials section, enter the host credentials of the machine on which the Administration Server of the source WebLogic Domain is installed. Select one of the following options:
  - **Preferred Credentials:** The preferred credentials stored in the Management Repository are used. The Preferred Credentials option will be available only if it has already been defined in Enterprise Manager. For more information on setting up Preferred Credentials, see *Managing Preferred Credentials*.
  - **Named Credentials:** The credentials stored in the Management Repository is used. To use the Named Credentials stored in the Management Repository, you must have already registered each of the preferred credential types with a unique name. Select the desired Named Credentials from the list available.
11. Click **Next**. In the Review page, you can review the information and click **Save and Upload**.
12. The Job Name is displayed at the top of the page. Navigate to the Job Activity page and check the job status. Once it has been completed, navigate to the Software Library page and confirm if all the components of the profile (WebLogic Domain, MWHome, and Profile) have a **Ready** Status. You can also verify if the profile has been successfully created by navigating to the Middleware Provisioning page. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Middleware Provisioning** and check if the profile is displayed there.

You can now use this profile as the source for future WebLogic Domain installations.

### 15.1.1.2 Using the Out-of-the-box Provisioning Profiles

When you install Enterprise Manager, a self update archive with three out-of-the-box provisioning profiles with different heap size settings (Small: 1.5 GB, Medium: 3 GB, and Large: 6 GB) is available. You can download these profiles from the Self Update console.

To use the out-of-the-box provisioning profiles, follow these steps:

1. Log in to Enterprise Manager as a user with the `EM_CLOUD_ADMINISTRATOR` role.
2. From the **Setup** menu, select **Extensibility** and then select **Self Update**. The Self Update Console with the list of entity types that can be downloaded appears.

**Figure 15–3 Self Update Console**

**Self Update** Page Refreshed Jan 25, 2013 1:16:55 PM PST

Oracle periodically provides new functionality and updates for existing features in Enterprise Manager. The Self Update home allows administrators to receive notifications and view, download, and apply such updates. While these updates are retrieved automatically, a manual check can be made at any time.

**Status**

Connection Mode: **Error** Last Download Time: Jan 24, 2013 4:25:29 PM PST Last Apply Time: Jan 22, 2013 12:07:04 PM PST  
 Most Recent Refresh Time: Jan 24, 2013 3:01:04 PM PST Last Download Type: Middleware Profiles and Gold Images Last Apply Type: Plug-in  
 Last Successful Refresh Time: N/A

**Actions** Open Check Updates Show Applicable Updates

| Type                                | Available Updates | Downloaded Updates | Applied Updates | Description                                                                                                                                                                                                 |
|-------------------------------------|-------------------|--------------------|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Extensibility Development Kit       | 0                 | 0                  | 1               | Extensibility Development Kit (EDK) is a standalone tool designed to help development of Enterprise Manager extensions (for example: Plug-ins)                                                              |
| Management Connector                | 0                 | 0                  | 1               | Management Connectors are components that integrate different enterprise frameworks into the Enterprise Manager Console                                                                                     |
| Middleware Profiles and Gold Images | 1                 | 1                  | 0               | A collection of Software Components used for provisioning of Oracle Application Server homes.                                                                                                               |
| Plug-in                             | 0                 | 0                  | 14              | Plug-in extends Enterprise Manager to manage newer target type as well as to bring vertical functionality                                                                                                   |
| Provisioning Bundle                 | 0                 | 0                  | 1               | Provisioning bundle is a collection of deployment procedures, software library entities, and other related artifacts that cater to the provisioning and patching of various Oracle and non-Oracle Products. |

**Past Activities: Middleware Profiles and Gold Images**

| Actions   | Status    | Product Version | Target                           | Version    | Administrator | Start Time                  | Elapsed Time(Sec) |
|-----------|-----------|-----------------|----------------------------------|------------|---------------|-----------------------------|-------------------|
| Download  | Succeeded | 10.3.6.0        | Oracle Supplied Fusion Middle... | 12.1.0.3.0 | SYSMAN        | Jan 24, 2013 4:25:29 PM PST | 0.32              |
| Available | Succeeded | 10.3.6.0        | Oracle Supplied Fusion Middle... | 12.1.0.3.0 | SYSMAN        | Jan 24, 2013 4:25:29 PM PST | 0.58              |

3. Select the **Middleware Profiles and Gold Images** folder and click **Open**. The **Middleware Profiles and Gold Images Updates** page appears. You will see "1" as an available update.
4. Select the update and click **Download**. The **Schedule Download** dialog appears.
5. Specify if the update is to be downloaded immediately or at a later date.
6. Click **Select**. An Enterprise Manager job is created to download the update to the Software Library. Enterprise Manager starts downloading the archive from the Oracle Enterprise Manager store. Wait for the download to complete. (When in offline mode the system starts reading from the specified location.) When the download is complete, the status indicates that the update has been downloaded.



**Note:**

- The out-of-the-box profiles have been certified to work with WebLogic Server 10.3.6.0.
- As the sample provisioning profiles do not contain the Oracle Middleware Home Gold Image component, you must ensure that the Middleware Homes (WebLogic 10.3.6.0) and JRF binaries of the same version have been created on the selected destination host(s) when they are provisioned.
- Scale-up operations are supported only with non-SSL profiles.

## 15.1.2 Creating a Middleware Home

A Middleware Home consists of the WebLogic Server Home (such as WebCenter Home, WebTier Home, and so on) and optionally, one or more Oracle Homes. The Middleware Home must be created on all the hosts that are part of the PaaS Infrastructure Zone. See [Section 9.3, "Creating a PaaS Infrastructure Zone"](#) for details. If the Middleware Home has not been created on a host, that host cannot participate in the Middleware Pool.

---

**Note:** You can also manually install the Middleware Home and use the **Discover Promote Oracle Home Target** job to discover it in Enterprise Manager.

---

To create a Middleware Home, you must run the **Provision from Oracle Middleware Home Gold Image** deployment procedure.

### Prerequisites

Before running this deployment procedure, you must meet the following prerequisites:

- The user must have Write permissions on:
  - The Working Directory on all destination hosts.
  - The Middleware Home on all destination hosts.
- An Oracle Middleware Home Gold Image must be present in the Software Library. For details on creating this gold image, see [Section 15.1.2.2, "Creating an Oracle Middleware Home Gold Image"](#).

### 15.1.2.1 Provision from Oracle Middleware Home Gold Image Deployment Procedure

**Note:**

- If you have manually created and discovered a Middleware Home, this step can be ignored.
- You cannot have more than one Middleware Home of the same WebLogic Server version on a single host.

You can clone a Middleware Home from a gold image present in the Software Library. This gold image must have been created earlier by pointing to an existing Middleware Home. To clone a Middleware Home from a gold image, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Middleware Provisioning**.
2. From the Middleware Provisioning Deployment Procedures section, select the Provision Middleware procedure from the list and click **Launch**.
3. In the Middleware Provisioning: Source page, select the **Provision from Oracle Middleware Home Gold Image** option. Click the **Search** icon next to the Home Gold Image from Software Library field and select a gold image from the list.
4. In the Source Information section, the selected component and its location is displayed. Click **Next**.
5. In the Middleware Provisioning: Destinations page, specify the destination hosts on which the Middleware Home is to be cloned.

**Figure 15–6 Middleware Provisioning (Gold Image): Destination**

6. Click **Add Hosts** and select a host from the list.
7. In the Select Destination Locations section, specify the location of the Middleware Home on the Host machine. The Create a New Middleware Home option is selected by default.
8. In the Middleware Home Directory field, enter the full path to the directory in which the Middleware Home is to be created.
9. In the Working Directory field, specify the directory on the destination Host on which the cloning related files are temporarily stored. This directory must have sufficient space to store the files. If this directory is not present, it will be created. When the cloning operation has been completed, the directory and its contents will be deleted.
10. In the JDK Home Location field, enter the absolute path to the JDK directory to be used on the destination Host. This field can be edited only if the JDK Home in the source domain is in an external location. If the JDK Home in the source domain is internal and installed in the Middleware Home, this field cannot be edited.

---

**Note:** If there are several destination hosts, the location of the Middleware Home, Working Directory, and JDK Home Location is the same across all the hosts.

---



11. In the Middleware Provisioning: Credentials page, specify the credentials for the destination host. They can be:
  - **Preferred Credentials:** The preferred credentials stored in the Management Repository are used. The Preferred Credentials option will be available only if it has already been defined in Cloud Control.
  - **Named Credentials:** The credentials of a named profile stored in the Management Repository is used. To use the Named Credentials stored in the Management Repository, you must have already registered each of the preferred credential types with a unique name. Select the desired Named credential from the list available in Credential Name. If you have created all necessary named credentials, you can use them now. If they have not been created, you can create them using this deployment procedure.
12. In the Schedule page, specify a Deployment Instance name. If you want to run the procedure immediately, then retain the default selection, that is, One Time (Immediately). If you want to run the procedure later, then select One Time (Later) and provide time zone, start date, and start time details. You can set the notification preferences according to deployment procedure status. If you want to run only prerequisites, you can select **Pause the procedure after the necessary prerequisite checks have been completed** to pause the procedure execution after all prerequisite checks are performed. Click **Next**.
13. On the Review page, review the details you have provided for the deployment procedure. If you are satisfied with the details, then click **Submit** to run the deployment procedure according to the schedule set. If you want to modify the details, click the **Edit** link in the section to be modified or click **Back** repeatedly to reach the page where you want to make the changes.
14. In the Procedure Activity page, view the status of the execution of the job and steps in the deployment procedure. Click the **Status** link for each step to view the details of the execution of each step. You can click **Debug** to set the logging level to Debug and click **Stop** to stop the procedure execution.

### 15.1.2.2 Creating an Oracle Middleware Home Gold Image

You can create an Oracle Middleware Home Gold Image and save it in the Software Library. You can then use this gold image as the source for future Middleware Home installations.

#### Prerequisites

- The Management Agent must be running on the Administration Server.
  - You must have the host credentials for the Administration Server running on the source machine.
  - The disk space required to create a gold image is calculated as follows:  

$$\text{Disk Space} = \text{Middleware Home Size} + \text{Space for Temporary Scripts}$$
1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Software Library**.
  2. From the **Actions** menu, click **Create Folder** to create a folder in which the gold image you are creating is to be stored. We recommend that you do not use an out-of-box folder to store the gold image.
  3. Select the folder you have created and from the **Actions** menu, select **Create Entity**, then select **Component**.

4. A Create Component popup window appears. From the Select Subtype drop-down list, select the Oracle Middleware Home Gold Image component and click **Continue**.
5. In the Describe page, the Parent Directory under which the gold image will be created and the Subtype are displayed. Enter a name and description for the profile and enter the values for the Product Version, Product, and Vendor attributes. For example, enter the Product Version as 10.3.6, Product as Middleware Home, and Vendor as Oracle.
6. Click **Next**. Create Oracle Middleware Home Gold Image: Configure page appears. Click the **Search** icon next to the Middleware Home field and select an existing Middleware Home from which the gold image is to be created.

**Figure 15–7 Create Oracle Middleware Home Gold Image: Configure**

The screenshot shows the 'Create Oracle Middleware Home Gold Image: Configure' page. At the top, there's a 'Software Library' header with 'Describe', 'Configure', and 'Review' tabs. Below the tabs, the title 'Create Oracle Middleware Home Gold Image: Configure' is displayed. On the right, there are buttons for 'Back', 'Step 2 of 3', 'Next', 'Save', and 'Cancel'. The main content area includes a 'Parent Directory' field with 'Components' as the subtype. Below this, a message states: 'In order to create a Oracle Middleware Home Gold Image, you must first select an existing Middleware Home. The home selected would be archived and stored in the software library for future cloning operations.' There are input fields for 'Middleware Home' (with a search icon), 'Host' (with a redacted value), and 'Working Directory' (set to 'C:\Temp\fmwProvSrc'). A 'Source Information' section contains a table with two columns: 'Component' and 'Location'. The table lists 'Middleware Home' at 'C:\Midthome' and 'WebLogic Server 10.3' at 'C:\Midthome\wlserver'. Below the table is the 'Host Credentials' section, which has radio buttons for 'Preferred Credentials' (selected) and 'Named Credentials'. A 'Credential Set' dropdown is set to 'Normal Host Credentials'.

| Component            | Location             |
|----------------------|----------------------|
| Middleware Home      | C:\Midthome          |
| WebLogic Server 10.3 | C:\Midthome\wlserver |

7. In the Working Directory field, specify the directory on the destination Host on which the files required for creating the gold image are temporarily stored. If this directory is not present, it will be created. When the gold image has been created, the contents of this directory will be deleted.
8. In the Host Credentials section, enter the host credentials of the machine on which the Middleware Home is located. Select one of the following options:
  - **Preferred Credentials:** The preferred credentials stored in the Management Repository are used. The Preferred Credentials option will be available only if it has already been defined in Enterprise Manager.
  - **Named Credentials:** The credentials stored in the Management Repository is used. To use the Named Credentials stored in the Management Repository, you must have already registered each of the preferred credential types with a unique name. Select the desired Named credential from the list available in Credential Name.
9. Click **Next**. In the Review page, you can review the information and click **Save and Upload**.
10. The Job Name is displayed at the top of the page. Navigate to the Job Activity page and check the job status. Once it has been completed, navigate to the Software Library page and confirm if all the components of the gold image have a Ready Status. You can now use this gold image for future cloning operations.

### 15.1.3 Configuring and Saving the Deployment Procedure

You can configure and save a Middleware Provisioning deployment procedure. This Configured Deployment Procedure can then be used by the EM\_SSA\_ADMINISTRATOR while creating a Middleware Service Template.

To customize the deployment procedure, follow these steps:

1. Log into Enterprise Manager as a user with EM\_SSA\_ADMINISTRATOR role.
2. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Middleware Provisioning**. The Middleware Provisioning page appears.
3. The profiles that have been uploaded are listed in the Profile section. The profile you select must contain only one cluster with one WebLogic Server.
4. Select the Provision Middleware Deployment Procedure from the Deployment Procedures region and click **Launch**.
5. Select **Provision** from WebLogic Domain Provisioning Profile option. Click the **Search** icon and select one of the profiles as the image to use for the oracle home.
6. Enter a name for the customized deployment procedure and click **Save**.

The MWaaS feature will not work if any variables in the deployment procedure are locked.

7. After the deployment procedure has been saved, click **Cancel** to exit the current deployment procedure. You will see the new customized deployment procedure listed in the Deployment Procedures table in the Middleware Provisioning page. The saved deployment procedure can now be used while creating a service template.

---

**Note:** The Configured Deployment Procedure you have created must contain only one cluster with one WebLogic Server.

---

## 15.2 Getting Started with MWaaS for Virtual Hosts

This section helps you get started by providing an overview of the steps involved in setting up the Middleware as a Service (MWaaS) and provision a service instance on one or more Oracle VM Zones. Before you set up MWaaS, you must download and deploy the required plug-ins. For more details, see [Section 3.3, "Deploying the Required Plug-ins"](#).


**Table 15–2** *Getting Started with MWaaS for Virtual Hosts*

| Step | Task                                                                                                                                                                  | Role                |
|------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------|
| 1    | Define roles for administrators and self service users. See <a href="#">Section 3.4, "Defining Roles and Assigning Users"</a> .                                       | Super Administrator |
| 2    | Optionally, you may need to configure Privilege Delegation Settings on your managed hosts. <a href="#">Section 3.6, "Configuring Privilege Delegation Settings"</a> . | Super Administrator |
| 3    | Configure the storage location for the Software Library. See <a href="#">Section 3.1, "Setting Up the Software Library"</a>                                           | Cloud Administrator |
| 4    | Register Oracle VM Manager. See <a href="#">Section 4.3, "Registering the Oracle VM Manager"</a> .                                                                    | Cloud Administrator |
| 5    | Discover Virtual Servers. See <a href="#">Section 4.5, "Discovering a Virtual Server"</a> .                                                                           | Cloud Administrator |

**Table 15–2 (Cont.) Getting Started with MWaaS for Virtual Hosts**

| Step | Task                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Role                       |
|------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------|
| 6    | Register Storage Servers. See <a href="#">Section 4.9, "Registering Storage Servers"</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Cloud Administrator        |
| 7    | Configure Networks. See <a href="#">Section 4.7, "Setting Up Networks"</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Cloud Administrator        |
| 8    | Create Virtual Server Pool. See <a href="#">Section 4.14, "Creating a Virtual Server Pool"</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Cloud Administrator        |
| 9    | Create Oracle VM Zone. See <a href="#">Section 4.15, "Creating a Zone"</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      | Cloud Administrator        |
| 10   | Create a PaaS Infrastructure Zone. See <a href="#">Section 9.3, "Creating a PaaS Infrastructure Zone"</a> .                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       | Cloud Administrator        |
| 11   | <p>Deploy an Oracle Virtual Assembly (ova) and create a Deployment Plan. See <a href="#">Section 8.8.4, "Provisioning Guest Virtual Machines Using Oracle Virtual Assemblies (OVA)"</a>.</p> <p>For more information on assemblies, see the Oracle Virtual Assembly Builder page<br/> <a href="http://www.oracle.com/technetwork/middleware/ovab/overview/index.html">http://www.oracle.com/technetwork/middleware/ovab/overview/index.html</a> on the Oracle Technology Network.</p> <p>For instructions on building assemblies, see the Oracle Virtual Assembly Builder Installation Guide.</p> | Self Service Administrator |

[Figure 15–8](#) shows the sequence of steps required to enable middleware as a service (virtual hosts).

**Figure 15–8 Getting Started with MWaaS (Virtual Hosts)**

After you followed the steps listed in [Figure 15–8](#), do the following:

- Set up Middleware Cloud Self Service Portal by following the steps listed in [Section 16.1, "Setting Up the Middleware Cloud Self Service Portal"](#)
- Review [Section 16.1.2, "Setting Up the Middleware Cloud Self Service Portal for Virtual Hosts"](#)

- Ensure that the prerequisites listed in [Section 16.8, "Creating a Middleware Virtual Service Template"](#) are met.

---

## Setting Up the MWaaS Self Service Portal

This chapter describes the procedure to set up and enable the Middleware Self Service Portal. It contains the following sections:

- [Setting Up the Middleware Cloud Self Service Portal](#)
- [Creating a Middleware Pool](#)
- [Configuring Request Settings](#)
- [Creating a Load Balancer Configuration](#)
- [Setting Up Quotas](#)
- [Creating a Middleware Service Template](#)
- [Configuring Chargeback](#)

### 16.1 Setting Up the Middleware Cloud Self Service Portal

The `EM_SSA_ADMINISTRATOR` can configure and set up the middleware self service portal by creating middleware pools, quotas, request settings, and service templates. To set up the MWaaS Self Service Portal, follow these steps:

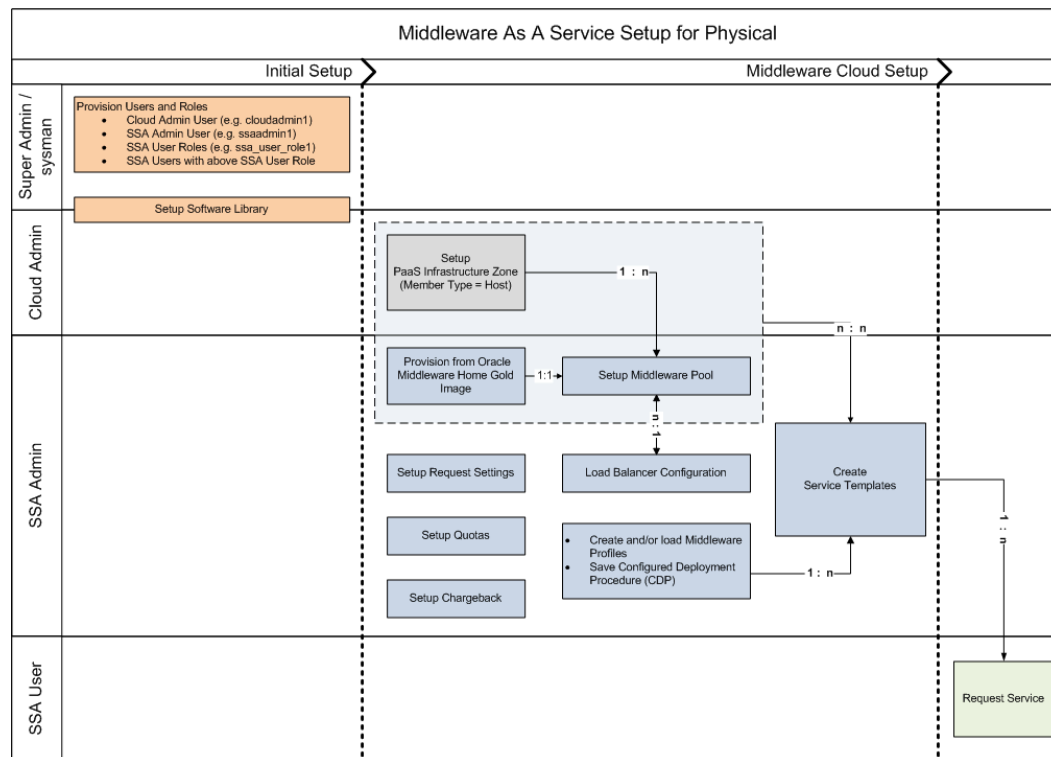
1. Follow the steps in [Section 15.1, "Getting Started with MWaaS for Physical Hosts"](#) to enable MWaaS.
2. Set up one or more PaaS Infrastructure zones. See [Section 9.3, "Creating a PaaS Infrastructure Zone"](#) for details.
3. Create a Middleware Pool. See [Section 16.2, "Creating a Middleware Pool"](#).
4. Optionally, you can configure an external load balancer. See [Section 16.3, "Creating a Load Balancer Configuration"](#).
5. Configure request settings. See [Section 16.4, "Configuring Request Settings"](#).
6. Define quotas for each self service user role. See [Section 16.5, "Setting Up Quotas"](#).
7. Create service templates. See [Section 16.6, "Creating a Middleware Service Template"](#).
8. Optionally, you can configure the Chargeback Service. See [Section 16.9, "Configuring Chargeback"](#).

#### 16.1.1 Setting Up the Middleware Cloud Self Service Portal for Physical Hosts

After you have set up MWaaS, you must configure and set up the MWaaS Cloud Self Service Portal by creating middleware pools, setting up quotas, defining request settings, configuring the load balancer, creating service templates, and configuring

chargeback. [Figure 16–1](#) diagram shows the initial setup and the relationship between the various entities involved in setting up the Cloud Middleware Self Service Portal.

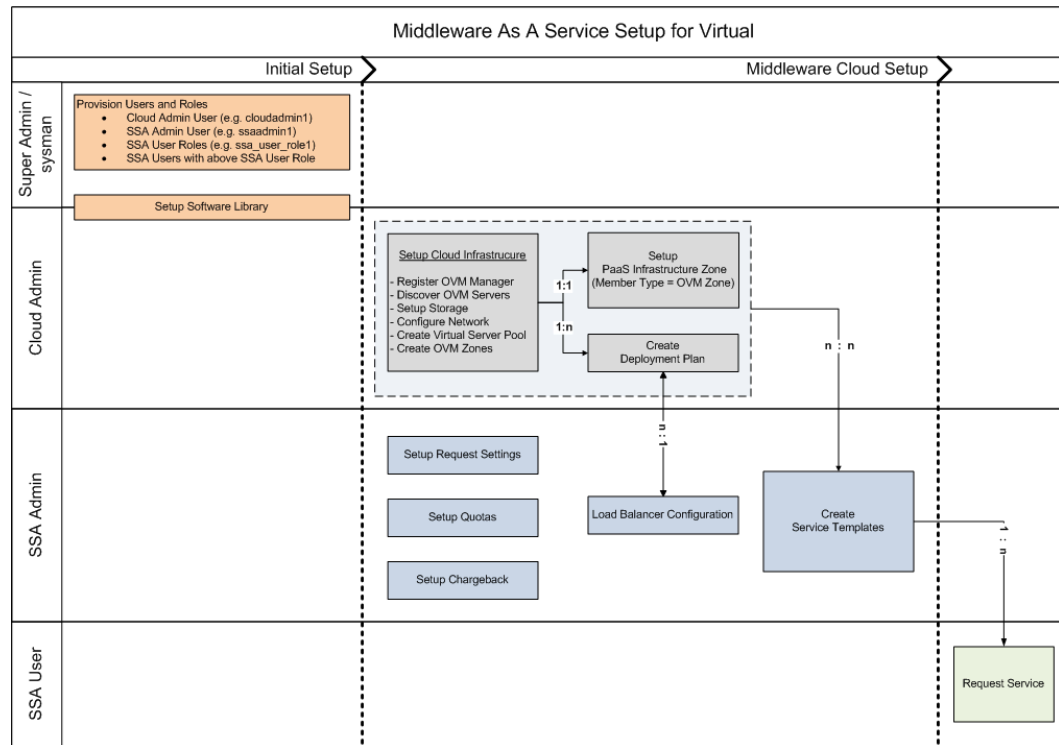
**Figure 16–1 Setting Up MWaaS (Physical Hosts)**



### 16.1.2 Setting Up the Middleware Cloud Self Service Portal for Virtual Hosts

After you have set up MWaaS, you must configure and set up the MWaaS Cloud Self Service Portal by creating middleware pools, setting up quotas, defining request settings, configuring the load balancer, creating service templates, and configuring chargeback. [Figure 16–3](#) diagram shows the initial setup and the relationship between the various entities involved in setting up the Cloud Middleware Self Service Portal.



**Figure 16–2 Setting Up MWaaS (Virtual Hosts)**

## 16.2 Creating a Middleware Pool

A Middleware Pool contains a set of resources that can be used to provision a service instance within a PaaS Infrastructure Zone. To create a Middleware Pool, follow these steps:

1. Log in to Enterprise Manager as a user with the `EM_SSA_ADMINISTRATOR` role.
2. From the **Setup** menu, select **Cloud**, then select **Middleware**. The Middleware Cloud Self Service Portal Setup page appears.
3. Click **Middleware Pools**. The existing middleware pools are listed on the Middleware Pools page.

**Figure 16–3 Middleware Cloud Self Service Portal Setup: Middleware Pools**

| Middleware Cloud Self Service Portal Setup |                                                                                                                                                                                                                                                                                 |                          |         |                                           |           |       |                          |         |             |           |        |         |   |  |
|--------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------|---------|-------------------------------------------|-----------|-------|--------------------------|---------|-------------|-----------|--------|---------|---|--|
|                                            |                                                                                                                                                                                                                                                                                 |                          |         | Page Refreshed May 8, 2013 5:44:14 PM PDT |           |       |                          |         |             |           |        |         |   |  |
| Middleware Pools                           | <b>Middleware Pools</b><br>A middleware pool is a set of Oracle homes that can be used to provision user requested service instances. All Oracle homes within a middleware pool must be of the same version and platform, and must belong to the same PaaS Infrastructure Zone. |                          |         |                                           |           |       |                          |         |             |           |        |         |   |  |
| Load Balancer Configurations               | View <span>▼</span> <span>Create...</span> <span>Edit...</span> <span>Delete</span>                                                                                                                                                                                             |                          |         |                                           |           |       |                          |         |             |           |        |         |   |  |
| Request Settings                           | <table border="1"> <thead> <tr> <th>Pool Name</th><th>Owner</th><th>PaaS Infrastructure Zone</th><th>Targets</th><th>Description</th></tr> </thead> <tbody> <tr> <td>my_mwpool</td><td>SYSMAN</td><td>omszone</td><td>1</td><td></td></tr> </tbody> </table>                    |                          |         |                                           | Pool Name | Owner | PaaS Infrastructure Zone | Targets | Description | my_mwpool | SYSMAN | omszone | 1 |  |
| Pool Name                                  | Owner                                                                                                                                                                                                                                                                           | PaaS Infrastructure Zone | Targets | Description                               |           |       |                          |         |             |           |        |         |   |  |
| my_mwpool                                  | SYSMAN                                                                                                                                                                                                                                                                          | omszone                  | 1       |                                           |           |       |                          |         |             |           |        |         |   |  |
| Quotas                                     |                                                                                                                                                                                                                                                                                 |                          |         |                                           |           |       |                          |         |             |           |        |         |   |  |
| Service Templates                          |                                                                                                                                                                                                                                                                                 |                          |         |                                           |           |       |                          |         |             |           |        |         |   |  |
| Chargeback                                 |                                                                                                                                                                                                                                                                                 |                          |         |                                           |           |       |                          |         |             |           |        |         |   |  |

4. Click **Create**. The Create Middleware Pool: General page appears. Enter a name and description for the Middleware Pool.

5. Specify the Placement Policy Constraints. These constraints are used to set maximum limits on resource utilization for each host. In the Maximum Number of Java Servers (per host) field, enter the maximum number of Java Server instances that can be running on each host in the pool.

**Figure 16–4 Create Middleware Pool: General**

Create Middleware Pool : General

Back Step 1 of 2 Next Cancel

**General**

\* Name my\_mwpool

Description

**Placement Policy Constraints**

Placement policy constraints allow the self service administrator to set maximum ceilings for resource utilization. This provides protection for the members of the software pool in terms of resource consumption. For example, a production software pool might enforce more conservative limits, whereas a development software pool might enforce more liberal limits.

\* Maximum Number of Java Servers (per host) 3

☒ **TIP** The service instance will be provisioned on the first member that satisfies the placement constraints.

6. Click **Next**. The Create Middleware Pool: Targets page appears. In this page, you can add one or more Oracle homes to the Middleware Pool being created.
  1. In the Zone Name drop box, select the PaaS Infrastructure Zone in which the Middleware Pool is to be created. See [Section 9.3, "Creating a PaaS Infrastructure Zone"](#) for details.
  2. Specify the Version of the Oracle Home being added to the Middleware Pool.

---

**Note:** The hosts on the service instance is to be provisioned must contain a Middleware Home. This Middleware Home must be of the same version as the Middleware Home of the WebLogic Domain from which the provisioning profile was created.

---

3. Click **Add** and select the Oracle Home to be added to the Middleware Pool. Click **Submit** to create a software pool.

---

**Note:** You must select only one Oracle Home from a host being added to the Middleware Pool. The Oracle Home must be of the same platform and version you have selected earlier.

---

The newly created pool will appear in the Middleware Cloud Self Service Portal Setup: Middleware Pools page. This Middleware Pool can now be used as a target while creating the Service Template.

## 16.3 Creating a Load Balancer Configuration

Load Balancer Configurations are used to configure and route web traffic to middleware service instances. You can configure external pre-configured load balancers such as Oracle HTTP Server (OHS), Oracle Traffic Director (OTD) or a third party load balancer. If you want to configure a third party load balancer, you must use a custom script to update the load balancer configuration. For more details on this script, see [Section 16.3.1, "Creating a Custom Script"](#). A Middleware Pool or a Deployment Plan can be associated with only one of the load balancer configurations.

### Prerequisites


- The external load balancer being configured must be installed on an Enterprise Manager monitored host and is used to route Web traffic to one or more MWaaS service instances. When a MWaaS service instance is created, updated (scaled up or scaled down), or deleted, the load balancer configuration is automatically updated.
- To route HTTPS Web traffic, you must do the following:
  - Create an OTD configuration (This is required only if you are configuring an Oracle Traffic Director type of load balancer).
  - Create or install a certificate.

For more details configuring SSL authentication, refer to the Oracle Traffic Director Administrator's Guide and the Oracle HTTP Server Administrator's Guide.

To configure a load balancer, follow the steps below:

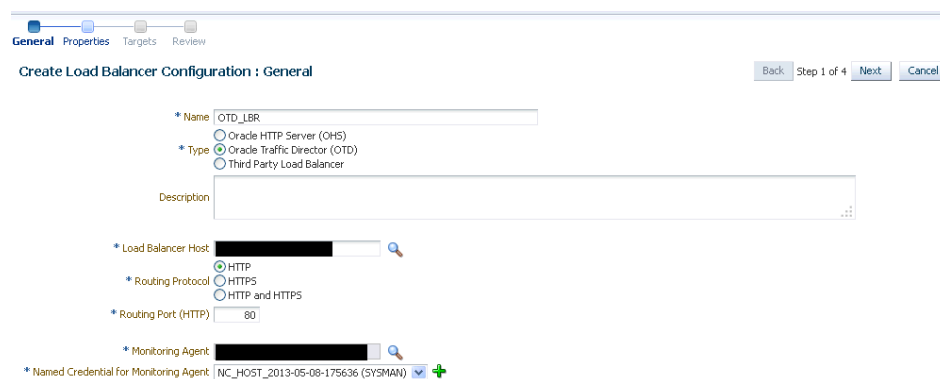
1. Log into Enterprise Manager as a user with the EM\_SSA\_ADMINISTRATOR role.
2. From the **Setup** menu, select **Cloud**, then select **Middleware**. The Middleware Cloud Self Service Portal Setup page appears. Click Load Balancer Configurations. The following page appears.

**Figure 16–5 Load Balancer Configurations**



3. Click **Create** to create a new load balancer configuration.

**Figure 16–6 Load Balancer Configuration: General**



4. Enter the following details:
  - Name of the load balancer configuration.
  - Select the type of load balancer being configured. This can be Oracle HTTP Server (OHS), Oracle Traffic Director (OTD), or a Third Party Load Balancer.

- Select the host on which the load balancer is installed. This must be an Enterprise Manager monitored host.
- Specify the Routing Protocol (HTTP, HTTPS, or both) and the Routing Port number.

---

**Note:** If you are configuring the Oracle HTTP Server, the Routing Port is required only for the http Routing Protocol. In this case, the port you specify must be the default port of the default virtual host.

---

- Select the Monitoring Agent (Management Agent) and specify the Named Credentials.
5. Click **Next**. The Create Load Balancer Configuration: Properties page appears.


**Figure 16–7 Create Load Balancer Configuration: Properties**

If you have selected the Load Balancer Type as:

- **Oracle HTTP Server:** Enter the following details:
  - **Instance Home:** The directory in which the configuration data and binary files pertaining to Oracle HTTP Server instance is stored.
  - **Component Name:** The component in the specific instance that is to be configured.
- **Oracle Traffic Director:** Enter the following details:
  - **Oracle Home:** The directory in which the Oracle Traffic Director binaries are installed.
  - **Admin Protocol:** Specify the protocol (HTTP or HTTPS) to be used by the Oracle Traffic Director Administration Server to forward requests. By default, this is HTTPS.
  - **Admin Listen Port:** Specify the Administration Server listen port number.
  - **SNMP Port:** The SNMP port configured for Oracle Traffic Director. The default port number is 11161.
  - **Instance Node:** Specify the Oracle Traffic Director Instance node name.
  - **Configuration Name:** This field is optional and is required only if you have selected the HTTPS or HTTP and HTTPS Admin Protocol. If you select the HTTP Admin Protocol, and do not specify a value in this field, the OTD Configuration Name is created as config1 by default.

- **Certificate Name:** This field is optional. If you selected the HTTPS or HTTP and HTTPS Admin Protocol, enter the name of the self signed certificate.
  - **Virtual Server Name:** The name of the virtual server. If the virtual server does not exist, a virtual server with the name specified here will be created. In this release, host patterns are not defined for each virtual server in the configuration.
  - **Privileged Named Credentials:** Specify the Named Credentials for the Administration Server.
- **Third Party Load Balancer:** Select a script to configure the load balancer.
6. Click **Next**. The Create Load Balancer Configuration: Targets page appears.

**Figure 16–8 Create Load Balancer Configuration: Targets**



7. Click **Add** to add one or more middleware pools (See [Section 16.2, "Creating a Middleware Pool"](#)) or deployment plans (See [Deploying an Assembly](#)) with which the load balancer will be associated.

---

**Note:** A Middleware Pool or a Deployment Plan can be associated with only one load balancer configuration.

---

8. Click **Next**. Review the configuration and click Submit to create the Load Balancer Configuration. The newly created load balancer configuration appears in the Middleware Cloud Self Service Portal Setup: Load Balancer Configurations table. Click the **Load Balancer Name** link in the table to view the load balancer configuration details. This shows the targets associated with the load balancer configuration and middleware service instances that are using the configuration.

### 16.3.1 Creating a Custom Script

You can use a custom script to configure the load balancer. This script must be present in the Software Library as a directive. The directive must follow these guidelines:

- **LBR\_Work\_Directory:** Specifies the directory in which the `lbr.properties` file will be staged. The name-value pairs present in the `lbr.properties` file are extracted by the custom script and these values are used to configure the load balancer.
- The following properties must be present in the `lbr.properties` file:
  - **THIRDPARTY\_ROUTING\_PROTOCOL:** This can be `http`, `https`, or `both`.
  - **THIRDPARTY\_ROUTING\_HTTP\_PORT:** Enter the `http` port number.
  - **THIRDPARTY\_ROUTING\_HTTPS\_PORT:** Enter the `https` port number.
  - **LBR\_HOST:** The host on which the load balancer is installed.

- **UNIQUE\_URI\_REFERENCE:** The value specified here is used to create routing rules.
- **CONFIGURATION\_ACTION:** The actions that can be performed are:
  - \* **CREATE:** To create a service instance.
  - \* **ADD\_SERVERS:** To scale up or add one or more servers.
  - \* **DELETE\_SERVERS:** To scale down or delete one or more servers.
  - \* **REMOVE:** To delete a service instance.
- **HTTP\_WLS\_MANAGED\_SERVER\_PORT\_LIST:** A comma separated host:port (Managed Servers) list where port is a http port. For example, abc:123, xyz:456.
- **HTTPS\_WLS\_MANAGED\_SERVER\_PORT\_LIST:** A comma separated host:port (Managed Servers) list where port is an SSL port. For example, abc:123, xyz:456.

## 16.4 Configuring Request Settings

You can configure the request settings by specifying when a request can be made, its duration, and so on. To configure the request settings, follow these steps:

1. Log into Enterprise Manager as a user with the **EM\_SSA\_ADMINISTRATOR** role.
2. From the **Setup** menu, select **Cloud**, then select **Middleware**. The Middleware Cloud Self Service Portal Setup page appears.
3. Click the **Request Settings** tab. Specify the following:
  - **Future Reservation Length:** How far in advance a request can be scheduled. You can select:
    - **No Restriction:** A new service request can be scheduled on any date later than the current date.
    - **Restricted Reservation:** A new service request can be scheduled in the future within the period specified here. Any future requests scheduled beyond the Restricted Reservation date will be denied.

For example, if the Restricted Reservation is 1 year, and the current date is June 20, 2013, you can schedule a future request on or before June 20, 2014.
  - **Request Archive Retention:** Period after which the requests are automatically archived. You can select:
    - **No Restriction:** If this option is selected, requests are not archived.
    - **Restricted Retention:** Requests that older than the specified duration will be archived. The date on which a request is archived is based on the last modified date of the request.

For example, if a request has been created on June 10, 2013 and the **Restricted Retention** period specified is 10 days. Suppose there are some associated child requests as follows:

      - Reschedule Request: Submitted on June 12
      - Start Instance: Submitted on June 14
      - Stop Instance: Submitted on June 17
      - Start Instance: Submitted on June 18
      - Delete Instance: Submitted on June 20

Since the date on which the request was last modified is June 20, the request will be purged on June 30th (Restricted Retention = 10 days).

- **Default Retirement Period:** The maximum period for which the service instance can be retained. You can select:
  - **No Restriction:** Requests do not have a end date and can be extended as required.
  - **Restricted Duration:** The period after which the service instance is to be released.

For example, if the Restricted Duration is 15 days, and the request begin date is June 5, 2013, the service instance will be retired or released on June 25, 2013.

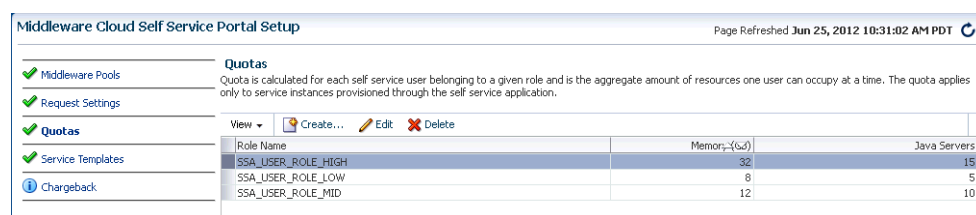
4. Click **Apply**.

## 16.5 Setting Up Quotas

Quota is the aggregate amount of resources that can be granted to each self service user belonging to a certain role. This quota applies only to the service instances provisioned through the Middleware Self Service Portal. To define a quota, follow these steps:

1. Log into Enterprise Manager as a user with the EM\_SSA\_ADMINISTRATOR role.
2. From the **Setup** menu, select **Cloud**, then select **Middleware**. The Middleware Cloud Self Service Portal Setup page appears.
3. Click **Quotas**. Quota is the aggregate amount of resources that can be granted to each self service user belonging to a certain role. This quota applies only to the service instances provisioned through the Middleware Self Service Portal.

**Figure 16–9 Middleware Cloud Self Service Portal Setup: Quotas**



| Role Name          | Memory (MB) | Java Servers |
|--------------------|-------------|--------------|
| SSA_USER_ROLE_HIGH | 32          | 15           |
| SSA_USER_ROLE_LOW  | 8           | 5            |
| SSA_USER_ROLE_MID  | 12          | 10           |

4. Click **Create**. In the Create New Quota window, select the Role Name for which the quota is to be granted. Specify the quota limits for the role as follows:
  - **Memory:** The maximum memory that can assigned to the instances owned by a specific user.
  - **Application Servers:** The number of Application Servers that can be owned by a user at any point in time.

---

**Note:** Only roles that have the EM\_SSA\_USER privileges are displayed in the Select Roles dialog box.

---

## 16.6 Creating a Middleware Service Template

A service template is a standardized service definition that can be used by self service users to provision a service instance into one or more PaaS Infrastructure zones. You can create the following types of service templates. See:

- [Creating a Middleware Physical Service Template](#)
- [Creating a Middleware Virtual Service Template](#)

## 16.7 Creating a Middleware Physical Service Template

A service template can be used to provision an instance of a service into a set of selected PaaS Infrastructure zones.

### Prerequisites

- Set up and enable MWaaS as described in [Section 15.1, "Getting Started with MWaaS for Physical Hosts"](#).
- A Middleware Pool must have been created.
- A PaaS Infrastructure Zone must have been created.
- A copy of a configured Middleware Provisioning deployment procedure with appropriate values must be present in the Software Library.
- The Oracle Home version in the deployment procedure and the Middleware Pool must be the same.
- If Named Credentials are used, the user creating the service template must be the owner of the Named Credentials.

To create a Middleware Physical Service Template, follow these steps:

1. Log into Enterprise Manager as a user with the EM\_SSA\_ADMINISTRATOR role.
2. From the **Setup** menu, select **Cloud**, then select **Middleware**. The Middleware Cloud Self Service Portal Setup page appears.
3. Click the **Service Templates** tab. You can create service templates that can be used to provision an instance of a service in a set of selected zones.
4. Click **Create**. In the Select Service Template window, choose Middleware Service Template for Physical option and click Select to launch the Create New Service Template wizard. The Create Service Template: Name page appears.

**Figure 16–10 Create Service Template: General**

Service Templates

General Configuration PaaS Infrastructure Zones Roles Review

Create Middleware Service Template for Physical : General

Back Step 1 of 5 Next Cancel

The following service template will be made available in the self service portal. Associate a name and description that can help self service users make the right choice for their service template.

\* Name new\_service\_template

Description

\* Deployment Procedure

5. Enter a name and description for the service template. The description must be unique and provide information on the type of service template being created.



6. Click the **Search** icon next to the Deployment Procedure field and select a Configured Middleware Provisioning deployment procedure that specifies the configuration of the WebLogic Domain. The configured deployment procedure can contain only one cluster with one WebLogic Server. You can scaleup or scale down by adding or deleting WebLogic servers to this cluster. Click **Next**.

---

**Note:** Before you create a service template, you must have saved a copy of the Configured Middleware Provisioning Deployment Procedure with appropriate values.

---

7. The Create New Service Template: Configuration page appears. Enter the following details:

- **WebLogic Username:** Enter a value of minimum 8 characters as the name of the user logging into the WebLogic server. This is a mandatory field.
- **WebLogic Password:** Enter a value of minimum 8 characters as the password for the WebLogic server user. This is a mandatory field.
- **Port Range Start - End:** Enter the port range for the WebLogic Server.
- **Expected Memory Consumption:** Specify the memory required for each additional WebLogic Server instance.

This parameter must be specified (cannot be left blank) and the value must match the `-Xmx` value specified in the profile. For example, if the profile specifies `-Xmx512m`, the memory required value must be set to 0.5. Similarly, a profile value of `-Xmx1024m` would need a memory required value of 1, and so on. Refer to the **Note** in [Section 15.1.1, "Creating Provisioning Profiles"](#) for more details.

- **Expected CPU Consumption:** Specify the CPU required for each additional WebLogic Server instance.
- **Prescript and Postscript for Service Instance Creation:** Select the script that is to be executed before and after a service instance has been created.
- **Prescript and Postscript for Service Instance Deletion:** Select the script that is to be executed before and after a service instance has been deleted.

The scripts must meet the following guidelines:

- Directive that is created should be a PERL directive.
- An input parameter named `INPUT_FILE_LOCATION` must be defined.
- The path to the location of a property file is specified in the `INPUT_FILE_LOCATION`.
- The property file contains information about the service instance and the domain.
- **Hosts for Executing Scripts:** Specify one or more target hosts on which the scripts will be executed.
- **Named Credential for execution of scripts:** Specify the named credentials for the hosts.
- **Execution of Scripts on Admin Server Host:** Enter **1** if the scripts are to be executed on the Administration Server.

- **Expose WebLogic and FMW Console URLs to SSA User:** Enter 1 to expose the URLs to the EM\_SSA\_USER.

---

**Note:** This field must be set to 1 only if the EM\_SSA\_USER needs to have complete control over the WebLogic and Fusion Middleware console. The EM\_SSA\_USER will then be able to perform all configuration changes and life cycle management operations from the console.

---

- **Enable Load Balancer:** Enter 1 to enable the load balancer configuration for all services instances that are created with this service template.
8. Click **Next**. The Create New Service Template: PaaS Infrastructure Zones page appears.

**Figure 16–11 Create Service Template: PaaS Infrastructure Zones**

9. Click **Add** to select the zones into which that the service instances can be provisioned.
10. Click the **Search** icon in the Software Pool field and select a Middleware Pool from the list.

---

**Note:** The Middleware Pool that you select must contain Oracle homes with the same WebLogic Server version as specified in the Configure Deployment Procedure you have selected in the Create Service Template: Name page.

---

11. Click **Next**. In the Create Service Template: Roles page, click **Add** to select the self service user roles to which this service template will be available. All users belonging to the selected role can use this service template.

If you need to add more PaaS Infrastructure zones, repeat steps 10 and 11.

12. Click **Submit**. The newly created service template will appear in the Middleware Cloud Self Service Portal Setup: Service Templates page. You can click on the Service Template Name link to view additional details.

To edit an existing MWaaS service template, follow these steps:

1. Click **Edit** in the Service templates page to launch to the Edit Service Template wizard.
2. In this mode, you can modify the description, user name, password, and port range of the WebLogic Server instance being provisioned, add or remove PaaS

Infrastructure Zones, Software Pools into which the template is being deployed, and add or remove users who can use this template.

## 16.8 Creating a Middleware Virtual Service Template

### Prerequisites

- Set up and enable MWaaS as described in [Section 15.2, "Getting Started with MWaaS for Virtual Hosts"](#).
- A deployment plan must be present in the Software Library. To create a deployment plan, you must launch the assembly deployment wizard and save the assembly as a deployment plan on an Oracle VM Zone. See [Section 8.8.4, "Provisioning Guest Virtual Machines Using Oracle Virtual Assemblies \(OVA\)"](#) for details. To launch assembly deployment, from **Oracle VM Zone** target menu, select **Deploy**, then select **Assembly**.
- The deployment plan (assembly) must contain a WebLogic Domain, Managed Server, and an Administration Server.
- The WebLogic Domain must be configured on a single host, with a single cluster and one Managed Server.
- The WebLogic Server version must be 10.3.5.0 or 10.3.6.0.
- A PaaS Infrastructure Zone must have been created.
- A Middleware Pool must have been created.

To create a Middleware Virtual Service Template, follow these steps:

1. Log into Enterprise Manager Grid Control as a user with the EM\_SSA\_ADMINISTRATOR role.
2. From the **Setup** menu, select **Cloud**, then select **Middleware**. The Middleware Cloud Self Service Portal Setup page appears.
3. Click the **Service Templates** tab. You can create service templates that can be used to provision an instance of a service in a set of selected zones.
4. Click **Create**. In the Select Service Template window, choose Middleware Service Template for Physical option and click Select to launch the Create New Service Template wizard. The Create Service Template: Name page appears.

**Figure 16–12 Create Middleware Service Template for Virtual: General**

5. Enter a name and description for the service template. The description must be unique and provide information on the type of service template being created. Click **Next**.

**Figure 16–13 Create Middleware Service Template for Virtual: Configuration**

| Name                                                  | Value | Description                                                                                       |
|-------------------------------------------------------|-------|---------------------------------------------------------------------------------------------------|
| WebLogic Username *                                   |       | Username for the WebLogic Server                                                                  |
| WebLogic Password *                                   |       | Password for the WebLogic Server                                                                  |
| Re-enter WebLogic Password *                          |       | Re-enter Password for the WebLogic Server                                                         |
| Root Password                                         |       | Enter Root Password to override (Virtual Only)                                                    |
| Agent Install User Password                           |       | Enter Agent Install User Password to override (Virtual Only)                                      |
| Host credential name for virtual flow                 |       | Enter name of the host credential to use for provisioning operations (Virtual Only)               |
| Port Range Start                                      |       | Lower Bound for choosing Free Ports (Physical only).                                              |
| Port Range End                                        |       | Upper Bound for choosing Free Ports (Physical only).                                              |
| Expected Memory Consumption (For Use By Placement L   |       | Memory required by each additional instance in GB (Physical only)                                 |
| Expected CPU Utilization (For Use By Placement Logic) |       | CPU required by each additional instance in percentage (Physical only)                            |
| Prescript for Service Instance Creation               |       | Prescript for Service Instance Creation                                                           |
| Postscript for Service Instance Creation              |       | Postscript for Service Instance Creation                                                          |
| Prescript for Service Instance Deletion               |       | Prescript for Service Instance Deletion                                                           |
| Postscript for Service Instance Deletion              |       | Postscript for Service Instance Deletion                                                          |
| Hosts for Executing Scripts                           |       | Use ',' to separate multiple target hosts                                                         |
| Named Credential for execution of scripts             |       | This will be used for the above hosts                                                             |
| Execution of Scripts on Admin Server Host             |       | To execute scripts on Admin Server besides the hosts entered above, enter 1 for 'Yes', 0 for 'No' |
| Expose WebLogic and FMW Console URLs to SSA User      |       | Enter 1 for 'Yes', 0 for 'No'                                                                     |
| Enable Load Balancer                                  |       | Enter 1 to enable, 0 to disable Load Balancer                                                     |

The Create New Service Template: Configuration page appears.

Enter the following details:

- **WebLogic Username:** Enter a value of minimum 8 characters as the name of the user logging into the WebLogic server. This is a mandatory field.
- **WebLogic Password:** Enter a value of minimum 8 characters as the password for the WebLogic server user. This is a mandatory field.

**Note:** The WebLogic Username and Password specified here must be the same as what has been specified in the Deployment Plan. If an incorrect password is entered, the deployment procedure will fail.

- **Root Password:** The root user password for the virtual hosts that will be provisioned.
- **Agent Install User Password:** The password required to deploy the agent on the virtual hosts that will be provisioned.
- **Host credential name for virtual flow:** The name of host credentials to be used for provisioning operations on these virtual hosts.
- **Port Range Start - End:** Enter the port range for the WebLogic Server.
- **Expected Memory Consumption:** Specify the memory required for each additional WebLogic Server instance.

This parameter must be specified (cannot be left blank) and the value must match the -Xmx value specified in the profile. For example, if the profile specifies -Xmx512m, the memory required value must be set to 0.5. Similarly, a profile value of -Xmx1024m would need a memory required value of 1, and so on. Refer to the **Note** in [Section 15.1.1, "Creating Provisioning Profiles"](#) for more details.

- **Expected CPU Consumption:** Specify the CPU required for each additional WebLogic Server instance.

- **Prescript and Postscript for Service Instance Creation:** Select the script that is to be executed before and after a service instance has been created.
- **Prescript and Postscript for Service Instance Deletion:** Select the script that is to be executed before and after a service instance has been deleted.

The scripts must meet the following guidelines:

- Directive that is created should be a PERL directive.
  - An input parameter named `INPUT_FILE_LOCATION` must be defined.
  - The path to the location of a property file is specified in the `INPUT_FILE_LOCATION`.
  - The property file contains information about the service instance and the domain.
- **Hosts for Executing Scripts:** Specify one or more target hosts on which the scripts will be executed.
  - **Named Credential for execution of scripts:** Specify the named credentials for the hosts.
  - **Execution of Scripts on Admin Server Host:** Enter 1 if the scripts are to be executed on the Administration Server.
  - **Expose WebLogic and FMW Console URLs to SSA User:** Enter 1 to expose the URLs to the `EM_SSA_USER`.

---

**Note:** This field must be set to 1 only if the `EM_SSA_USER` needs to have complete control over the WebLogic and Fusion Middleware console. The `EM_SSA_USER` will then be able to perform all configuration changes and life cycle management operations from the console.

---

- **Enable Load Balancer:** Enter 1 to enable the load balancer configuration for all services instances that are created with this service template.
6. Click **Next**. In the Create New Service Template: Zones page, click **Add**. Select the PaaS Infrastructure Zone into which that the service instances can be provisioned. The PaaS Infrastructure Zone you select must contain an Oracle VM Zone. See [Section 9.3, "Creating a PaaS Infrastructure Zone"](#). The Deployment Plan associated with the Oracle VM Zone is displayed. Click **Add** to return to the previous page. The table lists the following details:
- **PaaS Zone:** Name of the PaaS Infrastructure Zone.
  - **Oracle VM Zone:** The Oracle VM Zone present in the PaaS Infrastructure Zone.
  - **Is Exalogic:** Indicates if this zone has been deployed on an Exalogic rack.
  - **Deployment Plan:** The deployment plan associated with the OVM Zone.

You must have created a deployment plan by launching the assembly deployment wizard and saving the assembly as a deployment plan on an Oracle VM Zone. See [Section 8.8.4, "Provisioning Guest Virtual Machines Using Oracle Virtual Assemblies \(OVA\)"](#) for details

---

**Note:** When you request a service with the virtual service template, the service instance is dependent on the deployment plan configuration selected here. For example, if the assembly is created from a source domain with 20 managed servers, the `EM_SSA_USER` can scale up the service instance by a maximum of 20 servers.

---

7. Click **Next**. In the Create Service Template: Roles page, click **Add** to select the SSA user roles to which this service template will be available. All users belonging to the selected role can use this service template.
8. Click **Submit**. The newly created service template will appear in the Middleware Cloud Self Service Portal Setup: Service Templates page. You can click on the **Service Template Name** link to view additional details.

## 16.9 Configuring Chargeback

Click the Chargeback tab to add the middleware pools to the Chargeback Application. On this page, you can determine chargeback zones, services, items, and prices. After Chargeback is set up, you can view charges and plan details. For more details on setting up the Chargeback Service, see [Chapter 22, "Chargeback Administration"](#).

### 16.9.1 Sample Scenario

A PaaS Infrastructure Zone is a group of homogeneous resources such as hosts, or other targets. Each resource in a zone represents a location at which a service instance is to be deployed.

Before you enable or setup MWaaS, you must create a PaaS Infrastructure Zone which allows you to define the placement policy constraints for a specified set of targets and the users to whom this zone will be available. You can set up a zone which contains targets in different locations such as the West Coast, East Coast, and so on.

To set up the portal for the west coast zone, follow these steps:

1. Create a PaaS Infrastructure Zone named *west coast zone* by following the instructions in [Section 9.3, "Creating a PaaS Infrastructure Zone"](#).
2. After the zone has been created, you must create the Middleware Pool. The Middleware Pool is a set of resources used to provision a service instance for a zone.

Create a Middleware Pool (*west coast pool*) and add one or more targets (WebLogic Oracle Homes).

3. Create a service template, associate the west coast zone and the west coast pool, and add users with the `EM_SSA_USER` role to whom this service template is to be available. The service template is then made available to one or more SSA users who can then request services, create data sources, deploy, redeploy services, and so on.

---

## Using the MWaaS Self Service Portal

This chapter provides Self Service Application users with instructions on using the MWaaS Self Service Portal to request and monitor middleware services. It contains the following sections:

- [Using the MWaaS Self Service Portal](#)
- [Requesting a Middleware Service](#)
- [Viewing the Middleware Service Home Page](#)
- [Uploading a Java EE Component to the Software Library](#)
- [Viewing the Application Home Page](#)

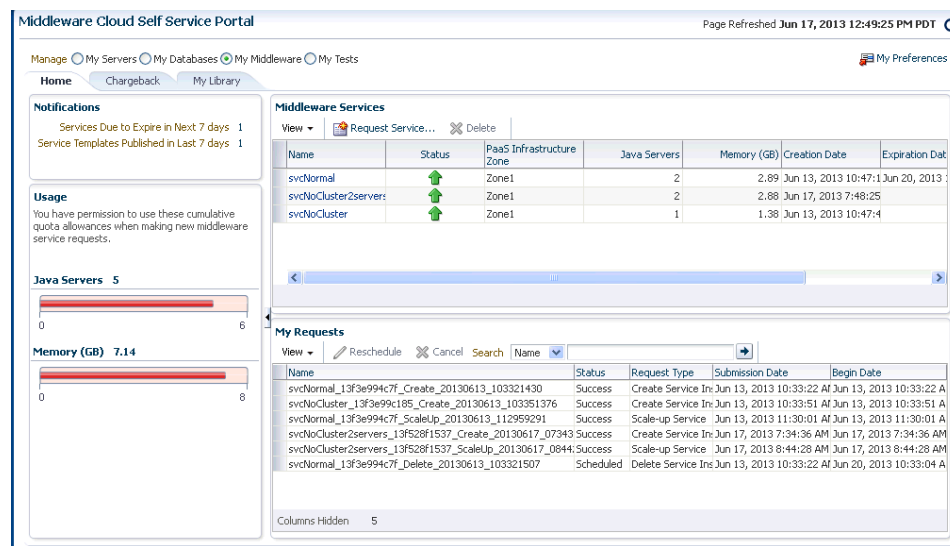
### 17.1 Using the MWaaS Self Service Portal

The Middleware Self Service Portal allows users with the `EM_SSA_USER` role to view, monitor, and request middleware services. Users with the `EM_SSA_USER` role can request middleware services using any of the predefined service templates. They can deploy, undeploy applications, create data sources, scale up, and scale down services with minimal or zero intervention by the `EM_SSA_ADMINISTRATOR` user.

To view the MWaaS Self Service Portal, follow these steps:

1. Log into Enterprise Manager as a user with the `EM_SSA_USER` role.
2. Click the **My Middleware** radio button. The Middleware Cloud Self Service Portal appears.

Figure 17–1 Middleware Cloud Self Service Portal



### 3. This page contains the following tabs:

- **Home:** The Home tab contains the following regions:
  - **Notifications:** This region shows the number of services that are due to expire in the next 7 days and the number of the services templates published within the last 7 days.
  - **Usage:** The total quota (servers and memory) used so far, including all the services that the user has created.
  - **Middleware Services:** This region shows the list of services owned by the SSA user. For each service, the status, the zone being used by the service, the creation date, expiration date, and days until expiration are displayed. To request a service, click **Request Service**. See [Section 17.2, "Requesting a Middleware Service"](#) for details. Click on the **Name** link to drill down to the Service Home page.
  - **My Requests:** This region shows a list of Service Creation, Service Deletion, Scale Up, and Scale Down requests. For each request, the status, type, start date, and submission date of the request is displayed. Select columns from the **View** menu to view more details about the request. You can select a scheduled request and click **Reschedule** to reschedule the request. You can select a scheduled request and click **Cancel** to delete the request. If a request that is scheduled such as **Create** gets canceled, all other associated requests are also canceled.
- **Chargeback:** This tab displays all the chargeback break-up data for the service. The charge incurred will be calculated from the service requested date to current date. This page contains the following regions:
  - **Charge Trend:** This graph show the usage and charge trends over time.
  - **Details:** You can aggregate the chargeback data for the service by resources of metrics. You can export this data to an external application such as a spreadsheet.



- **Charge Plans:** The charge plans that have been assigned to each PaaS Infrastructure Zone are displayed. Select a zone to view the pricing. The rate for each metric in the zone for a specific time period is displayed.
- **My Library:** This tab displays a list of Java EE application components that are available to the SSA user. You can do the following:
  - **Upload:** You can upload a new Java EE application component to the Software Library. See [Section 17.4, "Uploading a Java EE Component to the Software Library"](#) for details.
  - **Delete:** Select a Java EE component from the list and click **Delete** to remove it from the library.
  - **Share / Stop Sharing:** You can share a component with one or more users and roles. Select the component and click **Share**. The Select and Share: Users and Roles window appears. From the **Type** drop list, select User or Role and click **Search**. The users or roles with which the component can be shared. Select one or more users or roles and click **Share**. The selected component is shared with the specified users and roles. To stop sharing the component, select the component and click **Stop Sharing**.
- **My Preferences:** Click the **My Preferences** link to set up your preferences such as the Default Service Type, Default Email Address for Notifications, Default Locale, Default PaaS Infrastructure Zone, and so on. To set the Middleware Cloud Self Service Portal as the default page that appears when you log in to the Self Service Portal, select **Middleware** in the **Default Service Type** drop down list and click **Apply**. For more details on setting the preferences, see [Section 7.1.7, "Viewing My Preferences"](#).

**Figure 17–2 My Preferences (Middleware)**

**General Preferences**

Email Address  [Test Email](#)

Locale

☐ Request Notifications

Request Status

☐ Instance Expiry Notifications

Days before first notification

Days before second notification

☐ Quota Notifications

Quota Threshold (%)

Default Service Type

---

**My Servers Preferences**

Default Server Root Password

Confirm Default Server Root Password

Default Zone for New Requests  🔍 ✖

Default Source Software  🔍 ✖

---

**My Database Preferences**

Default PaaS Infrastructure Zone  🔍 ✖

---

**My Middleware Preferences**

Default PaaS Infrastructure Zone  🔍 ✖

---

**My Tests Preferences**

## 17.2 Requesting a Middleware Service

You can deploy a middleware service to a specific zone with a selected service template. To request a middleware service, follow these steps:

1. Click **Request Service** in the Home page of the Middleware Cloud Self Service Portal.
2. The Select Service Template window appears. A list of service templates available in the zones that the user can access is displayed. Choose a service template and click **Select** to launch the New Service Request page.

**Figure 17-3 New Service Request**

**Middleware Cloud Self Service Portal**

**New Service Request** Submit Cancel

Use this page to request a new middleware service. When you click Submit, the request will be submitted as a job, and will be listed in the My Requests table so that you can track the progress.

Service Template: smallHeapTemplate10.3.6  
Description: 1.5 GB

\* Service Name:

\* PaaS Infrastructure Zone: Zone1

**Schedule**

Start Date: ☒ Immediately ☐ Later  (GMT-08:00) Los Angeles - Pacific Time (PT)

End Date: ☒ Indefinitely ☐ Until

**Create WebLogic User Account (optional)**

You can optionally specify a WebLogic user account to be created on your service after it is provisioned. You will be able to access the WebLogic Server Administration Console using the user name and password specified here.

WebLogic User Name:   
WebLogic Password:   
Confirm WebLogic Password:

**Deploy Application (optional)**

You can optionally select an application to be deployed to your service after it is provisioned. Select an application component from the Software Library or an application archive file from your desktop.

Application Name:

☒ Select an application component from the Software Library  
Application Component:

☐ Select an application archive from your desktop  
Archive:  Browse...

Deployment Plan:  Browse...

3. Enter a name for the service and select the PaaS Infrastructure Zone in which the new service is to be deployed.
4. You can specify the schedule for the service request.
  - **Start Date:** The service can be started **Immediately** or **Later**. If you select the **Later** option, specify the date from which the service is to be available.
  - **End Date:** The service can be available **Indefinitely** or till a specific date. If you select the **Until** option, specify the date till which the service is to be available. The service will be deleted on the specified date and will no longer be available.
5. In the Create WebLogic User Account (optional) region, you can optionally specify a WebLogic user account to be created on your service after it is provisioned. Enter the WebLogic User Name and WebLogic Password. You will be able to access the WebLogic Server Administration Console using the user name and password specified here.
6. In the Deploy Application region, you can optionally choose to deploy an application to the new service being created. Enter the following details:
  - **Application Name:** Enter a name for the application that is to be deployed.
  - **Application Component:** Select a Java EE application component from the Software Library or an application archive from your desktop.
  - **Deployment Plan:** Select an xml file that contains the deployment options. This file overrides the deployment plan (if any) specified with the application component.

---

**Note:** If you have set the Expose WebLogic and FMW Console URLs to SSA User field to 1 while Creating a Service Template, you will be prompted for the WebLogic Username and Password here.

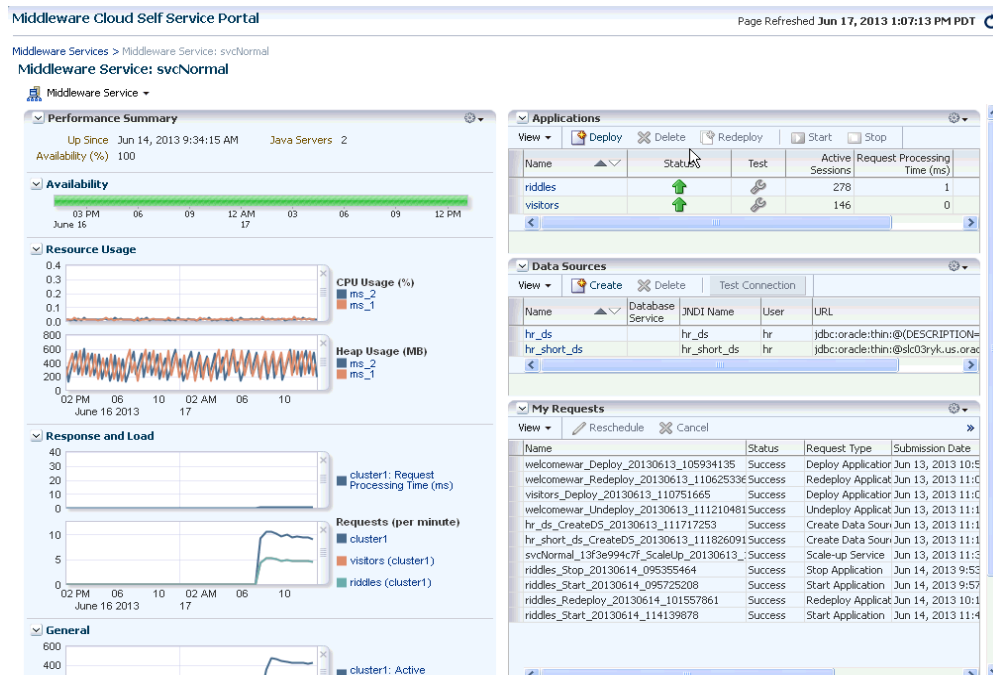
---

7. Click **Submit** to create the service request. This request will now appear in the My Requests table in the Home page of the Middleware Cloud Self Service Portal. You can monitor the status of this request in this table. When the request has been completed, you can refresh the page to view the new service in the Services table.

## 17.3 Viewing the Middleware Service Home Page

The Middleware Service Home page provides detailed information about the services you have requested using the Request Service option. To view the Middleware Service Home page, click the **Name** link in the Middleware Services region in the Middleware Cloud Self Service Portal: Home tab.

**Figure 17–4 Middleware Service Home**



This page contains the following regions:

- **Performance Summary:** Shows the availability status of the service, the number of servers, and the performance metrics for the service. The metrics displayed are the metrics for the underlying WebLogic cluster associated with this service.

- **Resource Usage:** Shows the CPU and Heap Usage charts for each server in the cluster.

**Note:** The server names are based on the name of the first server configured in the middleware provisioning profile. For example, if the first managed server is named as ms\_1, after scale up, the subsequent servers are named as ms\_2, ms\_3 and so on.

- **Applications:** Lists all the applications deployed to this service. For each application, the name, status, and the key metrics such as Active Sessions, Request Processing Time, and Request per Minute are displayed. Click the **Name** link to drill down to the Application Home page. You can do the following:
  - **Deploy:** Click **Deploy** to deploy the application to the service. See [Section 17.3.1, "Deploying and Redeploying an Application"](#) for details.
  - **Delete:** Select an application and click **Delete** to undeploy an application.

- **Redeploy:** If you have an updated Java EE application component, you can select the application and click **Redeploy** to redeploy the application. See [Section 17.3.1, "Deploying and Redeploying an Application"](#) for details.
- **Start / Stop:** Select an application and click **Start** or **Stop** to start or stop the application.
- **Test:** Select an application and click the **Test** icon. A list of URLs is displayed in the Application URLs dialog box. Click the URL to navigate to the associated page.
- **Data Sources:** This region shows all the data sources that are available for this service. For each data source, the Name, Database Service Name, JNDI name and the URL for the database is displayed. If the data source has been created using a database service, the name is displayed in the Database Service column. Click on the name to drill down to the Database Service Home page. You can create new data sources, edit or delete existing data sources, and test the connection to any specified data source.
- **Response and Load:** Shows the request processing time and number of requests per minute for the current service.
- **My Requests:** This table shows all the requests that have submitted for this service. These include requests such as Deploy, Undeploy, Redeploy, Scale Up, Scale Down, Start, Stop, Create, Edit, and Delete Data Source.

### 17.3.1 Deploying and Redeploying an Application

You can deploy a preconfigured Java EE application to a Service Instance. To deploy an application, follow these steps:

1. Log into Enterprise Manager as a user with the EM\_SSA\_USER role and navigate to the Middleware Cloud Self Service Portal page.
2. Click on a **Name** link in the Middleware Services region. The Middleware Service Home page appears.
3. Click **Deploy** in the Applications table. The Deploy Application page appears.

**Figure 17–5 Deploy Application**

4. Enter a name for the application.
5. You can choose to:
  - **Select an application component from the Software Library:** Click the **Search** icon and select a Java EE application component from the Software Library.
  - **Select an application archive from the Software Library.** Click **Browse** and select an application archive from your desktop.

6. In the Deployment Plan field, you can optionally specify an xml file that will override the default values of the application component's deployment plan.
7. Click **Deploy**. After the request to deploy the application has been created, you will return to the Middleware Service Home page where a confirmation message indicates that the application has been deployed.

#### 17.3.1.1 Redeploying an Application

You can redeploy an application that has been deployed to a Service Instance. To redeploy an application, select a deployed application from the Applications table and click **Redeploy**. In the Redeploy Application page, select the Application Component that is to be redeployed, the deployment plan if any, and click **Redeploy**. You will see a confirmation message indicating that the application has been redeployed.

#### 17.3.1.2 Deleting an Application

You can delete (undeploy) an application that has been deployed. Select the application to be deleted from the Applications table and click **Delete**. A confirmation message is displayed. Click **Yes** to undeploy the application.

### 17.3.2 Scaling Up and Scaling Down a Service

You can scale up or scale down a service. On the Middleware Service Home page, select the following options from the Middleware Service menu:

- **Scale Up:** You can scale up a service by adding one or more servers. The number of servers you can add is limited by your quota. Click **Scale Up** and enter the number of servers by which the service is to be scaled up. Click **Submit** to submit the scale up request and add more servers.
- **Scale Down:** Click **Scale Down** and enter the number of servers by which the service is to be scaled down. Click **Submit** to submit the scale down request and reduce the number of servers.

### 17.3.3 Creating a Data Source

You configure database connectivity by adding data sources to your WebLogic domain. WebLogic JDBC data sources provide database access and database connection management. Before an application can access a database, it must be configured to use the JDBC data source.

1. Log into Enterprise Manager as a user with the EM\_SSA\_USER role and navigate to the Middleware Cloud Self Service Portal page.
2. Click **Create** in the Data Sources table in the Middleware Service Home page. The Create Data Source page appears.

**Figure 17–6 Create Data Source**

**Create Data Source** Create Cancel

Use this page to create a data source on this middleware service.

\* Data Source Name

JNDI Name

\* Database Driver Class

Select the database by choosing a database service or by entering the database URL.

☐ Select Database Service

Database Service

URL

☒ Specify Database URL

\* URL

Specify the database credentials, then test the connection if desired.

\* Database User Name

\* Password

\* Confirm Password  Test Connection

3. Enter the Data Source Name and the JNDI Name.
4. Select the Database Driver Class.
5. If you have a database service instance configured, you can select the Database Service and the corresponding URL is displayed. If not configured, you must manually enter the Database URL.
6. Specify the Database User Name and Password.
7. Click **Create** to create the data source.

#### 17.3.3.1 Editing a Data Source

You can edit the properties of a data source. Click the Name link in the Data Sources table in the Middleware Service Home page. The Edit Data Source page appears. You can edit the properties of the data source. If some of the properties are changed, the data source needs to be redeployed and applications using this data source need to be restarted.

## 17.4 Uploading a Java EE Component to the Software Library

You can upload pre-configured ear, war, jar, or rar files to the Software Library. These files are uploaded to the Software Library as Java EE Application components, which are then deployed to your service. To upload a file to the Software Library, click the **My Library** tab. The Upload to Software Library field appears.

**Figure 17–7 Upload to Software Library**

Enter the name and description of the file being uploaded. You can then choose to upload files as follows:

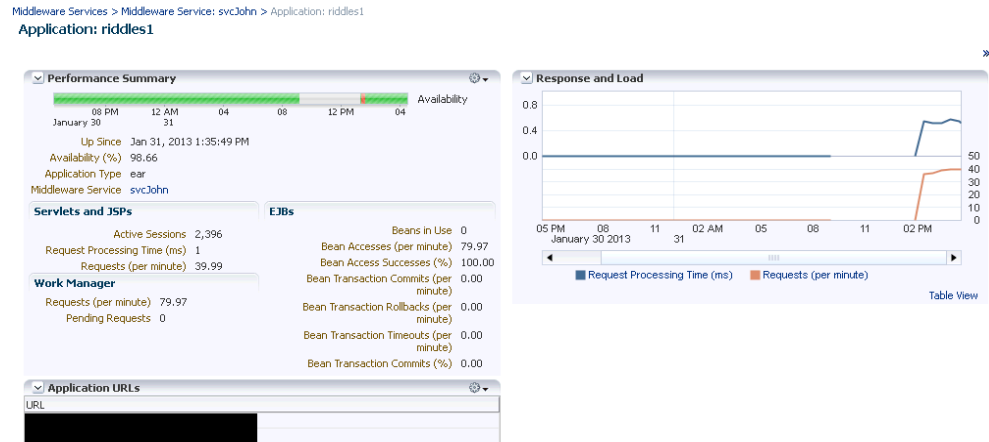
- **Individual Files:** If you choose the **Upload Files Individually** option, you need to specify the following:
  - **Archive:** Select the archive file to be uploaded to the Software Library. This can be an ear, war, jar, or rar file. You can also optionally specify the following:
    - \* **Deployment Plan:** This is an xml file containing the deployment options for the Java EE Application component.
    - \* **Pre Deploy Script:** This is a script containing WLST commands. The Management Agent runs this script on the Administration Server associated with each instance before the application is deployed. You can use this script to create data sources, JMS end points, and any other resources that might be needed by the application that is being deployed.
    - \* **Post Deploy Script:** This is a WLST script that is executed by Management Agent on the Administration Server after the application is deployed. You can use this script to perform any post deployment configuration. For example, if you need to roll back and undo the changes made by the pre deploy script, you can select this option.
- **Zip File:** If you select the **Upload Zip File** option, you must upload a zip file that contains the JavaEEAppComp.manifest file that describes the contents of the zip file.

Click **Upload** to upload the file to the Software Library.

## 17.5 Viewing the Application Home Page

To view the Application Home page, select an application, and click the **Name** link in the Applications table on the Middleware Service Home page.



**Figure 17–8 MWaaS Application Home**

This page displays performance and availability for the application, along with a chart that shows response and load. You can redeploy, delete (undeploy), start, or stop the application using options from the **Application** menu.



## Monitoring a MWaaS Cloud

This chapter describes Enterprise Manager support for monitoring and administering an MWaaS Cloud. It contains the following sections:

- [Viewing the Middleware Pool Home Page](#)

### 18.1 Viewing the Middleware Pool Home Page

The Middleware Pool Home page provides detailed information of the Middleware Pool. To view this page, click the **Pool Name** link in the Middleware Pools region of the Middleware Cloud Self Service Portal Setup: Middleware Pools page. The Middleware Pool Home page contains the following regions:

**Figure 18–1** *Middleware Pool Home*

The screenshot displays the 'my\_mwpool' Middleware Pool Home page. The page header shows the pool name and a refresh button. The main content area is divided into several sections:

- General**: Contains a summary table with columns for Description, Placement Policy Constraints, Java Servers, and Load Balancer.
 

| Summary                                              | Placement Policy Constraints                 | Java Servers          | Load Balancer                                     |
|------------------------------------------------------|----------------------------------------------|-----------------------|---------------------------------------------------|
| Description: Parent PaaS Infrastructure Zone omszone | Maximum Number of Java Servers (per Host): 1 | Total Java Servers: 0 | Load Balancer Name: OTD_LBR                       |
| Oracle Homes: 1                                      |                                              |                       | Load Balancer Type: Oracle Traffic Director (OTD) |
| Oracle Home Version: 10.3.5.0                        |                                              |                       |                                                   |
- Service Templates**: A table with columns Name, Load Balancer Enabled, and Description. It shows 'No data to display.'
- Service Instances**: A table with columns Name, Status, Expiration Date, Template Name, and Oracle Homes. It shows 'No data to display.'
- Members**: A table with columns Name, Type, Host, and Home Location. It lists one member: 'WebLogicServer10\_3\_5\_0\_slc02oys\_8325' with a home location of 'middleware\_p'.

- **Summary:** This region shows the following information:

- **Parent PaaS Infrastructure Zone:** The zone with which the Middleware Pool is associated. Click on the link to drill down to the PaaS Infrastructure Zone Home page.
- **Oracle Homes:** The total number of Oracle Homes in the Middleware Pool.
- **Oracle Home Version:** The version of the Oracle homes.
- **Placement Policy Constraints:** The total number of WebLogic servers in the Middleware Pool.
- **Total Java Servers (per host):** The maximum number of WebLogic servers that can be running on each host.
- **Load Balancer:** This region shows the following details:
  - **Load Balancer Configuration:** The name of the load balancer configuration associated with the Middleware Pool. Click on the link to view the load balancer details.
  - **Load Balancer Type:** The type of load balancer. This can be Oracle HTTP Server, Oracle Traffic Director, or Third Party Load Balancer.
- **Service Templates:** The number of service templates with which the Middleware Pool is associated. For each service template, the name, a column indicating if the load balancer configuration has been enabled, and the description of the template is displayed.
- **Service Instances:** The service instances running on the Middleware Pool are displayed.
- **Members:** This region shows a list of members in the Middleware Pool. The name of the target, the target type, the host on which it is running, the location of the Oracle Home, and the location of the Middleware Home is displayed.

# Part V

---

## Enabling Testing as a Service (TaaS)

This section provides an overview of Testing as a Service (TaaS), describes the configuration steps, and the testing self service portal.

This part contains the following chapters:

- [Chapter 19, "About Testing as a Service"](#)
- [Chapter 20, "Setting Up Testing as a Service"](#)
- [Chapter 21, "Using the Testing as a Service Portal"](#)



---

## About Testing as a Service

Testing as a Service (TaaS) is a cloud-based platform for delivering automated application testing services. It is a self-service solution designed for private clouds that:

- Orchestrates the testing process end-to-end by automating the provisioning of test labs including applications under test and test tools.
- Executes load and/or functional test scripts against the application.
- Provides rich application monitoring and diagnostics data for analysis.
- Has sophisticated chargeback facility for metering and charging the usage of the testing cloud by end-users.

Delivered through Enterprise Manager Cloud Control, TaaS makes it possible to significantly reduce testing time and costs without compromising quality and enables organizations to be more agile in delivering critical business applications to their users.

Key TaaS features include:

- Self-service portal for executing application load and functional tests.
- Test library with full access and security controls which stores all the test assets (scripts, databank files, scenario description and user configurable parameters) available to end users.
- On-demand automated provisioning of complete test labs that includes ability to deploy complex multi-tier applications, test tools, and test scripts. Pre-built assembly for Oracle Load Testing, as well as custom test driver.
- Sharing of cloud hardware resource pools to help maximize hardware utilization while abiding by security policies.
- Integrated and rich application monitoring and diagnostics for middle tier and data tier, to ensure detection of bottlenecks and problems for the application under test.
- Resource metering and chargeback for tests that require to be charged to cost centers. These metering capabilities allow for tracking and charging for compute resources used.

### 19.1 TaaS Roles

There are three roles used in TaaS: Test Administrators, Test Designers, and Testers. The following sections describe these roles in more detail.


- [Test Administrators](#)

- [Test Designers](#)
- [Testers](#)

### 19.1.1 Test Administrators

TaaS enables Test Administrators to manage the application infrastructure for testing activities, define user activities for self-service testing, and manage test environments and test assets. See [Figure 19–1](#).

**Figure 19–1 Testing Cloud Administration Activities**



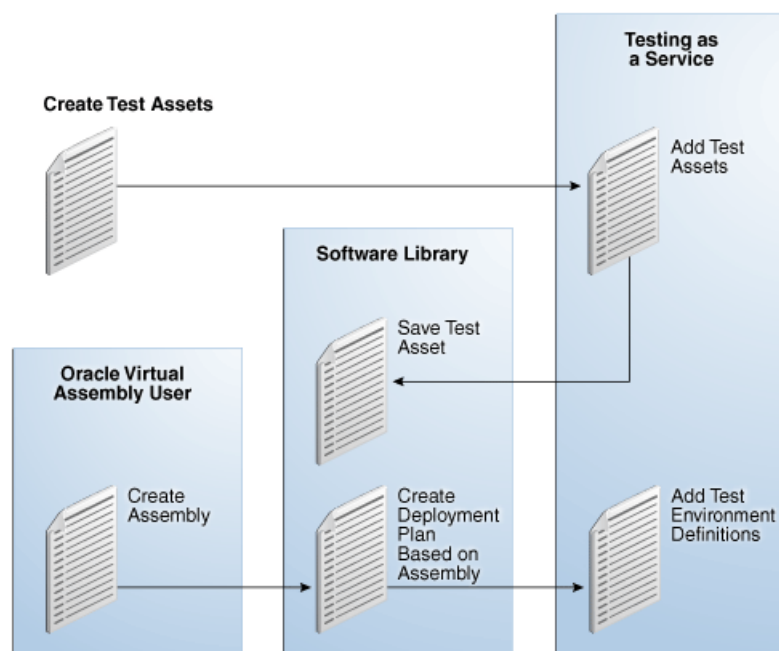
From TaaS, the Test Administrator can:

- Set Up Applications
- Set Up Test Drivers
- Set Up Quotas
- Publish to User Roles
- Set Up Chargeback Service

### 19.1.2 Test Designers

TaaS enables Test Designers to create test assets and test environments that can be made available to users to perform testing activities. Test Designers create test scripts and testing scenarios by using the test drivers native capabilities. These scenarios include which parameters are configurable at run-time. See [Figure 19–2](#).



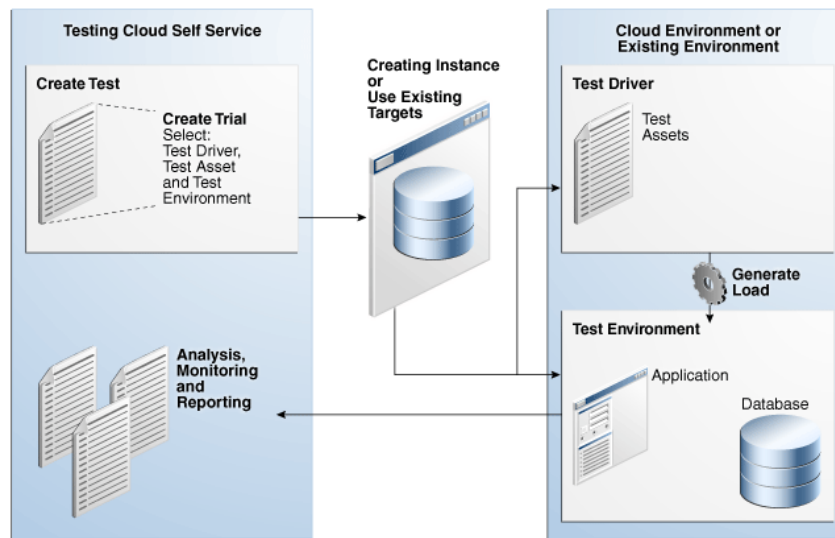
**Figure 19–2 Testing Cloud Designer Activities**

From TaaS, the Test Designer can:

- Create Test Assets
- Create Test Environments

### 19.1.3 Testers

TaaS allows Testers to create tests from test assets and test environments that have been published to them by the Test Administrator. The Tester runs trials based on the test asset and environment, monitors ongoing trials, and analyzes the results of completed trials. The Tester can apply patches to an environment and rerun the trials to compare the results and determine how changes to the environment affect key metrics and overall performance. See [Figure 19–3](#).

**Figure 19–3 Testing Cloud End User Activities**

From TaaS, the Tester can:

- Verify Available Test Assets and Environments
- Create Tests
- Create and Run Trials
- Share Tests
- Monitor On-Going Trials
- Report and Compare Trial Data from Previous Trials
- View Activity and Usage Information
- View Deployment Information

## Setting Up Testing as a Service

This chapter provides an overview of the steps involved in setting up testing as a service (TaaS).

Many of these setup tasks are also required to use Infrastructure as a Service; if you have already set up IaaS, you will have already completed many of these tasks. This chapter contains the following sections:

- [Getting Started](#)
- [Creating Test Administrators](#)
- [Creating an Assembly Component](#)
- [Creating a Deployment Plan](#)

### 20.1 Getting Started

The table below provides an overview of TaaS setup tasks.

**Table 20–1** *Getting Started with TaaS*

| Step | Task                                                                                                                                                                   | Role                |
|------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------|
| 1    | Install the TaaS Plug-In<br>See <a href="#">Section 3.3, "Deploying the Required Plug-ins"</a>                                                                         | Super Administrator |
| 2    | Create Testers<br>See <a href="#">Section 20.2, "Creating Test Administrators"</a>                                                                                     | CLOUD_ADMIN         |
| 3    | Create a Virtualization (VT) Assembly with the Required Assembly Zip files<br>See <a href="#">Section 20.3, "Creating an Assembly Component"</a>                       | CLOUD_ADMIN         |
| 4    | Install Oracle VM Manager (OVM)<br>This is part of Setting the Cloud Infrastructure. See <a href="#">Chapter 3, "Setting Up the Cloud Management Infrastructure"</a> . | CLOUD_ADMIN         |
| 5    | Install VT Plug-in from Enterprise Manager<br>See <a href="#">Section 3.3, "Deploying the Required Plug-ins"</a>                                                       | CLOUD_ADMIN         |
| 6    | Register OVM<br>See <a href="#">Section 4.3, "Registering the Oracle VM Manager"</a>                                                                                   | CLOUD_ADMIN         |
| 7    | Register Oracle VM Server (OVS)<br>Reimage steps<br>See <a href="#">Section 4.5, "Discovering a Virtual Server"</a>                                                    | CLOUD_ADMIN         |

**Table 20–1 (Cont.) Getting Started with TaaS**

| Step | Task                                                                                                                                                          | Role        |
|------|---------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------|
| 8    | Add Storage<br>See <a href="#">Section 4.10, "Creating a Storage Repository"</a>                                                                              | CLOUD_ADMIN |
| 9    | Refresh Mount Locations                                                                                                                                       | CLOUD_ADMIN |
| 10   | Add Storage Repository<br>See <a href="#">Section 4.10.1, "Performing Administrative Operations on a Storage Repository"</a>                                  | CLOUD_ADMIN |
| 11   | Present Server<br>See <a href="#">Section 4.11, "Presenting the Storage Repository"</a>                                                                       | CLOUD_ADMIN |
| 12   | Create Server Pool<br>See <a href="#">Section 4.14, "Creating a Virtual Server Pool"</a>                                                                      | CLOUD_ADMIN |
| 13   | Create Zone<br>See <a href="#">Section 4.15, "Creating a Zone"</a>                                                                                            | CLOUD_ADMIN |
| 14   | Import Assembly into Storage Repository<br>See <a href="#">Section 4.12, "Importing Assemblies, Templates, and Virtual Disks into the Storage Repository"</a> | CLOUD_ADMIN |
| 15   | Set Up Network<br>See <a href="#">Section 4.7, "Setting Up Networks"</a>                                                                                      | CLOUD_ADMIN |
| 16   | Generate MAC Addresses<br>See <a href="#">Section 4.7.1, "Generating MAC Addresses"</a>                                                                       | CLOUD_ADMIN |
| 17   | Create a Deployment Plan<br>See <a href="#">Section 20.4, "Creating a Deployment Plan"</a>                                                                    | CLOUD_ADMIN |

## 20.2 Creating Test Administrators

Follow these steps to create Test Administrators:

1. Log in to Enterprise Manager as a Super Administrator or a user who has privileges to create a new administrator.
2. From **Setup**, select **Security**, then select **Administrator**.
3. Create the TAASADMIN administrator with the EM\_SSA\_ADMINISTRATOR role.
4. Create the TAASDESIGNER administrator with the EM\_TAAS\_DESIGNER role.
5. Create the TAASUSER administrator with the EM\_SSA\_ROLE.

## 20.3 Creating an Assembly Component

To create a VT assembly, perform the following steps:

1. From the Enterprise menu, select **Provisioning and Patching**, then select **Software Library**.
2. Select the Root directory. Select **Actions** then select **Create Folder**.
3. Select the directory which has write access. From the **Actions** menu, select **Create Entity**, then select **Virtualization**.
4. Select Assembly as the type and click **OK**.

5. Provide the required value and click **Next**.
6. Select the upload mechanism. Click **OK**.
7. Click **Add** button to enter file name and click **OK**.
8. Continue to click **Next** until you are on the Review page and click **Save and Upload**.

If you are uploading a file from the OMS or Agent, then you need to Upload Files to a destination where you want the files to be uploaded.

## 20.4 Creating a Deployment Plan

Perform the following steps to create a Deployment Plan:

1. From Enterprise menu, select **Cloud**, then select **Infrastructure Home**.
2. In the right panel, select a zone and right click.  
**Note:** A Deployment Plan created from Zone is the only supported option. (If the deployment plan is created from Server Pool or Oracle VM Server, that deployment plan is not visible in TaaS.)
3. Select **Deploy Assembly**. Select the Assembly created in the Software Library or assembly directly imported in the storage repository. Click **Next**.
4. On the Assembly Deployment: Network and Agent page, select **Enable Agent Push Configuration**. The push means that the Enterprise Manager Agent will be pushed from Enterprise Manager after OVM creation.
5. Provide the Installation Base Directory, Platform, User Name, and Password.
6. Click **Save as Plan**. Provide the name of the plan and select the software library folder. Any other information to provide depends on the assembly you want to install.



## Using the Testing as a Service Portal

After the TaaS environment has been setup, Test Administrators and Test Designers must perform their up front tasks before the Tester can start testing. In general, the tasks are:

- **Test Administrator** gives access and quotas for environments on the cloud, test assets, and virtual user test limits for the Oracle Load Testing test driver.
- **Test Designer** creates test assets (pre-packaged test scripts and scenarios) and test environments (templates for application-under-test, based on Oracle Assemblies) for users to consume.
- **Tester** is the consumer of the test assets and test environments for testing created by the Test Designer. The Tester also plans, executes, and analyzes the test results.

**Table 21–1 Tasks in TaaS**

| Task                                          | How To                                                           | Role                                      |
|-----------------------------------------------|------------------------------------------------------------------|-------------------------------------------|
| Ensure Prerequisites Are Met                  | <a href="#">Prerequisites</a>                                    | Test Administrator, Test Designer, Tester |
| Add application                               | <a href="#">Setting Up Applications</a>                          | Test Administrator                        |
| Define Test Drivers                           | <a href="#">Setting Up Oracle Load Testing Test Drivers</a>      | Test Administrator                        |
| Define Custom Test Driver Types               | <a href="#">Setting Up Custom Test Driver Types</a>              | Test Administrator                        |
| Set Up Quotas                                 | <a href="#">Setting Up Quotas</a>                                | Test Administrator                        |
| Publish Test Environments and Test Assets     | <a href="#">Publishing Test Environments and Test Assets</a>     | Test Administrator                        |
| Set Up Chargeback Service                     | <a href="#">Setting Up Chargeback Service</a>                    | Test Administrator                        |
| Create Test Asset                             | <a href="#">Creating Test Assets</a>                             | Test Designer                             |
| Create Test Environments                      | <a href="#">Creating Test Environments</a>                       | Test Designer                             |
| Verify Available Test Assets and Environments | <a href="#">Verifying Available Test Assets and Environments</a> | Tester                                    |
| Create Tests                                  | <a href="#">Creating Tests</a>                                   | Tester                                    |
| Create and Run Trials                         | <a href="#">Creating and Running Trials</a>                      | Tester                                    |
| Monitor On-Going Trials                       | <a href="#">Monitoring On-Going Trials</a>                       | Tester                                    |
| Compare Trial Results                         | <a href="#">Comparing Trial Results</a>                          | Tester                                    |
| View Activity and Usage Information           | <a href="#">Viewing Activity and Usage Information</a>           | Tester                                    |

**Table 21–1 (Cont.) Tasks in TaaS**

| Task                        | How To                                         | Role   |
|-----------------------------|------------------------------------------------|--------|
| View Deployment Information | <a href="#">Viewing Deployment Information</a> | Tester |

## 21.1 Prerequisites

Ensure the following prerequisites are met when using the Testing Cloud Self Service Portal.

### Before You Do Anything

- Ensure Enterprise Manager Cloud Control is set up to manage the cloud environment.
- Ensure Self Service Application Plug-in is installed because Testing as a Service is part of the SSA plug-in.

### Prerequisites Specific to Test Designer

- Has the privileges to access the Testing as a Service (TaaS) Test Designer options. This user must have the `EM_TAAS_DESIGNER` role.
- Must have knowledge about what applications are to be tested using Enterprise Manager and their structure (such as database servers, application servers, and configurations for the servers.)
- Has created test scripts using the Oracle Application Testing Suite scripting tool OpenScript, the Oracle Load Testing scenario, and the self-contained .scnzip file if creating Oracle Load Testing test assets. For information on how to create and record scripts, refer to the Oracle Application Testing Suite documentation.
- Must be familiar with Enterprise Manager Cloud Control features.
- Has created Test Assets based on scripts and scenarios for Oracle Load Testing or other test drivers and Test Environments based on Oracle assemblies.

### Prerequisites Specific to Test Administrator

- Has the privileges to access the Testing Cloud Self Service Portal administrator options. This user must have the `EM_SSA_ADMINISTRATOR` role.
- Must have knowledge about what applications are being deployed, managed, and tested using Enterprise Manager and their structure (such as database servers, application servers, and configurations for the servers.)
- Must have basic DB administration knowledge and have a good understanding of virtualization technologies.
- Must be familiar with Enterprise Manager Cloud Control features.
- Ensure that the Test Designer has created Test Assets and Test Environments.
- Ensure that the Testers have the privileges to access the Testing Cloud Self Service Portal user options.
- Has assigned test assets and test environments to the Tester role.

### Prerequisites Specific to Testers

- Ensure that the Test Designer has created Test Assets and Test Environments.



- Ensure that the Test Administrator has assigned Test Assets and Test Environments to the Tester role (EM\_SSA\_USER).
- Have the privileges to access the Testing as a Service Tester options. This user must have the EM\_SSA\_USER role.
- Have the privileges to access the Testing Cloud Self Service Portal user options.
- Must be familiar with Enterprise Manager Cloud Control features.

## 21.2 Test Administrator Tasks

The following tasks are performed by the Test Administrator.

- [Setting Up Applications](#)
- [Setting Up Oracle Load Testing Test Drivers](#)
- [Setting Up Custom Test Driver Types](#)
- [Setting Up Quotas](#)
- [Publishing Test Environments and Test Assets](#)
- [Setting Up Chargeback Service](#)

### 21.2.1 Setting Up Applications

The Testing Cloud Self Service Portal allows Testing Administrators to manage the application infrastructure for testing activities, define user activities for self-service testing, and manage test environments and test assets. Specific responsibilities include:

- Manage the test drivers.
- Manage user roles and access to test drivers (specify which users have access to which plug-in.)
- Manage access to test environments and test assets based on users and roles.
- Assign the OVM and Zone privileges to the following roles and users:
  - OVM - View privilege to Test Designer and Test Administrator
  - Zone - Operator privilege to Test Designer and Test Administrator

Since TaaS cannot know before hand the zones and Oracle Virtual Machines present in user deployments, the Super Administrator, Cloud Administrator, or respective target owner will decide which zones and Oracle Virtual Machines are to be used.

**Note:** Ensure the prerequisites are met before using the Testing Cloud Self Service Portal (see [Prerequisites](#)).

Follow this process to set up applications to be tested.

1. Log in to Enterprise Manager Cloud Control using Testing Cloud Self Service Portal Testing Administrator credentials.
2. From the **Setup** menu, select **Cloud**, then select **Testing**. The Testing Administrator page appears.
3. Click **Applications** located in the left panel to define applications to be tested.
4. Click **Add** to define the name and description of the application to make available for testing.

5. Define the name and description and click **Save**.

## 21.2.2 Setting Up Oracle Load Testing Test Drivers

**Note:** Ensure the prerequisites are met before setting up Oracle Load Testing Test Drivers (see [Prerequisites](#)).

The Testing as a Service Portal allows you (Testing Administrator) to Register an Oracle Load Testing test driver to be used by Testers during testing. The Test Driver is the package of software that Enterprise Manager uses to generate workload against test system based on the Deployment Plan. To register an Oracle Load Testing Driver, follow these steps:

1. Log in to Enterprise Manager Cloud Control using Testing Cloud Self Service Portal Administrator credentials.
2. From the **Setup** menu, select **Cloud**, then select **Testing**. The Testing Administrator page appears.
3. Click **Test Drivers** located in the left panel. The Test Driver page appears.
4. From the **Register** menu, select **Oracle Load Testing**.
5. Enter a meaningful name for this Test Driver.
6. Click the magnifying glass for selecting a deployment plan on which the test driver is based. In the resulting dialog box, select the Deployment Plan that should be used for the Oracle Load Testing test driver and click **Select**.

Ensure that Test Designers and Administrators create the deployment plan "on a zone". Also, ensure they have right privileges to this zone. For additional information, see [Section 3.2, "Setting Up Self Update"](#). For information regarding deployment plans, see *Oracle Enterprise Manager Lifecycle Management Administrator's Guide*.

7. Enter a description for the Test Driver and click **Register**.
8. Select a Test Driver in the list then click on the component to drill down to the Deployment Plan details.

## 21.2.3 Setting Up Custom Test Driver Types

**Note:** Ensure the prerequisites are met before setting up Custom Test Driver Types (see [Prerequisites](#)).

The Custom Test Driver allows testers to use test tools from third party vendors together with TaaS. Use Custom Test Driver Types to define and register test-driver types which meet your specific testing needs.

Oracle Load Testing is a *type* of driver, and since it is created by Oracle, we (Oracle) know the specifics of the driver. However, when you create a test driver, you need to direct your scripts on how to launch and monitor your tests. You can register your own type of test drivers as long as you can build a deployment plan around your assembly which contains the testing product and software.

**Note:** If you are using third party test automation tools, these tools need to be registered as Custom Test Drivers.

To register a custom test driver, follow these steps:

1. Log in to Enterprise Manager Cloud Control using Testing Cloud Self Service Portal Administrator credentials.

2. From the **Setup** menu, select **Cloud**, then select **Testing**. The Testing Cloud Self Service Setup page appears.
3. Click **Custom Types** located in the left panel. The Custom Test Driver Types page appears
4. Click **Create**.
5. On the Add Custom Test Driver Type dialog, provide the following information:
  - a. Enter a meaningful name and description.
  - b. The Working directory is the folder on the virtual machine (VM) created from your test-driver where the operations will occur.
  - c. The Results directory is the folder on the VM created from your test-driver and where the results of your tests will appear, TaaS will zip up the entire folder, send it to Enterprise Manager.
  - d. Pre-Execution Directive is a file which performs pre-execution steps for your test. Click **Browse** to locate the file.
  - e. Execution Directive is a file which launches the test. Click **Browse** to locate the file.
  - f. Execution Monitoring Directive is a script which TaaS will run on the test driver, to determine if the test is running or complete (and its outcome). Click **Browse** to locate the file.

When the test is complete, TaaS reports the outcome of the test. TaaS looks at the value returned by this script to determine the test outcome (passed, failed, warning).

**Note:** For TaaS to know the outcome of the trial, ensure that your script prints one of the following outcomes: passed, failed or warning, as the last line to the standard output.

- g. Post-Execution Directive is a file which performs the clean up steps after your test runs. Click **Browse** to locate the file.
- h. Custom properties file is a property file which you specify values used at run-time.

Once you provide these inputs, your type is defined and you can now register a new driver of this type.

6. To register a driver, on the Test Drivers page, select **Custom** from the Register menu.
7. Provide a meaningful name for this Test Driver.
8. Click the magnifying glass for selecting a deployment plan on which the test driver is based. In the resulting dialog box, select the Deployment Plan that should be used for the Custom test driver and click **Select**.
9. Enter a description for the Test Driver and click **Register**.
10. Select a Test Driver in the list then click on the component to drill down to the Deployment Plan details.

## 21.2.4 Setting Up Quotas

**Note:** Ensure the prerequisites are met before setting up quotas (see [Prerequisites](#)).

The Testing as a Service Portal allows you (Testing Administrator) to Setup usage quotas for Testers by specific roles. To set up quotas, follow these steps:

1. Log in to Enterprise Manager Cloud Control using Testing Cloud Self Service Portal Administrator credentials.
2. From the **Setup** menu, select **Cloud**, then select **Testing**. The Testing Administrator page appears.
3. Click **Quota** located in the left panel. The Quotas page appears.
4. Click **Create**.
5. Click the magnifying glass for selecting roles. In the resulting dialog box, select the role to set quotas for and click **Select**.
6. Enter the number of Guest Virtual Machines, CPUs, Memory size, Storage size, and number of Virtual Users to allocate to the selected role and click **Create**.

### 21.2.5 Publishing Test Environments and Test Assets

**Note:** Ensure the prerequisites are met before publishing test environments and test assets (see [Prerequisites](#)).

The Testing as a Service Portal allows you (Testing Administrator) to Setup usage quotas for Testers by specific roles. To publish test environments and test assets, follow these steps:

1. Log in to Enterprise Manager Cloud Control using Testing Cloud Self Service Portal Administrator credentials.
2. From the **Setup** menu, select **Cloud**, then select **Testing**. The Testing Administrator page appears.
3. Click **Publish** located in the left panel. The Publish page appears.
4. Click **Add**.
5. From the **Add** menu in the Select Test Assets, Test Environments, Test Drivers region, select **Test Environment**. In the resulting dialog box, select the environment to publish and click **Select**.
6. From the **Add** menu in the Select Test Assets, Test Environments, Test Drivers region, select **Test Asset**. In the resulting dialog box, select the test asset to publish and click **Select**.
7. From the **Add** menu in the Select Test Assets, Test Environments, Test Drivers region, select **Test Driver**. In the resulting dialog box, select the test driver to publish and click **Select**.
8. Click **Add** in the Select Roles region. In the resulting dialog box, select the role to publish the test asset and environment to and click **Select**.
9. Click **Publish**.

### 21.2.6 Setting Up Chargeback Service

**Note:** Ensure the prerequisites are met before setting up the chargeback service (see [Prerequisites](#)).

Follow this process to set up applications to be tested.

1. Log in to Enterprise Manager Cloud Control using Testing Administrator credentials.

2. From the **Setup** menu, select **Cloud**, then select **Testing**. The Testing Administrator page appears.
3. Click **Chargeback** located in the left panel to define the chargeback service.
4. Click **Configure Charge Plan** to define Extended Charge Plans. Plans may be defined by configuration or usage. If defining plans by usage, charges may be based on universal metrics (CPU, memory, storage) or by metrics specific to the targets. Use extended charge plans to set individual prices. For information regarding chargeback, see *Oracle Enterprise Manager Cloud Administration Guide*.
5. Click **Configure Targets** to assign Charge Plans. Charge plans must be assigned to targets. If targets are grouped into a zone, then a charge plan can be assigned to the zone, and will apply to all targets within that zone. If there are multiple zones, then different charge plans may be assigned to each one. The Self-Service Application (SSA) administrator can determine chargeback pricing by assigning existing charge plans to the different zones.
6. Click **Configure Cost Center** to set up Cost Centers (Optional). Cost Centers may be set up to aggregate costs among groups of users, but are not required. If Cost Centers are set up, the chargeback reports by Cost Center are only available to SSA administrators.

## 21.3 Test Designer Tasks: Creating Test Assets and Test Environments

The following tasks are performed by the Test Designer:

- [Creating Test Assets](#)
- [Creating Test Environments](#)

### 21.3.1 Creating Test Assets

A Test Asset is a collection of test artifacts (scenarios, scripts, tests) that can be executed on a test environment. A test asset defines the actions that need to get executed to reproduce the workload against the test environment. The Testing Cloud Self Service Portal allows the Test Designer to create Test Assets based upon the available (Oracle Load Testing- or Custom - Test Drivers) Testing drivers defined by the Testing Administrator.

The Test Designer should know the application under test and its relevant Key Performance Indicators (KPIs). The Test Designer should also know the business use-cases for the test scenarios being run. Specific responsibilities include:

- The Test Designer should be skilled in Oracle OpenScript and Oracle Load Testing if using the Oracle Load Testing Test Driver. If a Custom Test Driver is used, the Test Designer should be skilled using the corresponding test scripting tool.
- The Test Designer should be skilled as DBA.
- Create Test Assets using the Testing as a Service portal.
- Create Test Environments using the Testing as a Service portal.
- If the Test Designer is required to create environments, the Test Designer must know how to create assemblies.

**Note:** Ensure the prerequisites are met before creating test assets (see [Prerequisites](#)).

Follow this process to create test assets.

1. Log in to Enterprise Manager Cloud Control using Test Designer credentials.

2. From the **Enterprise** menu, select **Cloud**, then select **Self Service Portal**. The Self Service page appears.
3. Click **My Tests**. The My Library tab appears.
4. Click **Test Assets** to define the test assets to be added to the library.
5. From the **Create** menu, select the type of asset you want to add, either Oracle Load Testing or Custom. Test Asset Properties page appears.
6. For the Custom asset type, select the type of Custom asset.  
Enter a name and description for the test asset.
7. Click the **Application** drop-down list and select the application to be tested from the list.
8. Enter any tags to use as keywords to further identify this test asset. These tags make test assets easier to identify and to understand the purpose of the asset.
9. Click **Browse** and select the asset file you want to upload. The file type for Oracle Load Testing is SCNZIP. The file type for Custom can be whatever file type understood by the test-driver involved.
10. Select or clear the **Parameters Exposed for Testing** options. These are parameters that are specified as part of the Oracle Load Testing scenario. Parameters selected in this dialog box will be exposed within TaaS and the parameters can be changed for the TaaS test runs.
11. Click **OK**. The Test Asset is added to the library. Click on a test asset name to drill down to the Test Asset Details page.

### 21.3.2 Creating Test Environments

A Test Environment is a collection of entities that represents the entire application stack, thus an environment consists of the following elements:

- **Application under test (AuT):** Contains the software bits and application data. For example, for a deployed Oracle E-Business Suite application, end users can modify settings, configurations, and parameters of the application.
- **Infrastructure under test (IaT):** These are entities the Application under test depends upon, for example, the cloud setup/infrastructure in a cloud deployment. Typically an infrastructure is something created for consumption of end users by administrators.

**Note:** Ensure the prerequisites are met before creating test environments (see [Prerequisites](#)).

Follow this process to create test environments.

1. Log in to Enterprise Manager Cloud Control using Testing Cloud Self Service Portal Test Designer credentials.
2. From the **Enterprise** menu, select **Cloud**, then select **Self Service Portal**. The Self Service page appears.
3. Click **My Tests**. The My Library tab appears.
4. Click **Test Environments** to define the test environments to be added to the library.
5. Click **Create**. The Test Environment Properties page appears.
6. Enter a name and description for the test environment.

7. Click the **Application** drop-down list and select the application to be tested from the list.
8. Enter any tags to use as keywords to further identify this test environment. These tags make test environments easier to identify and to understand the purpose of the environment.
9. Click the magnifying glass icon and select the **Deployment Plan** to use for this test environment. The Deployment plan must be created on a zone to which the Test Designer has privileges on Oracle Virtual Machine and Operator. For information regarding deployment plans, see *Oracle Enterprise Manager Lifecycle Management Administrator's Guide*.
10. Click the **Select web-application** drop-down and select the appliance that will act as the application or web server when the environment is deployed in the cloud.
11. Click **Save**. The test environment is added to the library. Click on a test environment name to drill down to the Test Environment Details page. Click on a deployment plan name to drill down to the Deployment Plan Details page.

## 21.4 Using the Testing Environment

Once the Test Administrator and Test Designer have performed their respective tasks, the Tester can now start testing.

The Tester should know the application under test and its relevant Key Performance Indicators (KPIs). The Tester should also know the business use-cases for the test scenarios being run. Specific responsibilities include:

- Running the tests by using the test assets and environments created by the Test Designer and assigned to them by the Testing Administrator.
- Reserving resources from the cloud, cancel upcoming reservations, update configurations of resources that have been reserved (CPU, memory, and so on).
- Editing test scenarios set up by the Test Designer to specify run time test parameters and schedule the tests for execution.
- Performing the test run analysis to compare various metrics that are relevant to the application and business use-case.

The following tasks are performed by the Tester.

- [Creating Tests](#)
- [Creating and Running Trials](#) (Create Trial)
- [Monitoring On-Going Trials](#)
- [Comparing Trial Results](#)
- [Browsing the Testing Home Page](#)
- [Share Tests](#)
- [Viewing Activity and Usage Information](#)
- [Viewing Deployment Information](#) (Browse My Deployments)

### 21.4.1 Verifying Available Test Assets and Environments

The Testing Cloud Self Service Portal allows Testers to verify which test assets and test environments have been assigned to them by the Testing Administrator.



**Note:** Ensure the prerequisites are met before verifying available test assets and environments (see [Prerequisites](#)).

Follow this process to verify test assets and test environments assigned to your role.

1. Log in to Enterprise Manager Cloud Control using Tester credentials. You are logged into Self Service Portal.
2. Click **My Tests**. The Testing Cloud Self Service Portal Home tab appears.
3. Click the **My Library** tab. The My Library tab shows the test environments and test assets published to you by the Testing Administrator.
4. Click **Test Environments**.
5. Click on an environment to view the details about the environment. Click on the component to drill down to the Assembly Structure details.
6. Click **Test Assets**.
7. Click on a test asset to view the details about the asset.

## 21.4.2 Creating Tests

TaaS allows Testers to create tests from test assets and environments to perform testing activities. Testers create tests using test assets and environments that have been assigned to them by the Testing Administrator.

**Note:** Ensure the prerequisites are met before creating Tests (see [Prerequisites](#)).

Follow this process to create Tests.

1. Log in to Enterprise Manager Cloud Control using Testing Cloud Self Service Portal Tester credentials. You are logged into Self Service Portal.
2. Click **My Tests** in the Self Service Portal.
3. Click the **My Tests** tab in the Testing Cloud Self Service Portal.
4. From the **Create** menu, select the type of test you want to create, either **Oracle Load Testing** or **Custom**. The Create Test page appears.
5. For the Custom test, select the Test Driver type.  
Enter a name and description for the test.
6. Click the magnifying glass icon, select the test asset to be used in this test, and click **Select**.
7. Select the test asset, review the details, and click **Continue**.
8. Select the Source and click the magnifying glass icon and select the test environment to be used in this test.
9. Select the test environment and click **Select**. The Mapping URLs settings control shows all the URLs found in the load-testing scenario contained in the asset.
10. From the options column, select **From Appliance**. You will see an appliance (a virtual machine) that will be used to replace the "source URL" that was in the load testing scenario.
11. Click the **Test Driver** drop-down and select the test driver to use to run the test in the cloud.
12. Click **OK**. The test is added to the library that specifies an asset, an environment, and a test-driver. Click on a test name to drill down to the details page.



### 21.4.3 Creating and Running Trials

The Testing Cloud Self Service Portal allows Testers to create and run trials from the test instances that have been created in the portal. The Tester runs trials on the test asset and environment, monitors ongoing trials, and analyzes the results of completed trials.

**Note:** Ensure the prerequisites are met before creating and running trials (see [Prerequisites](#)).

Follow this process to create and run trials.

1. Log in to Enterprise Manager Cloud Control using Testing Cloud Self Service Portal Tester credentials. You are logged into Self Service Portal.
2. Click **My Tests** in the Self Service Portal.
3. Click the **My Tests** tab in the Testing Cloud Self Service Portal.
4. Select a test to run for this trial.
5. Click **Create** in the **Trials** section. The Create Trial page appears.
6. Enter a name and description for the trial.
7. Specify the Application Under Test instance details.
  - Creating new instance: Users can create new virtual machines on the cloud for the application under test.
  - Re-using an existing instance: If users have created instances for this test or this application before, they can choose to re-use the same set of virtual machines instead of creating new ones.

Re-using an existing instance is faster.
8. Enter a password of your choice for the Root Password and reenter to confirm it if creating a new instance.

---

**Important:** This value is not retrievable. DO NOT FORGET THIS PASSWORD.

---

9. Specify the **Stop After** and **Delete After** settings and click **Continue**.
10. Enter the number of virtual users to run for this trial in the **Trial Virtual Users** column of each script.

You can set these values **ONLY** if the Test Designer (who created the asset) has chosen to allow users to override the values. Otherwise, the application uses the information available in the asset and these controls are grayed out.

11. Specify the **Ramp Up Rate** settings for this trial.

The Ramp Up Rate is the rate by which the load increases during the load test. You can set these values **ONLY** if the Test Designer (who created the asset) has chosen to allow users to override the values. Otherwise, the application uses the information available in the asset and these controls are grayed out.

12. Set the **Test End Conditions** settings for the trial and click **Continue**.

You can set these values **ONLY** if the Test Designer (who created the asset) has chosen to allow users to override the values. Otherwise, the application uses the information available in the asset and these controls are grayed out.

If you do not set the test end conditions, the test runs until it is stopped manually.

13. Select the test asset, review the details, and click **Continue**.
14. Enter an instance name for the test driver.  
  
You can create a new virtual machine for the test driver or re-use an existing instance. If the existing instance is being used by another trial, it will not be available.
15. Enter a password of your choice for the Master Password and reenter to confirm it. This is the password Testing Cloud Self Service Portal will set for the Oracle Load testing Virtual Machine and all accounts for Oracle Load Testing.
16. Specify the **Stop After** and **Delete After** settings and click **Continue**.
17. Specify the **Schedule** for starting the trial run and click **OK**. A new trial will be created and submitted for execution. The initial execution of a trial will take some time to deploy the Virtual Machines, copy necessary software, and start the test.

#### 21.4.4 Monitoring On-Going Trials

The Testing Cloud Self Service Portal allows Testers to monitor running trials and create charts of the key performance indicators for the application-under-test.

**Note:** Ensure the prerequisites are met before monitoring on-going trials (see [Prerequisites](#)).

Follow this process to monitor on-going tests.

1. Log in to Enterprise Manager Cloud Control using Testing Cloud Self Service Portal Tester credentials. You are logged into Self Service Portal.
2. Click **My Tests** in the Self Service Portal.
3. Click the **My Tests** tab in the Testing Cloud Self Service Portal.
4. Select the test to view the trial details. The trial status information appears in the Trials pane.
5. Click on the trial name link in the Trials pane to navigate to the trial-home page.

You can monitor the progress of the trial by clicking **Status**. If the status of the trial is In Progress, click the trial name which will take you to the trial home and the remaining trial information.

6. Expand the metric nodes and select metrics or counters to plot them in the chart.

#### 21.4.5 Comparing Trial Results

The Testing Cloud Self Service Portal allows Testers to compare trial results and create charts of the key performance indicators for the application-under-test.

**Note:** Ensure the prerequisites are met before comparing trial results (see [Prerequisites](#)).

Follow this process to compare trial results.

1. Log in to Enterprise Manager Cloud Control using Testing Cloud Self Service Portal Tester credentials. You are logged into Self Service Portal.
2. Click **My Tests** in the Self Service Portal.
3. Click the **My Tests** tab in the Testing Cloud Self Service Portal.

4. Select the test to view the trial details. The trial status information appears in the Trials pane.
5. Select two trials of the same test that have been previously run.  
To select 2 trials of the same test, select one trial, then click Shift/Control to select the other trial. The trials must belong to the same test.
6. Click **Compare Results**. You can compare the metrics of the two trials to view the differences.  
When in this mode, selecting any metric will chart two lines - one for each of the selected trials. You can also compare the Oracle Load Testing session summary report side-by-side.

### 21.4.6 Share Tests

To avoid duplication of effort, you can share tests with other users. Note that the person who created and shared the tests will be charged for any resource usage associated with the shared tests.

To share tests, perform the following steps:

1. Log in to Enterprise Manager Cloud Control using Testing Cloud Self Service Portal Tester credentials. You are logged into Self Service Portal.
2. Click **My Tests** in the Self Service Portal.
3. Click the **My Tests** tab.
4. Select a test from the list of Tests on the My Tests page.
5. Click **Share**.
6. Select a user to share the Test with.
7. Click **Share**.
8. Click **Log Out**.

To verify that the test is shared:

1. Log in as the user whom you shared the test with.
2. Verify that you have access to the shared test.

### 21.4.7 Viewing Activity and Usage Information

The Testing Cloud Self Service Portal allows Testers to view summary information for test assets and test environments available to you, your quota as defined by the Testing Administrator for your role, your current utilization of hardware resources, and submitted and scheduled trials.

**Note:** Ensure the prerequisites are met before viewing activity and usage information (see [Prerequisites](#)).

Follow this process to view activity and usage information.

1. Log in to Enterprise Manager Cloud Control using Testing Cloud Self Service Portal Tester credentials. You are logged into Self Service Portal.
2. Click **My Tests** in the Self Service Portal.
3. Click the **Home** tab in the Testing Cloud Self Service Portal. This tab shows summary information for test assets and test environments available to you, your

quota as defined by the Testing Administrator for your role, your current utilization of hardware resources, and submitted and scheduled trials.

### 21.4.8 Viewing Deployment Information

The Testing Cloud Self Service Portal allows Testers to create tests from test assets and test environments to perform testing activities using the Testing Cloud Self Service Portal. Testers create tests using test asset and environment that have been assigned to them by the Testing Administrator. The Tester runs trials on the test asset and environment, monitors ongoing trials, and analyzes the results of completed trials.

**Note:** Ensure the prerequisites are met before viewing deployment information (see [Prerequisites](#)).

Follow this process to view activity and usage information.

1. Log in to Enterprise Manager Cloud Control using Testing Cloud Self Service Portal Tester credentials. You are logged into Self Service Portal.
2. Click **My Tests** in the Self Service Portal.
3. Click the **My Deployments** tab in the Testing Cloud Self Service Portal. This tab shows a table of all environments and test-drivers deployed for you.
4. Select a test environment or test driver instance. The details pane shows the appliances within the selected test environment or test driver instance.
5. Click on an appliance within the selected test environment or test driver instance to navigate to the details view of the appliance.

### 21.4.9 Browsing the Testing Home Page

The users HOME tab displays a summary of what the user has done, including resources consumed for testing and tests created by the user. From the HOME tab, users can quickly access previously executed test trials in order to compare results or to do more in depth post-run analysis.

The Testing home page provides information regarding the statistics of ongoing testing. General statistics are provided, as well as test assets and environments, ongoing test and trials, Oracle Load Testing, and the Infrastructure. It also provides an overview of the servers and infrastructure resources consumed by the user, like the number of servers provisioned, number of CPUs, as well as RAM and Storage.

When you log in as Test Designer, you only see the My Library tab. When you log in as Test Administrator, you see the Home tab. When you log in as a Super Administrator, you see both the Home and My Library tabs.

The sections on the Home tab are:

- General - Total counts of various artifacts available in the system, for example, its “total tests present in system”, “total trials ran across all trials”, and so on.
- Test Assets/Environments - Accumulation graph of test assets, environments, and machines in use in the Testing Cloud. Count and trend data are included. This is the Test Administrator’s overview of the entire TaaS portal, regardless of Testers.
- Tests and Trials - Daily totals of the number of tests created or daily totals of the number of trial runs.
- Infrastructure - Daily consumption of infrastructure resources: CPU, Memory, and Disk space storage.

- Oracle Load Testing (OLT) - Daily total of the number of OLT Virtual Users (VU) that were run. Oracle Load Testing uses OpenScript scripts to simulate users accessing an application under test.

The vertical axis shows daily usage totals based upon the View selection in each of the regions.

The My Library tab provides details on both Test Assets and Test Environments.

1. Log in to Enterprise manager using Tester credentials.
2. Click **My Tests** radio button.
3. Click the **Home** tab.

You will see a summary of test assets and test environments available to you, your quota as defined by the Test Administrator for your role, your current utilization of hardware resources, and the submitted and scheduled trials.



# Part VI

---

## Using Consolidation Planner and Chargeback

This section describes the Consolidation Planner and Chargeback features in Enterprise Manager.

It contains the following chapters:

- [Chapter 23, "Using Consolidation Planner"](#)
- [Chapter 22, "Chargeback Administration"](#)





---

## Chargeback Administration

This chapter describes the Chargeback application and details the administrative tasks necessary to set up and run the application. It then presents various ways to access and use the information Chargeback collects and aggregates.

The chapter contains the following sections:

- [Overview of Chargeback](#)
- [Setting Up Chargeback](#)
- [Accessing Chargeback Information](#)

### 22.1 Overview of Chargeback

This section gives a brief overview of the most important aspects of Chargeback. It covers the following topics:

- [Why Implement Chargeback](#)
- [Enterprise Manager Chargeback](#)
- [Understanding Targets and Metrics](#)
- [About Chargeback, Self Service Applications, and Zones](#)

#### 22.1.1 Why Implement Chargeback

The purpose of Chargeback is to gather data on resource use, allocate charges for the use of these resources, and present the results in a comprehensible format.

Chargeback is used to allocate the costs of IT resources to the people or organizations who consume them. While it can be applied in situations where IT resources are dedicated, it is particularly relevant in situations where resources are shared, as without some way to meter and charge for consumption there will be a tendency for users to use more resources than they need. This problem is exacerbated in cloud environments where users are able to provision their own resources using self-service.

Consider the following primary use cases:

- Increasingly, organizations are adopting consolidation platforms such as Oracle Exadata and Oracle Exalogic, where there are multiple applications running on a shared hardware platform. There needs to be a way to meter the resources consumed by each application so charges can be applied accordingly. This, in turn, helps to measure the business value of IT investments so that priority can be given to the applications with the greatest return on investment (ROI).

- With the popularity of cloud computing on the rise, where users have the ability to self-service provision resources such as databases and virtual machines, it is highly desirable to be able to meter the resources each cloud consumer uses so that they can be charged based on their consumption. Cloud consumers likewise benefit by rationing their consumption to that which provides optimal business value, thus contributing to the overall ROI of the organization's IT investment.

To sum up, adopting a Chargeback (or Showback) model can deliver significant benefits to both IT and Line of Business users:

- Chargeback benefits users by placing them in control of their IT costs.
- Chargeback benefits IT by driving up utilization rates while reducing the number of physical servers they need to manage.

### 22.1.2 Enterprise Manager Chargeback

*Chargeback*, as the name implies, is a tool of accountability. The application's primary uses can generally be described as follows:

- Provide resource usage metering by aggregating and normalizing the enormous amount of metric data Enterprise Manager collects.
- Provide IT a means to "charge" a currency amount to internal organizations that use resources.
- Provide internal organizations and users with reports detailing their consumption and charges.

Chargeback has three basic metrics against which to compute resource consumption: CPU usage, and memory and storage allocation. These metrics comprise a universal charge plan that can be applied to any target type configured for Chargeback.

While CPU, memory and storage can be used for Chargeback across a variety of target types, there may be situations where target-specific charges are required. In this instance, an extended charge plan can be used. The extended charge plan provides greater flexibility to Chargeback administrators, enabling them to:

- Define target type-specific charges
- Define fixed, configuration, and usage-based rates
- Override or adjust universal plan rates

Chargeback provides an out-of-box extended plan that you can use as a template for creating your own extended plans tailored to your environment. The shipped extended plan serves an Oracle VM Guest target type where charge items and applicable rates vary depending on machine size. The sample plan derives from the Cloud Infrastructure self-service portal (from the **Setup** menu, select **Cloud**, then select **Infrastructure**).

When charges for the consumption of resources are calculated, these charges must be assigned somewhere. In Chargeback, the assignment of these costs is to a cost center. Cost centers are typically organized in a business hierarchy and may correspond to business units—sales, engineering, human resources, and so forth. The cost center hierarchy is defined within the Chargeback application and is used for rollup and drill-down purposes when reporting on usage and charges.

### 22.1.3 Understanding Targets and Metrics

Enterprise Manager collects metric and configuration information on targets and stores it in the Management Repository. These collections occur with varying

frequencies. When a target is added for Chargeback, the application extracts the data from the repository and transforms it for use in Chargeback-dedicated tables. This process of extracting and transforming data is the Chargeback data collection process, which is scheduled to run every 24 hours. Chargeback maintains its own data collection tables because it needs to perform various transformations on the data, and to provide a historical data reference.

Chargeback uses information that Enterprise Manager collects as a basis for metering. Chargeback metrics can be configuration-based, usage-based, or availability-based.

- Configuration-based metrics refer to things such as number of CPUs or database version; these metrics are considered static and thus have a daily aggregation.
- Usage-based metrics refer to things such as CPU utilization or memory consumption; these metrics have an hourly aggregation.
- Availability-based metrics measure the time a target is available, otherwise known as a target's uptime. It is the number of minutes in a given hour a target is considered available. The charge is prorated for downtime.

Charges can also be applied at a fixed rate.

Target types for which charge rates can be defined for specific metrics include host, virtual machine, WebLogic Server, database instance, and pluggable database (PDB). Composite target types can also be added to Chargeback. These include Cluster Database, Oracle WebLogic Cluster, Oracle WebLogic Domain, Oracle VM Zone, Oracle VM Server Pool, and so forth. Adding groups and systems automatically includes all members that are valid Chargeback targets. After a target has been added to Chargeback, a collection job will run once a day against target tables in Enterprise Manager to extract information for use by Chargeback.

---

**Note:** Changes in Chargeback setup are not immediately reflected in reports. For example, if you change a cost center assignment, it may take up to 24 hours (until the next data collection job runs) for reports to show the change. If you have SYSMAN Super Administrator privileges, however, you can initiate data collection on-demand from the **Targets** tab. To do this, select **On-demand data collection** from the **Action** menu.

---

Chargeback targets can be added in dedicated or shared mode.

- Dedicated means a consumer or a group of consumers are using a target where charges accrue against a single cost center.
- Shared means consumers belong to different groups; charges accrue against different cost centers.

All supported target types can be added as dedicated. Database and WebLogic Server also can be configured as shared.

- A shared database makes multiple database services available to sharing groups, which use different service names; charges are by the service.
- A shared WebLogic Server makes multiple J2EE applications available to sharing groups; charges are by the application.

## 22.1.4 About Chargeback, Self Service Applications, and Zones

A key component of cloud computing is availability of self service applications. Enterprise Manager includes SSAs for Oracle VM, database, and middleware.

1. To access the respective portal, from the **Enterprise** menu, select **Cloud**, then select **Self Service Portal**.
2. On the Self Service Portal page, select the respective radio button.
  - **My Servers**—to request and manage servers in an Oracle VM zone.
  - **My Databases**—to request and manage database services in a PaaS Infrastructure zone.
  - **My Middleware**—to request and manage middleware services in a PaaS Infrastructure zone.
  - **My Tests**—to request and manage a virtual testing process using cloud and other resources.

Note that each self service portal home page includes a **Chargeback** tab that shows the charges on the SSA targets requested by the SSA user for the current SSA service. It also shows the total charge for the SSA user in the specific service type, which can be on a single PaaS Infrastructure zone or multiple PaaS Infrastructure zones for the database service.

While performing cloud setup, you can link directly into Chargeback to configure the respective zone. Note that zone setup is a prerequisite.

1. From the **Setup** menu, select **Cloud**, then select as follows:
  - **Infrastructure**—to set up VM types.
  - **Database**—to set up database resource pools.
  - **Middleware**—to set up middleware resource pools.
2. Each respective home page has a **Chargeback** selection on the left. Click it to display an informational page on Chargeback setup that includes links into Chargeback to perform setup.

Chargeback supports three zones:

- PaaS Infrastructure (includes DBaaS and MWaaS)
- Oracle Exalogic (depends on Exalogic zone support in the SSA plug-in)
- Oracle VM

Within Chargeback, the PaaS Infrastructure zone hierarchy appears on the **Targets** tab under the Cloud category and the other two zones under the Servers, Storage and Network category, as follows:

---

**Note:** A PaaS Infrastructure zone can contain a database instance, a RAC instance, a Composite application, or any combination of these three, as represented below.

---

| Zones               | Zone Hierarchy                                                                                                                                                                                                                                                              |
|---------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PaaS Infrastructure | PaaS Infrastructure Zone<br>Single Instance Database<br>Cluster Database<br>Database Instance1<br>Database Instance2<br>Oracle Composite Application<br>WebLogic Domain<br>WebLogic Cluster<br>Oracle WebLogic Server1<br>Oracle WebLogic Server2<br>Oracle WebLogic Server |
| Oracle Exalogic     | Oracle Exalogic Zone<br>Oracle Assembly Instance<br>Oracle Tier Instance<br>Oracle VM Guest1<br>Oracle VM Guest2                                                                                                                                                            |
| Oracle VM           | Oracle VM Zone<br>Oracle VM Server Pool<br>Oracle VM Guest1<br>Oracle VM Guest2                                                                                                                                                                                             |

The following characteristics apply to all target type zones:

- Adding a zone for Chargeback automatically includes all child instances.
- Removing a zone from Chargeback automatically removes all child instances.
- A child instance with no explicitly assigned charge plan inherits the plan of its nearest ancestor above it in the hierarchy.
- A child instance with no explicitly assigned cost center inherits the cost center of its nearest ancestor above it in the hierarchy.
- Querying on charges at a given level within the hierarchy will include all charges for the levels below (either directly or indirectly).
- New members discovered or added to a zone are automatically recognized by Chargeback at the next data collection cycle.

The SSA user is automatically associated with the SSA target as the cost center, when the SSA target is added for Chargeback. See [Section 22.2.4](#) to learn about assigning cost centers. When assigning charge plans, a plan assigned to a zone applies to all members within the zone that do not otherwise have a plan assigned. See [Section 22.2.5](#) to learn about assigning charge plans.

## 22.2 Setting Up Chargeback

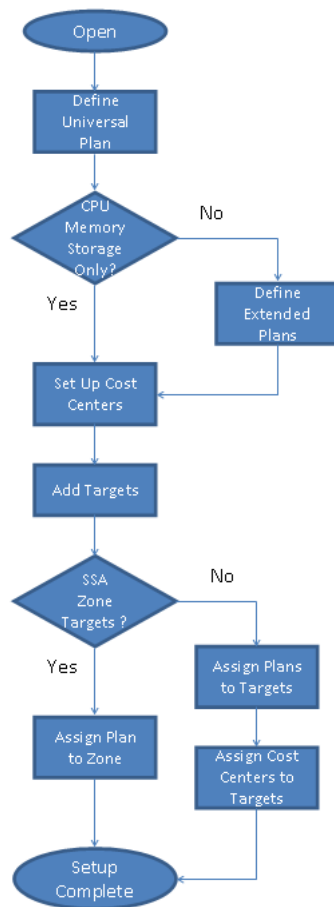
This section describes the following tasks, which must be completed to implement and support Chargeback services:

- [Working with Charge Plans](#)

- [Creating Cost Centers](#)
- [Adding Targets for Chargeback](#)
- [Assigning Cost Centers to Targets](#)
- [Assigning Charge Plans to Targets](#)
- [Configuring Chargeback Settings](#)

The person performing these tasks must have the necessary privileges. These privileges can be granted using the EM\_CBA\_ADMIN role. [Figure 22–1](#) shows a simple workflow of steps required to set up Chargeback.

**Figure 22–1 Chargeback Setup Workflow**



Following initial setup, you perform tasks on a regular basis, such as:

- Adding new targets for Chargeback
- Updating charge plans
- Assigning charge plans and cost centers

Other events happen automatically:

- SSA target charges are rolled up to higher-level cost centers
- New targets that are members of a composite target (a VM Guest in a VM zone, for example) are discovered and added for Chargeback

The daily data collection job is responsible for capturing these changes, in addition to collecting metric and configuration data.

## 22.2.1 Working with Charge Plans

As described in the Overview, there are two types of charge plan, the universal charge plan and an extended charge plan.

The universal charge plan establishes rates for three basic metrics (CPU, memory, and storage). For example, charge a dollar a day for a CPU (generic), two dollars a week per GB of memory, and a dollar a week per GB of storage. Optionally, different rates can be applied per CPU architecture. For example, charge different rates for a PA-RISC, a SPARC, and an x86. You can create multiple revisions of the universal charge plan for use in the future.

An extended charge plan enhances the universal plan to include target-specific metrics. You may want to implement charges that relate to specific characteristics of a target. An extended charge plan affords that flexibility. The target type determines the items for which rates can be charged.

In defining an appropriate charge plan for use with a given target type, you have to decide on a case-by-case basis whether the three basic metrics of CPU, memory, and storage are sufficient to charge for resource use. If the answer is no, then you have to define an extended plan for the target type.

In any case, you want to define the universal charge plan for general use. You can then indicate in an extended charge plan how the universal plan rates apply.

An extended charge plan consists of groups of target-type specific charge items. Target-specific charge templates are available for the following target types:

- Host
- Oracle WebLogic Server (dedicated and shared)
- Oracle VM Guest
- Database Instance (dedicated and shared)
- Pluggable Database (PDB)

With an extended plan you can base charges on target-specific charge items. For example, you may want an additional charge for a host running Windows OS, or a database with the Partitioning option installed. In addition to configuration- and usage-based charge items, you can include a fixed charge. For example, charge \$100 a month for each host, regardless of consumption.

You may find that even with an extended plan, the charge items available for a given target type are insufficient for your particular needs. In this case, you can introduce custom charge items through use of CLI verbs. For additional information, see Chapter 29, "Chargeback and Metering EMCLI Verbs."

This section covers the following topics:

- [Charge Plans and Effective Dates](#)
- [Using Conditions in Charge Plans](#)
- [Creating and Revising the Universal Charge Plan](#)
- [Creating an Extended Charge Plan](#)
- [Revising Extended Charge Plans](#)

### 22.2.1.1 Charge Plans and Effective Dates

When you create a charge plan (universal or extended), you establish the plan's effective date, which is always the first of a month. The plan revision remains in effect until you create a new plan revision whose effective date determines the previous plan revision's termination.

Say, for example, you want to set up a charge plan for a calendar year, in which a new plan revision goes into effect each quarter. You create an initial plan revision that goes into effect January 1. The plan shows an effective date of Jan. 1, yyyy - Onward. You create a plan revision that goes into effect April 1. The initial plan revision now shows an effective date of Jan. 1, yyyy - Mar. 31, yyyy, and the new plan revision, Apr. 1, yyyy - Onward. You create two more plan revisions, one with an effective date of July 1 and the other with an effective date of October 1. You now have four plan revisions with effective dates as follows:

- Jan. 1, yyyy - Mar. 31, yyyy
- Apr. 1, yyyy - Jun. 30, yyyy
- Jul. 1, yyyy - Sep. 30, yyyy
- Oct. 1, yyyy - Onward

Only current and future plan revisions appear in the Charge Plans navigation pane. You can retrieve expired plan revisions by selecting **Historical Revisions** from the **View** menu.

### 22.2.1.2 Using Conditions in Charge Plans

Conditions introduce considerable flexibility when creating extended charge plans. You can create variations known as configurations that allow you to determine a set of rates to charge based on target configuration. You base configurations on a property or attribute of a target, such as the size of a virtual machine. The sample extended plan, for example, sets conditions on a VM Guest target type, based on the machine size of the virtual machine, allowing for different sets of charges to apply depending on machine size. There is also an implied default configuration, which is the charge plan itself if no configurations are set up, or the charge plan of record if there is no match for a configuration condition. The sample extended plan thus has four configurations, one each for machine sizes small, medium, and large, plus the default "Other."

Similarly, you can place a condition on an individual charge item that determines if a charge will be applied. For example a condition set on Partitioning means charge only if partitioning is used.

Use the Like condition operator with string-based target type and charge items to match a range of string values. For example, set a Like condition on a database instance to match on all 11g releases (Like 11%). Or use it as a wildcard (Like %) to match on any string of characters (or none). Use an underscore (Like \_) to match any single character. Use backslash to escape a wildcard character (Like \%). Use a double backslash to escape the escape character (Like \\%).

When you have multiple configurations, you have to set the order of precedence to determine how conditions are applied. Suppose you have the following sequence of configurations based on the host operating system:

- OS Like %Linux%
- OS Like %Linux 6%
- Default



The second condition (Linux 6) will never be tested because the first will always take precedence when the OS is any version of Linux. Thus, you have to make the test for Linux 6 first to base charges specifically on the Linux 6 OS.

Contrast the Like operator with the Equals operator, where the match must be exact, and there is no wildcard recognition. So, for example, Equals 11.2% does not match 11.2.0 or 11.2.1. You must specify Equals 11.2.0 or Equals 11.2.1. Note also that all string matching is case-sensitive. Equals UPPER is not the same as Equals upper. There is also a Not Equals operator with which you can exclude specific string values.

Consider these other condition examples:

- For a shared database, set a condition to base the charge on the nodes of a RAC.
- For a dedicated database instance, use the SQL Executions charge item to set a charge rate on SQL statement executions per second.

Conditions give you flexibility. For example, you may want to create many charge plans with no conditions and assign individually to targets. Or, you can create a single charge plan with many conditions, and assign to many targets.

### 22.2.1.3 Creating and Revising the Universal Charge Plan

Set up the universal charge plan as follows:

1. How to proceed depends on the revision date:
  - If the effective date under Universal Charge Plan starts in the past, select the revision in the navigation pane, then select **Revision** from the **Create** menu.
  - If the effective date under Universal Charge Plan is current or in the future, select the revision in the navigation pane and click **Set Rates**.
2. In the dialog that opens, select the effective date of the revision. Note that a change to rates currently in effect is retroactive to the beginning of the cycle. Click **OK**.
3. In the Charge Plan editor, set the rates and charge frequency for the three basic metrics.

Select a row, enter the rate and select the frequency from the drop-down menu.

The rates set here are effective from the beginning of the current month, and have no impact on any previous month.

4. Make additional entries for different CPU architectures, as required.
  - a. Click the **Add** button.
  - b. Select a CPU from the drop-down menu and click **OK**.
  - c. In the new row added to the table, enter the rate and select the charge frequency from the drop-down menu.
5. When done, click the Save button in the upper right hand corner to complete the universal charge plan revision.

### 22.2.1.4 Creating an Extended Charge Plan

Set up an extended charge plan as follows:

1. From the **Create** menu, select **Plan**. The Charge Plan editor opens.
2. Enter a name for the plan. The name must be unique but the plan can have multiple, date-based revisions.
3. Click **Add**. The Add Target Types dialog opens.

4. Select from the list of supported target types. As a plan can serve more than one target type, you can multiselect in the dialog. Click **OK**.

Selected target types appear in a table with one configuration (default) indicated.

5. Click **Set Up Configurations** to open the Set Up Configurations dialog, then click **Add** to open the Add Configuration dialog.

Add a condition to set up a configuration for the target type selected. The list of conditions available derives from the target type. Select from the list, choose an operator, and set the condition value. You can specify a value or click the search icon to select the condition value from a list, or to search for some other value.

You might, for example, set a condition of machine architecture for a host target type, where there are two options: sparc sun4u and AMD64. You can then use the **Move Up** and **Move Down** buttons to establish the order of precedence in matching conditions, if applicable.

Click **OK** when done with configuration setup.

As [Figure 22–2](#) shows, the charge plan table now displays separate columns for conditional configurations and the default configuration called Other.

**Figure 22–2 Setting Up a Target Configuration**

| Host                                                                                                                   |                                               |                  |       |
|------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------|------------------|-------|
| <span>+</span> Add Item... <span></span> Edit Item... <span>✕</span> Remove Item <span></span> Setup Configurations... |                                               |                  |       |
| Item                                                                                                                   | Charge by Configuration: Machine Architecture |                  |       |
|                                                                                                                        | AMD64                                         | sparc sun4u      | Other |
| ▼ Universal Rate Adjustments                                                                                           |                                               |                  |       |
| ▶ CPU Rate Factor                                                                                                      | 1x <span></span>                              | 1x <span></span> | 1x    |
| ▶ Memory Rate Factor                                                                                                   | 1x <span></span>                              | 1x <span></span> | 1x    |
| ▶ Storage Rate Factor                                                                                                  | 1x <span></span>                              | 1x <span></span> | 1x    |

Note that you can set only a single condition for a configuration. Suppose you want to change the condition from machine size to high availability for VM Guest. To do so, open the Set Up Configurations dialog and remove the machine size condition items. When you click **Add**, the high availability condition item is now available.

6. Click **Add Item** to open the Add Item dialog. Use this dialog to add charge items based on target type for billing purposes.

Select a charge item from the drop-down list. The list of items derives from the target type. For certain items, you can select the charge type from a drop-down list—flat rate, per unit, and so forth. You can also set conditions, the same as for configurations in Step 5. For a shared database, charges are per service. For a shared WebLogic Server, charges are per application. Make your selection and click **OK**.

The charge plan table now displays a row for the charge item just added. A column appears for each configuration.

7. In the charge item row, specify the rate and select the frequency for each configuration. Say, for example, you added a CPU count charge item for target type host, for which you want to charge a \$1 a week per unit for the SPARC architecture, \$.50 a week per unit for the AMD architecture, and \$.75 per unit a week for all other architectures. Make the appropriate entries in the respective columns as displayed in [Figure 22–3](#).

**Figure 22–3 Setting a Condition on a Charge Item**

| Host                                                                                                                                 |                                               |                |                |
|--------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------|----------------|----------------|
| <a href="#">+ Add Item...</a> <a href="#">✎ Edit Item...</a> <a href="#">✖ Remove Item</a> <a href="#">⚙ Setup Configurations...</a> |                                               |                |                |
| Item                                                                                                                                 | Charge by Configuration: Machine Architecture |                |                |
|                                                                                                                                      | AMD64                                         | sparc sun4u    | Other          |
| CPU Count                                                                                                                            | \$ 0.50 / Week                                | \$ 1.00 / Week | \$ 0.75 / Week |
| ▽ Universal Rate Adjustments                                                                                                         |                                               |                |                |
| ▶ CPU Rate Factor                                                                                                                    | 1x                                            | 1x             | 1x             |
| ▶ Memory Rate Factor                                                                                                                 | 1x                                            | 1x             | 1x             |
| ▶ Storage Rate Factor                                                                                                                | 1x                                            | 1x             | 1x             |

- Repeat Steps 6 and 7 to include additional charge items for the target type.
- Optional. Make rate adjustments to the universal charge plan basic metrics for this extended charge plan. Make adjustments by specifying a factor of the base rate, where 1, the default, represents the prevailing base rate. So, for example, enter 1.5 to increase the rate by 50 percent; enter .5 to decrease the rate by 50 percent. If you want to ignore a universal plan rate altogether, enter 0.

Expand the rate factor for a metric to see what the actual charge is. For example, if the rate adjustment for memory allocation is 1.5 and the rate is \$2 a month, the adjusted rate will be \$3 a month.

The metrics available for a target type determine which rate adjustments are possible. For example, you cannot make rate adjustments for a shared WebLogic Server.

Suppose, given the example, you want to adjust the CPU rate to 75 percent for the AMD architecture and to 125 percent for the SPARC, leaving the rate as is for the rest. Make appropriate entries as shown in [Figure 22–4](#).

**Figure 22–4 Extended Plan Showing Customized Charge Items**

| Host                                                                                                                                 |                                               |                    |                    |
|--------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------|--------------------|--------------------|
| <a href="#">+ Add Item...</a> <a href="#">✎ Edit Item...</a> <a href="#">✖ Remove Item</a> <a href="#">⚙ Setup Configurations...</a> |                                               |                    |                    |
| Item                                                                                                                                 | Charge by Configuration: Machine Architecture |                    |                    |
|                                                                                                                                      | AMD64                                         | sparc sun4u        | Other              |
| CPU Count                                                                                                                            | \$ 0.50 / Week                                | \$ 1.00 / Week     | \$ 0.75 / Week     |
| ▽ Universal Rate Adjustments                                                                                                         |                                               |                    |                    |
| ☑ CPU Rate Factor                                                                                                                    | 0.75x                                         | 1.25x              | 1x                 |
| CPU Usage                                                                                                                            | \$1.00 / CPU / Day                            | \$1.00 / CPU / Day | \$1.00 / CPU / Day |
| Rate Factor x CPU Usage                                                                                                              | \$0.75 / CPU / Day                            | \$1.25 / CPU / Day | \$1.00 / CPU / Day |
| ▶ Memory Rate Factor                                                                                                                 | 1x                                            | 1x                 | 1x                 |
| ▶ Storage Rate Factor                                                                                                                | 1x                                            | 1x                 | 1x                 |

- When done, click the **Save** button in the upper right hand corner to complete the extended charge plan revision for the target type.

You can now assign the extended charge plan to a target instance of the appropriate target type. Remember that extended charge plans are automatically effective from the beginning of the month.

### 22.2.1.5 Revising Extended Charge Plans

You can update an extended charge plan in the following ways:

- Make changes to the charge rates in effect for the current or a future cycle.

- Create a plan revision for the next or a later report cycle, based on an existing plan.

To make changes to the charge rates in effect for the current or a future cycle:

1. Select the plan revision in the navigation pane and click **Set Rates**.
2. Make adjustments to the charge items and rates in effect.
3. Click **Save** to update the plan revision.

Note the warning when changing charge rates for the current cycle that the changes are retroactive to the beginning of the cycle.

To create a plan revision for the next or a later report cycle, based on an existing plan revision:

1. Select a plan in the navigation pane, then select **Revision** from the **Create** menu.
2. In the dialog that opens, select the effective date of the revision. The default date is the first month after the most recently added revision. For example, if the current cycle is for May and a June revision already exists, July 01 is the default effective date. Click **OK**.
3. In the familiar create-like model, the configurations, charge items, and rate adjustments for the plan you selected in the navigation pane appear in the plan details table on the right.

Edit the plan details as desired:

- Add and remove target types.
  - Add and remove configurations.
  - Add, change, and delete charge items.
  - Make adjustments to metric rates.
4. When done, click **Save** to complete the plan revision.

## 22.2.2 Creating Cost Centers

Cost centers are vehicles for aggregating charges. A cost center can be an individual, a department within an organization, or a multitiered business hierarchy that spreads charges across the enterprise. You can create a business hierarchy by setting up cost centers and assigning users to the business units.

A cost center can originate in the following ways:

- [Setting the Cost Center Property on the Target Home Page](#)
- [Creating the Cost Center in Chargeback](#)
- [Importing a Business Hierarchy from an LDAP Server](#)

### 22.2.2.1 Setting the Cost Center Property on the Target Home Page

When you make a server request in the process of setting up a self service application, the Cost Center target property of the target, VM Guest for example, is automatically set to the user login of the person making the request. When the target is configured for Chargeback, the user login appears in the Cost Center column for that target on the **Targets** tab and is consigned to the default cost center.

Similarly, you can manually set a target's Cost Center target property to an appropriate value (user ID), which will subsequently be imported and appear in the Cost Center column on the **Targets** tab when the target is configured for Chargeback. The user ID will also be consigned to the default cost center.

To create a cost center in this fashion:

1. On the All Targets page, select the target in question and go to the target home page.
2. From the target type menu, select **Target Setup**, then select **Properties**.
3. Click **Edit** and enter a value for the **Cost Center** property, then click **OK**.

The value will appear as described when the target is configured for Chargeback.

Note that if you change the Cost Center property for a target that has already been added for Chargeback, the system does not recognize the new value.

### 22.2.2.2 Creating the Cost Center in Chargeback

A user who provisions a target using the self service portal appears in the default cost center as a consequence of adding the SSA target for Chargeback. These SSA users can subsequently be reassigned to appropriate cost centers as described in Step 4 of the process below. Reassignment can also be automated using LDAP.

To create a cost center manually:

1. From the **Enterprise** menu, select **Chargeback**.
2. Select the **Cost Centers** tab and then click the **Add** button.
3. Enter the appropriate information in the **New Cost Center** dialog:
  - The cost center takes a unique name by which it is known internally in the system.
  - The cost center also takes a meaningful name that appears in the user interface and generated reports.
  - Indicate the cost center's position in the tree hierarchy; that is, whether it should be root or nested. You can subsequently change its position in the hierarchy.

When you click **OK**, the new cost center appears in the table.

4. Cost centers typically equate to business organizations. When you select a cost center row, users who are part of the organization appear in the region below. Users can also be unassigned.

In addition to the cost centers you create, there also is a system default cost center that serves as a clearinghouse for unassigned users. Included in this group are the users designated as the Cost Center target property on the target home page. You can reassign these users to roll up charges to a higher-level cost center.

- a. First, select the default cost center row in the Cost Center table. This exposes the users consigned to the default cost center.
- b. Next, select a user row in the Users table.
- c. Click the **Assign to Cost Center** button that is now active.
- d. In the dialog that opens, designate the cost center to which to assign the user, then click **OK**.

The user now appears in the designated cost center.

5. As the cost centers list grows, you may find that you want to combine cost centers, sales under marketing, for example. To do so, select the cost center row, then select **Move** from the **Action** menu. The change is effective from the beginning of the current reporting cycle.

In the dialog that opens, indicate whether to include nested cost centers and choose where in the tree to move the cost center, then click **OK**.

6. At some point, you may decide to deactivate the cost center. To do so, select the cost center row and click the **Remove** button. Indicate whether to deactivate nested cost centers as well. The change is effective from the beginning of the current reporting cycle. Deactivation means the cost center can no longer be assigned to a target, but its association with targets remains in previous report cycles; that is, if you run a report for a previous cycle, the target assignment reflects the deactivated cost center.

Note the following about cost center moves, for example when a cost center (and its users) move to a different parent:

- The user assignment of a target resource remains the same: the PC assigned to JDoe remains assigned to JDoe.
- Likewise, the resource usage assignment remains the same, but the report rollup reflects the new cost center parent child relationship, unless the administrator changes the assignment. The cost center move effects a change in the usage assignment, which is associated with the cost center and user ID.

### 22.2.2.3 Importing a Business Hierarchy from an LDAP Server

You can import a business hierarchy from an LDAP server. LDAP servers supported include:

- Oracle Internet Directory
  - Microsoft Active Directory
  - Sun iPlanet
  - Novell eDirectory
  - OpenLDAP
1. From the **Enterprise** menu, select **Chargeback**.
  2. Select the **Cost Centers** tab.
  3. From the **Action** menu, select **LDAP Settings**.
  4. The **LDAP Settings** dialog opens; it takes the following values:
    - Select the LDAP server type from the drop-down list (default is Oracle Internet Directory).
    - Specify host name of the LDAP server.
    - Select what to use as the authentication ID (which attribute the LDAP server uses to authenticate: user name or e-mail address).
    - If a login is required, mark the check box and provide the credentials (where username is the LDAP server Principal, for example, cn=orcladmin).
    - Enter the port number of the LDAP server.
    - Check the box for a secure connection.
    - Check the box to enter a custom search base (base DN where users are stored; for example, dc=oracle,dc=com).
    - Select the level to which to maintain the cost center. This refers to the level within the cost center hierarchy to be synchronized with the LDAP server at the end of each report cycle. Note that any levels lower in LDAP than the

specified level are subsumed into the specified level in Chargeback. So, for example, if you specify level 5, and LDAP has 7 levels, LDAP levels 6 and 7 are subsumed into level 5 in Chargeback.

Click **OK**.

An LDAP cost center overwrites a manually created cost center of the same name while maintaining target assignments. You can opt to reassign targets to some other cost center.

A successfully imported LDAP cost center hierarchy generates a scheduled repeatable job to run at the beginning of each report cycle (typically the first of the month) to sync the hierarchy with its counterpart on the LDAP server. You can also perform this operation on an ad hoc basis by selecting **On-demand LDAP Sync** from the **Action** menu.

Note that the LDAP Sync job only synchronizes existing SSA users; that is, SSA users who are already known to Chargeback as a result of their SSA targets being added for Chargeback. The sync job does not automatically synchronize any other (SSA or other) users in the LDAP directory.

When an SSA target is added for Chargeback, the SSA user appears in the default cost center automatically. The LDAP Sync job then populates the cost centers hierarchy with default cost center SSA users, based on the management hierarchical structure, as determined by the **MANAGER** attribute in the LDAP directory. The sync job also overwrites the cost centers of any SSA users who were manually reassigned, based on the same management hierarchical structure.

### 22.2.3 Adding Targets for Chargeback

Decide on the targets whose resources you want to manage within Chargeback. You have to add targets to start collecting metrics against which to calculate charges. Note that if you want to add a Real Application Cluster (RAC) that has member instances already monitored in Chargeback, you must first remove those instances from Chargeback.

For certain target types, specifically, shared database instances, RACs, and multitenant container databases (CDBs), you have to enable metric collection before you can add targets of these types for Chargeback. For information on enabling metric collection, see [Section 22.2.3.1](#).

A CDB is a composite target. When you add a CDB instance, all of its PDBs are automatically discovered and reported as children of the CDB instance. Usage and charges are calculated at the PDB level.

You must have the **ADD\_CHARGEBACK\_TARGET** role assigned to add any target monitored in the Enterprise Manager instance that qualifies for Chargeback support. Without this privilege, the **Add Targets** button is disabled. The **VIEW\_CAT\_TARGET** role allows you to see Chargeback data related to a specific target. The **VIEW\_ANY\_CAT\_TARGET** role allows you to see Chargeback data related to any target.

1. From the **Enterprise** menu, select **Chargeback**.
2. Select the **Targets** tab and click the **Add Targets** button.

The Add Targets wizard opens, where you perform the following steps:

- a. [Selecting Targets for Chargeback](#)
- b. [Identifying Shared Targets](#)
- c. [Making Assignments](#)

#### d. Reviewing Target Setup

It may take up to 24 hours after a target is added to Chargeback for its metering data to appear. To stop collecting data on targets, select the target or targets involved and click the **Remove Targets** button.

---

**Note:** When you remove a target from Chargeback scope, its usage and charge data are dropped from the current report cycle. A warning to this effect appears as part of the removal process.

---

As new services for cluster members are discovered as part of the daily collection, they appear automatically in the tree hierarchy under the RAC node. The charge plan in effect is automatically assigned. Make cost center assignments accordingly.

#### 22.2.3.1 Configuring Enterprise Manager Metric Collection for Shared Database Targets and Multitenant Container Databases

For shared database instances, RACs, and CDBs, metric collection is disabled by default. You have to enable it before Chargeback can manage resources on these targets.

To effect change to metric data collections, you must be assigned the `MANAGE_TARGET_METRICS` role on the specific target.

1. From the **Targets** menu, select **Databases**.
2. Locate the target database instance in the table and open its home page (provide credentials if requested).
3. From the **Oracle Database** menu, select **Monitoring**, then select **Metric and Collection Settings**.
4. On the Metric and Collection Settings page, change the view to **All Metrics**.
5. Enable metric data collection for **EM Database Services**. Be sure you are enabling EM Database Services, not simply Database Services.

Do this for each cluster member in a RAC and for each CDB instance.

#### 22.2.3.2 Selecting Targets for Chargeback

Step 1 in the Add Targets wizard is to select the targets you want to add for Chargeback.

1. Click the **Add** button.
2. In the dialog that opens search for the target types you want to add. Use criteria to filter your search.
3. In the search results list, select the specific targets you want to add. Note that you can multiselect targets. Click **Select**.

The dialog closes and the selected targets appear in the table. Adding a composite target automatically adds its member targets. The metered flag indicates that metrics and configuration data are being collected for the target and used in charge calculations. For example, given a WebLogic domain composite target with two WebLogic Server member targets, the server targets have the metered flag, but the domain target does not.

4. Click **Next** to identify shared targets.



### 22.2.3.3 Identifying Shared Targets

Step 2 in the Add Targets wizard is to identify targets to be shared across cost centers.

The targets you selected for Chargeback in the first step appear in the table. Those eligible for sharing, in particular, databases, Real Application Clusters (RACs), and WebLogic Servers, have a check box at the right of the row. Select the check box next to each target you want to be shareable. Select the column header check box to make all eligible targets shareable.

Database targets share services; WebLogic Servers share applications. The Shared Entity column contains a link that displays a list of the services or applications that the respective database or WebLogic Server target shares.

A RAC can be shared by instance or by service. To share by instance, Add a RAC in dedicated mode. This way, each RAC member appears under the RAC node in the tree hierarchy and can be assigned to different charge plans and different cost centers. To share by service, add the RAC in shared mode. Services appear under the RAC node in the tree hierarchy and can be assigned to different cost centers, but must all be assigned to the same charge plan.

Database targets that appear in the bottom table are not eligible for sharing. To make a database instance or RAC shareable, you must enable Enterprise Manager metric collection on each database instance. If you add a database target before enabling metric collection, you will have to remove it and add it back in shared mode after enabling metric collection.

Note that all targets appearing on this page will be added for Chargeback; those not identified as shared will be added in dedicated mode

Click **Next** to make charge plan and cost center assignments.

### 22.2.3.4 Making Assignments

Step 3 in the Add Targets wizard is to make charge plan and cost center assignments.

Charge plans provide a basis for calculating the cost of resource consumption.

1. Select the target (or targets) in the table and click the **Assign Plan** button.
2. In the dialog that opens, select the universal or extended charge plan appropriate to target selection and click **OK**.

You can opt not to assign charge plans now and assign them after adding the targets for Chargeback. See [Section 22.2.5](#) for more information.

Cost centers provide a way to aggregate the charges accrued for resource consumption.

1. Select the target (or targets) in the table and click the **Assign Cost Center** button.
2. In the dialog that opens, select the cost center appropriate to target selection and click **OK**.

The dialog closes, and the cost center assignment appears in the target table row.

You can opt not to assign cost centers now and assign them after adding the targets for Chargeback. See [Section 22.2.4](#) for details.

When done, click **Next** to review target setup and complete the task of adding targets for Chargeback.

### 22.2.3.5 Reviewing Target Setup

Step 4 in the Add Targets wizard provides an opportunity to review your selections and assignments, before completing the task. Note that for database and WebLogic Server targets, the wizard also adds the associated host in this view. If satisfied, click **Submit**.

You return to the **Targets** tab where a message confirms that *n* number of targets were successfully added. This number includes hosts associated with database and WebLogic Server targets. For more information on what happens after targets are added, see the paragraphs after the four wizard steps in [Section 22.2.3](#).

## 22.2.4 Assigning Cost Centers to Targets

Cost centers are a way of distributing charges for resource usage. A cost center is automatically assigned to an SSA target as part of SSA setup (Cost Center target property). Use this manual process only for assigning cost centers to non-SSA targets.

---

---

**Note:** You also can perform cost center assignment when adding a target for Chargeback. See Step 3 of the Add Targets wizard described in [Section 22.2.3.4](#).

---

---

1. From the **Enterprise** menu, select **Chargeback**.
2. Select the **Targets** tab.
3. Select the targets to which you want to assign a cost center and click the **Assign Cost Center** button.
4. A dialog opens displaying a list of cost centers. Choose the cost center for the target (s) and click **OK**.

The table of target assignments refreshes to reflect the assignment.

When charges are calculated, unassigned targets are assigned to the default cost center. Note that assigning a cost center to a composite target (system or group) maps directly to target members. For a composite target other than a system or group, the cost center assignment is directly to the composite target, and to any children that otherwise have no direct cost center assignment.

Shared services under a RAC node can be assigned to different cost centers.

---

---

**Note:** For targets requested and provisioned through SSA, the Cost Center target property is set to the user login. Oracle recommends that to change the cost center assignment, you do so from the **Cost Centers** tab by selecting the user login in the default cost center and assigning to another cost center.

---

---

To unassign a cost center, select the target or targets involved and select **Unassign Cost Center** from the **Action** menu.

## 22.2.5 Assigning Charge Plans to Targets

When assigning charge plans to targets, you have the option of assigning plans individually to child members or assigning at an ancestor level and having the child members inherit the assigned plan.

An extended charge plan includes parameters for how to apply universal charge rates to an assigned target. You also can replace an extended plan by reassigning the universal plan to a target.

For SSA targets (zones), the person performing SSA setup assigns the charge plan to the zone as part of setup so that all targets within the zone inherit the plan. Targets within the zone can be assigned a different plan, if necessary, using the manual process described below. Note, however, that if a target within a zone already has an assigned plan, and the desire is to have all targets within the zone have the same plan assigned at the zone level, the plan already assigned must be removed from the target.

---

**Note:** You also can perform charge plan assignment when adding a target for Chargeback. See Step 3 of the Add Targets wizard described in [Section 22.2.3.4](#).

---

1. From the **Enterprise** menu, select **Chargeback**.
2. Select the **Targets** tab. A hierarchical list of targets appears showing assigned cost centers and charge plans.
3. Select the targets to which you want to assign a charge plan and click the **Assign Plan** button.
4. A dialog opens displaying a list of plans that apply only to the target type selected. Choose the plan you want to assign, then click **OK**.

The table of target assignments refreshes to reflect the assignment.

Note that assigning a charge plan to a composite target (system or group) maps directly to target members. For a composite target other than a system or group, the charge plan assignment is directly to the composite target, and to any children that otherwise have no direct plan assignment.

For shared services nested under a RAC node, charge plan assignment is the same for all services. Assigning to one applies to all. Any change in plan assignment is propagated across all services under the RAC node.

A RAC that is shared by instance can have different charge plans assigned to cluster members because they function in dedicated mode.

To unassign a charge plan, select the target or targets involved and select **Unassign Plan** from the **Action** menu.

## 22.2.6 Configuring Chargeback Settings

Click the **Chargeback Settings** link on the **Home** tab to configure global settings for currency symbol and uptime calculations.

### Currency Symbol

You can change the default currency (USD) by clicking in the currency symbol text box and entering the desired currency symbol. The new selection becomes the default currency across all charge plans; that is, the universal plan and all extended charge plans. All reports, including historical reports, reflect the new currency. Note that no rate conversion occurs when you change the currency; that is, the numbers stay the same. Change from dollars to euros and a one dollar charge becomes a one euro charge.

### Uptime Calculations

Select the appropriate radio button to ignore or include uptime in charge calculations. The default is to include uptime as a consideration. Note that to include uptime has an impact on all fixed and configuration-based charge calculations for all targets. Based on the number of targets and charge plans in play, this can affect performance, so use discretion when including uptime in the equation.

Chargeback prorates charges and discounts accordingly. So, for example, if a target was available 22.5 hours in a 24-hour period, the daily charge would be adjusted 1.5 hours. A change in the uptime setting is effective from the beginning of the current report cycle, but does not impact previous report cycles; that is, charges in historical reports are not prorated based on a change made to the setting in the current cycle.

### Example

To demonstrate the difference in factoring in uptime, consider the following charge plan configurations:

- A universal charge plan that charges two dollars a day per gigabyte of memory allocation, a dollar a day per gigabyte of storage, and no charge for CPU usage
- An extended charge plan for a Guest VM (medium), with a base charge of a dollar an hour and a memory allocation adjustment of x5

Given these charge plans, a user requisitions a Guest VM with 4 GB of memory and 100 GB of disk space for two consecutive days. The first day the uptime metric records 20 hours and the second day, 24 hours.

When uptime is *not* a factor in the calculations, the charge per day is 164 dollars, computed as follows:

- Forty dollars a day for memory ( $4 \text{ GB} * \$2 * 5$ )
- A hundred dollars a day for storage ( $100 \text{ GB} * \$1$ )
- A base charge of 24 dollars a day ( $24 \text{ hours} * \$1$ )

When uptime is a factor, charges for the first day are 137 dollars (rounded up), computed by subtracting the four hours of downtime from the per day charge ( $20/24 * 164$ ).

## 22.3 Accessing Chargeback Information

The payoff for setting up Chargeback is the wealth of information it puts at your fingertips. The information is available in a variety of formats, including pie charts and bar graphs. You can drill down to details by cost center, target type, and resource. Trending reports show charges and resource utilization over time, up to a year in fact, which can help in IT resource planning. In addition to the reporting function available to administrators within the application, self service users can view Chargeback information related to their targets within the self service portal, and Line of Business users can view Chargeback information in BI Publisher reports.

This section covers the following topics

- [Following Usage and Charge Trends](#)
- [Viewing a Target's Collected Metric Data](#)
- [Generating and Distributing Chargeback Reports](#)
- [Viewing Chargeback Information in the Self Service Portal](#)
- [Sharing Chargeback Data with Other Applications](#)

## 22.3.1 Following Usage and Charge Trends

Once you define charge plans, set up cost centers, and begin to monitor targets, trends in usage and charges emerge. The **Home** tab displays bar graphs that denote usage and charge trends over time. Each trend region has an **Option** link in the upper-right corner that opens a dialog where you can customize the display.


### Usage Trend Options

Customize the usage trend display as follows:

- **Time Period**—Select the desired time period to chart usage trends; default is the current year
- **Aggregation**—Select how to group usage across the selected time period; default is monthly
- **Usage Value Type**—Show usage as a percent (utilization) or as a consumptive value; default is utilization
- **Resource**—Show usage in terms of a particular metric (CPU, memory, or storage); default is the three metrics

Figure 22–5 displays a usage trend showing CPU, disk, and memory utilization over the past week, aggregated daily.

**Figure 22–5 Sample Usage Trend Graph**




### Charge Trend Options

Customize the charge trend display as follows:

- **Time Period**—Select the desired time period to chart charge trends; default is current year
- **Aggregation**—Select how to group charges across the selected time period; default is monthly
- **Group By**—Group charges by either cost center or target type; default is cost center
- **Cost Center**—Display charge trends for all cost centers, a particular cost center, or for unassigned users; default is all

Figure 22–6 displays a charge trend showing last week’s charges for the VM Guest target type, aggregated daily.

**Figure 22–6 Sample Charge Trend Graph**

### 22.3.2 Viewing a Target's Collected Metric Data

You can view a target's metrics in graph form or as a time-stamped report of raw metric data. You can export the latter view to an external application such as a spreadsheet. You also can check metric collection status.

Note that you might have to expand the metric data region at the bottom of the **Targets** tab to expose the subtabs referenced in the instructions below.

1. From the **Enterprise** menu, select **Chargeback**.
2. Select the **Targets** tab and then select a target to view.
  - Select the **Chart** subtab below the targets list, to view a wave graph of data metrics for the target. Select the metric to display from the drop-down list on the right. The metric selection is target-type dependent. For a shared WebLogic Server, the metrics are per application. For a shared database, the metrics are service-related. For a RAC, metrics are represented as a multiline graph grouped by cluster member, where each line maps to usage by a particular member for the selected service.
  - Select the **Data** subtab to view the raw numbers. In this view, you can filter the data and export it as a csv-formatted file. You also can change the focus to another target. For a RAC, the view shows the same service data for all cluster members within the RAC.
  - Select the **Collection Status** subtab to glean statistical data such as error count and when metric data was last collected. For a RAC, collection status is for all cluster members within the RAC regardless of service selected.

Figure 22–7 shows a chart tracking instance uptime for the selected target over a span of six days.

**Figure 22–7 Sample Target Metric Data**

Note that you can see additional details on the Job Activity page.

1. From the **Enterprise** menu, select **Job**, then select **Activity**.
2. Click Advanced Search and specify the following filtering criteria:
  - Select **Chargeback Data Loader** as Job Type.
  - Select **Targetless** as Target Type.
  - Select appropriate values for Status and Scheduled Start.
3. Click **Go**.

The log output for each job step captures any warnings or errors encountered.

### 22.3.3 Generating and Distributing Chargeback Reports

Chargeback reports are a powerful analytical tool for tracking resource usage and charge distributions.

Data collection occurs once a day. The daily data collection job for the current cycle is based on charge plan and cost center assignments. The reporting cycle defines the time period for which to calculate charges. The cycle is for the current month starting on the first day of the month.

---

**Note:** Changes in Chargeback setup are not immediately reflected in reports. For example, if you change a cost center assignment, it may take up to 24 hours (until the next data collection job runs) for reports to show the change. If you have SYSMAN Super Administrator privileges, however, you can initiate data collection on-demand from the **Targets** tab. To do this, select **On-demand data collection** from the **Action** menu.

---

To generate ad hoc reports:

1. From the **Enterprise** menu, select **Chargeback**.
2. Select the **Reports** tab.
3. Design your report from the following options:
  - Use the current report cycle or customize a date range to report on.

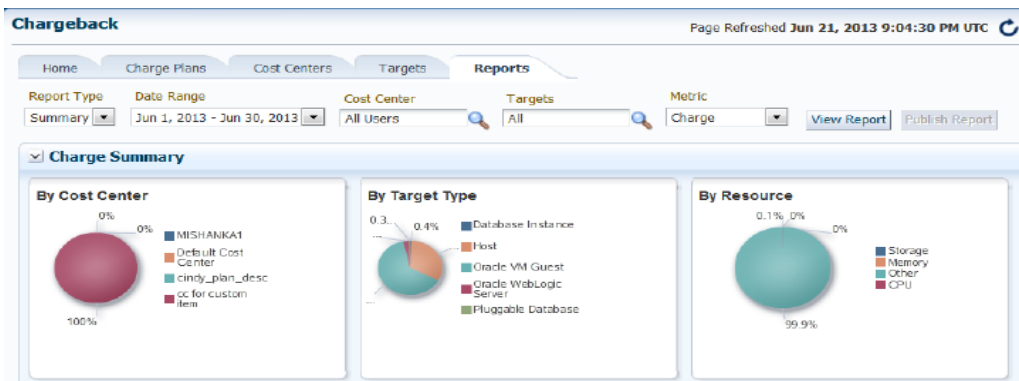
- Choose between summary and trend report types. A summary report presents a pie-chart breakdown, while a trend report uses a y-axis multiple bar chart to denote usage trends.
- Select specific cost centers or report on all users.
- Select specific targets or target types or report on all targets within all target types.
- Choose the metric to report on.

Click **View Report** to see the results.

The report displays color-coded graphs summarizing charges by cost center, target type, and resource, with details displayed in the table at the bottom. Click a color box link in the respective graph to recalculate the report contents for the color-coded selection, for example memory in the resource graph.

Figure 22–8 displays a summary report showing charges for the current reporting cycle for all cost centers and target types, with a breakdown by resource.

**Figure 22–8 Summary Report for a Cost Center and Target Type**



4. Filter the details by choosing from the drop-down list; the default is All. Use the query-by-example feature (icon to the right of the action buttons) to search report details. The icon acts as a toggle; clicking it alternately shows or hides text and selection boxes above the table columns. The feature is also available in the **View** menu. Enter search criteria in various combinations by selecting a date and by typing values in the respective columns. Press **Enter** to activate the search.
5. Click the **Export** button in the details region to export report contents to file.
6. Click **Publish Report** to make report contents public. This action integrates with BI Publisher, where you can:
  - Save reports in a variety of formats (Excel, PowerPoint, HTML, PDF)
  - Distribute generated reports to e-mail lists (users who do not have access to Enterprise Manager, for example) on a defined schedule

For information on BI Publisher setup, see the "Installing BI Publisher on Enterprise Manager" chapter in the *Enterprise Manager Cloud Control Advanced Installation and Configuration Guide*.

### Reporting on Uptime

When you choose to report on the uptime metric, the details display the fractions of hours in a day a target was available over the course of the reporting cycle. Note that reporting uptime does not prorate charges. The prevailing sense is to put in the hands



of the customer how best to make adjustments to charges based on the report results. To this end, export the results to a Microsoft Excel html-formatted file for further analysis and consideration.

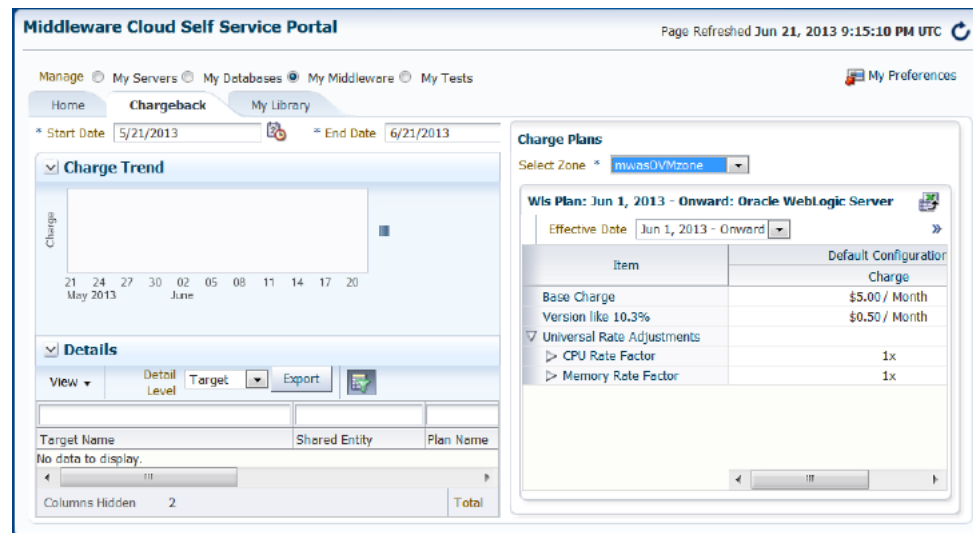
### 22.3.4 Viewing Chargeback Information in the Self Service Portal

You can access Chargeback information as it relates to self service applications from the self service portal.

1. From the **Enterprise** menu, select **Cloud**, then select **Self Service Portal**.
2. Select the portal type radio button (**My Servers**, **My Databases**, or **My Middleware**).
3. Click the **Chargeback** tab to view the following informational breakdown:
  - **Charge Trend** calculates charges across a date range for various metrics. If you hover over a metric, memory for example, the bar graph changes to reflect only that metric. Set a date range and click **Go** to recalculate charges.
  - **Details** below are a reflection of the charges and date range above. You can reorder the data by selecting from the **Detail Level** drop-down menu. You can also display additional information by exposing more columns.
  - **Charge Plans** allows you to select a zone to see the charge plan in effect, in this case, an extended charge plan defined for a specific target type. That target type, Oracle VM Guest for example, has certain charge items that were set as part of the plan definition and are displayed on the **Charge Items** tab. The charge plan determines the charges displayed in trends and details on the left. The other tab, **Universal Rate Adjustments**, shows adjustments if any made to the base plan rates (CPU, memory, and storage).

Figure 22–9 shows a sample Chargeback informational breakdown in the self service portal.

**Figure 22–9 Chargeback Breakdown for a Self Service Application**



### 22.3.5 Sharing Chargeback Data with Other Applications

Amassed Chargeback data is a valuable commodity that can be made available to other applications. There are several ways to do this:

- Export administrative reports (**Export** button on the **Reports** tab in Chargeback)
- Publish administrative reports via BI Publisher (**Publish** button on the **Reports** tab in Chargeback)
- Export details of charge trends for self service applications (**Export** button on the **Chargeback** tab in the self service portal)
- Execute the EM CLI verb `get_metering_data` to generate comma-separated output on usage and, optionally, charge data (see Chapter 29, "Chargeback and Metering EMCLI Verbs," for more information)

In addition, Enterprise Manager supports repository views to access information in the Management Repository for further processing and presentation. There are two views available (`MGMT$EMCT_CBA_CHARGE_HOURLY` and `MGMT$EMCT_CBA_CHARGE_DAILY`) to access Chargeback information in the Management Repository. See the *Enterprise Manager Cloud Control Extensibility Programmer's Reference* for more information.

---

## Using Consolidation Planner

This chapter covers use of the Consolidation Planner feature provided with Oracle Enterprise Manager Cloud Control.

This chapter includes the following sections:

- [Overview of Consolidation Planner](#)
- [Consolidation Constraints](#)
- [Using Consolidation Planner](#)

### 23.1 Overview of Consolidation Planner

Over the years, the typical enterprise data center will grow steadily due to the addition of servers required to satisfy the needs of the business. This growth typically results in excess servers that occupy rack space, consume a lot of power for cooling, and require system maintenance such as security and patching.

Depending on the procurement cycle or the specific hardware vendor agreement in effect, enterprises may acquire different types of server hardware and operating systems, inadvertently creating a confusing array of systems that administrators need to manage, administer, patch and upgrade. This in turn increases labor and ongoing maintenance and support costs against IT budgets. Enterprises look at consolidation as a way to migrate their disparate systems onto standardized operating systems and hardware, with the ultimate goal of reducing costs.

Enterprises also are increasingly investigating virtualization technologies such as Oracle Virtual Machine by moving from physical to virtual servers. This makes it possible to use the shared hardware infrastructure while getting the benefits of isolation that virtualization provides.

The goal of consolidation is to identify such under-utilized servers and find a way to consolidate them, enabling the enterprise to free up as many servers as possible while continuing to maintain service levels. Since servers have different levels of CPU capacity, Consolidation Planner uses computer benchmark data to normalize the CPU usage for a given hardware in the consolidation process. Specifically, Consolidation Planner uses the following CPU benchmarks for different classes of hardware:

- SPECint®\_base\_rate2006 for database hosts, application hosts, or mixed-workload hosts
- SPECjbb®2005 for middleware platforms

The Consolidation Planner feature enables you to match managed servers you want to consolidate with the generic physical machines, Oracle engineered systems (Exadata Database Machines or Exalogic Elastic Cloud systems), or Oracle Virtual Machine

(VM) servers they can be consolidated to. By leveraging metric and configuration data collected from managed target servers by Enterprise Manager Cloud Control, Consolidation Planner helps you determine the optimum consolidation scenarios that also reflect business and technical constraints in the consolidation process.

### 23.1.1 Key Concepts

The following concepts are central to using Consolidation Planner.

#### **Source Server**

An existing server that will be considered for consolidation.

#### **Destination Server**

An existing or yet-to-be-purchased server that a source server will be consolidated to. Can also be thought of as the consolidation target. These may be individual machines, virtual server pools, or an engineered system such as the Oracle Exadata Database Machine or an Exalogic Elastic Cloud system.

#### **Consolidation Project**

Defines the scope of a potential consolidation effort, including:

- The type of consolidation. In the current release, two types of consolidation schemes are supported:
  - P2V: From physical source servers to Oracle Virtual Machine (VM) destination servers
  - P2P: From physical source servers to physical destination servers
- The preliminary set of candidate source servers to consider consolidating *from*
- The preliminary set of candidate destination servers to consider consolidating *to*
- The duration over which data used to generate consolidation scenarios will be collected for the source servers
- The benchmark used to measure CPU capacities when determining how many source servers can be consolidated to a destination server

#### **Consolidation Scenario**

Each consolidation project contains one or more consolidation scenarios that are generated based on the inputs provided. Inputs provided to a scenario include:

- The source server resource requirements that a destination server must meet, including one or more of the following: CPU, memory, disk I/O, network I/O, and disk storage
- Any business, compliance or technical constraints that must be considered
- The destination servers to consider in the scenario

A set of pre-configured consolidation scenarios are provided, representing conservative, aggressive, and medium consolidation schemes. Each scenario is generated based on inputs you provide. Alternatively, you can create your own custom scenarios that best suit your situation. Once created, you can compare the various scenarios to determine which consolidation strategy best meets your requirements.

Each scenario also includes initial mappings between each source server and the destination server it may be consolidated to. You can choose to create mappings

manually, or allow Consolidation Planner to create them automatically. Once all inputs are specified, you can run the scenario and evaluate the results. Subsequently, you can rerun the scenario to re-evaluate the scenario based on the previously specified conditions with the latest available data. The results of the previous analysis will be over-written. You can also create a new scenario based on an existing scenario, where you tweak certain values to customize the new scenario.

## 23.2 Consolidation Constraints

Consolidation Planner allows you to specify various constraints that must be considered when creating consolidation projects and scenarios.

This section covers the following:

- [Source Server Constraints](#)
- [Destination Server Constraints](#)

### 23.2.1 Source Server Constraints

For source servers, you define constraints on the basis of compatibility or exclusivity.

#### **Compatible Servers**

Servers are considered compatible if they match on certain property and configuration values.

Properties include the following:

- Lifecycle Status
- Department
- Location

For example, there may be situations that mandate that specific servers remain within a specific location, such as data center location or geography.

Configuration items include the following:

- Network Domain
- System Vendor
- System Configuration
- CPU Vendor
- CPU Name
- Operating System

For example, there may be an accounting policy that the same system and CPU vendor be used.

You can choose to establish compatibility on none, some, or all of these target properties and configuration items.

#### **Mutually Exclusive Servers**

You can choose to exclude servers from consolidation scope because they violate certain Oracle best practices. Set either or both of the following conditions to exclude matching servers:

- Nodes of a RAC Database—do not consolidate nodes of the same RAC database to a single destination server
- Nodes of an Oracle Cluster—do not place nodes of an Oracle cluster in the same failure group

### 23.2.2 Destination Server Constraints

For destination servers, you can scope candidates to either new or existing candidates, but constraints are primarily expressed as a percentage of CPU and memory resource utilization; that is, how much of either resource type can maximally be used on a destination server.

## 23.3 Using Consolidation Planner

The steps in the consolidation planning process are:

1. Create a consolidation project. See [Section 23.3.1, "Creating a Consolidation Project"](#).
2. Define one or more consolidation scenarios within the project. You have two options:
  - Use a pre-defined consolidation scenario. See [Section 23.3.2, "Using a Pre-configured Consolidation Scenario"](#).
  - Create a custom consolidation scenario. See [Section 23.3.3, "Creating a Custom Consolidation Scenario"](#).
3. Evaluate your consolidation scenarios in the Consolidation Planner console to determine the consolidation strategy that best meets your needs. See [Section 23.3.5, "Evaluating Consolidation Scenarios"](#).
4. Modify the settings of your scenarios to generate different results. Continue this process until you have the most optimal scenario(s) for your situation.

Note that in this release of Enterprise Manager Cloud Control, consolidation scenarios are created for planning purposes only. Execution of scenarios - that is, the actual movement of software or data from source servers to destination servers - is *not* supported.

When consolidating multiple source servers to more than one destination server, the resource requirements of source servers are checked against the resource capacity of the destination servers. To consolidate all identified source servers to the least number of target servers, Consolidation Planner tries to identify a set of source servers that have known resource requirements that will fit into a destination server's corresponding available resources as tightly as possible.

For example, if the available memory in a target server is 2 GB, Consolidation Planner will try to find a set of source candidate servers with a sum of required memory as close to 2 GB as possible, then "fit" the source servers in the target server. The goal is to "fit" the source servers into the least number of destination servers. Note, however, that you can also choose to spread the source server loads evenly across the destination servers. In this case, Consolidation Planner tries to ensure that all the destination servers have the same or almost the same resource utilization after the consolidation.

This section includes the following:

- [Creating a Consolidation Project](#)

- [Using a Pre-configured Consolidation Scenario](#)
- [Creating a Custom Consolidation Scenario](#)
- [Other Scenario Creation Options](#)
- [Evaluating Consolidation Scenarios](#)
- [Managing Data Collections](#)

### 23.3.1 Creating a Consolidation Project

You will create a consolidation project for each consolidation effort, then add individual consolidation scenarios within it. You can then compare consolidation scenarios to determine which consolidation strategy makes the most sense.

After the project is defined, a Cloud Control job is submitted to collect available data for the specified servers from the Management Repository. Once the job finishes, the project becomes an active project. As long as the project is in an active state, data collection will continue.

1. From the **Enterprise** menu, select **Consolidation Planner**.
2. Click the **Create Project** button.
3. Enter the consolidation project name.
4. Select the consolidation type. Two types of consolidation schemes are supported: From physical source servers to Oracle VM servers (P2V), and from physical source servers to physical servers (P2P).

5. Establish operating system filtering criteria as appropriate.

Note that for either consolidation type, only source servers running on an operating system listed in the pull down menu can be consolidated.

6. Select an appropriate benchmark from the drop-down menu.
  - Specify SPECint®\_base\_rate2006 for database hosts, application hosts, or mixed-workload hosts
  - Specify SPECjbb®2005 for middleware platforms

Note that Consolidation Planner makes appropriate advisories for each selection.

7. Select the source servers to be added.
  - Click **Add Source Servers** to see a list of managed servers that could potentially be consolidated. Select the servers you want to add, then click **Select**.
  - Optionally click **Remove** to remove a server from the list.
8. Optionally select one or more existing servers to consolidate the source servers to.
  - If you are consolidating from physical servers to Oracle Virtual Servers (P2V), click **Add Existing Virtual Servers as Destinations** to view a list of existing VM-based Exalogic Elastic Cloud systems and Oracle Virtual Machine destination servers to consolidate the source servers to. Use the target type filter to differentiate the two. Select the servers you want to add, then click **Select**.
  - If you are consolidating from physical servers to physical servers (P2P), click **Add Existing Oracle Engineered System** to search for the Exadata Database Machine servers or Exalogic Elastic Cloud servers to consolidate the source servers to. Select the servers you want to add, then click **Select**.

9. Optionally set server I/O capacities for disk I/O request and network I/O volume capacities. Click **Specify Server I/O Capacity** to estimate these I/O capacities for all source and destination servers involved in the consolidation project. Note that you can subsequently fine-tune these estimates by editing the values for each server in the table.
10. Specify the duration over which data used to generate consolidation scenarios will be collected for the source servers specified in the project in the Data Collection region. This data will be used to determine the resource requirements that a destination server must meet.
  - Specify the minimum number of days to collect data. The default minimum value is 21 days. To use existing historical data to run and view consolidation scenarios immediately, set the minimum number of days to 0.
  - Specify the maximum number of days to collect data. The default maximum value is 90 days.
  - Specify when to begin the data collection process. Note that you can elect to suspend and resume data collection at any time from the **Actions** menu.
  - Optionally select **Continue Data Collection Over the Maximum Days** to purge the oldest day's data when data for a new day is added.
11. Click **Pre-configured Scenarios** if you want to use one or more of the out-of-the-box consolidation scenarios.

The pre-configured scenarios will be generated when the project is created using data collected for the source servers defined in the consolidation project.

You can also opt to create your own custom scenario once the consolidation project has been completed.

12. Click **Submit** when finished.

Once the project is created, it will show up in the Consolidation Planner console. Consolidation scenarios can then be defined for this project.

### 23.3.2 Using a Pre-configured Consolidation Scenario

When creating a consolidation project, you can optionally choose to generate up to three pre-configured consolidation scenarios to add to the project. These out-of-the-box scenarios represent conservative, aggressive, and medium consolidation schemes.

These scenarios are generated using data collected for the source servers defined in the consolidation project at the time the project is created. If no data is available when the project is created, the pre-configured scenarios will be automatically generated once the minimum amount of data has been collected.

1. During consolidation project creation, click **Pre-configured Scenarios**.
2. Select which scenarios you want to add to the project. Note that you can modify the name of each scenario.

By default, the scenarios are designated by the method used to aggregate daily source server resource usage:

- **Aggressive:** Aggregate data based on *average* daily usage per hour.

This typically results in a high consolidation (source server:destination server) ratio, because more source servers will “fit” into each destination server. But



because more source servers are involved, the odds that one or more will not meet the resource requirements are higher.

- **Conservative:** Aggregate data based on *maximum* daily usage per hour.

This typically results in a lower source server:destination server ratio, because fewer source servers will “fit” into each destination server.

- **Medium:** Aggregate data based on the 80 percentile usage.

This typically results in a source server:destination server ratio somewhere between Aggressive and Conservative aggregations.

Usage statistics are calculated based on the following criteria:

- **Resource Requirements:** The source server requirements, such as CPU, memory (GB) and disk capacity, that must be met by destination servers.
- **Applicable Dates:** The days of the week on which resource usage metrics are collected. Typically you will specify the day of highest resource usage.
- **Target Server Utilization Limit:** The maximum resource utilization (percentage) that can be used on destination servers.

3. Select destination candidates by choosing either of the following options:

- Click **Use New Servers** if you plan to use destination servers that do not exist yet, but you expect to be provisioned or purchased (also known as *phantom servers*).

Fill in the data for each server. Provide the estimated CPU capacity if available; otherwise, Consolidation Planner will determine requirements based on the CPU configurations provided earlier in the wizard.

Consolidation Planner will determine how many destination servers are required as part of the consolidation scenario generation.

- Click **Use Existing Servers** to use the destination candidates specified in the Consolidation Scope section of the consolidation project. For a P2P project, if you did not explicitly specify destination candidates, all source targets are potential destinations for consolidation.

Note that you can edit the CPU capacity and disk I/O fields.

4. Click **OK** when finished.

### 23.3.3 Creating a Custom Consolidation Scenario

You can create custom consolidation scenarios instead of or in addition to using the pre-configured scenarios. Multiple scenarios can be created within a project, enabling you to compare different scenarios before deciding on a solution. New consolidation scenarios are created within an existing consolidation project.

A wizard guides you through the creation process.

Note that you can also opt to use the pre-defined consolidation scenarios provided.

1. From the **Enterprise** menu, select **Consolidation Planner**.
2. Click on the consolidation project you will create the scenario within.
3. Click the **Create Scenario** button.
4. Specify the scenario details, such as scenario name.

5. Specify the source server resources to consider. Consolidation Planner will aggregate the specified resources to determine the total requirements.
  - Resource Type: The server requirements, such as CPU and memory (GB), that must be considered.
  - Scale Factor: Provide room for growth on the destination target for each source server. The resource requirement estimate uses the scale factor to pad the estimate for consolidation. So, for example, if the estimated requirement for a given source, based on usage data, is 2 GB of memory, which equates to a scale factor of 1, and you specify a scale factor of 1.1, 2.2 GB will be required to consolidate that source.
  - Applicable Days: The days of the week on which resource usage metrics are collected.
  - Resource Allocation: The method used to aggregate daily source server resource usage. Values are:
    - Aggressive: Aggregate data based on *average* daily usage per hour.  
This typically results in a high consolidation (source server:destination server) ratio, because more source servers will “fit” into each destination server. But because more source servers are involved, the odds that one or more will not meet the resource requirements are higher.
    - Conservative: Aggregate data based on *maximum* daily usage per hour.  
This typically results in a lower source server:destination server ratio, because fewer source servers will “fit” into each destination server.
    - Medium: Aggregate data based on the 80 percentile usage. This typically results in a source server:destination server ratio somewhere between Aggressive and Conservative aggregations.
  - The date ranges should define a period of time that is typical of standard resource requirements.
6. Click **Estimated Requirements** to view the estimated total resource requirements.

Resource requirements are shown based on the averaged hourly requirement. The displayed requirements reflect the scale factor (if any) specified for the resource. The 24-hour requirement pattern will be used as the minimum requirements that must be met by consolidation target(s).
7. Click **Next** to define consolidation constraints.

Specify business, corporate or technical constraints that must be enforced. These constraints will be used to:

  - Guide the server allocation process during automatic source-to-destination server mapping; or
  - Calculate violations if manual mapping between source and destination servers is used
8. Specify compatible servers.

Servers are considered compatible if they have the same specified target properties and configurations. If you are consolidating multiple source servers to a single target server, only compatible servers can be consolidated together on the same target server.
9. Specify mutually exclusive servers.

Certain types of source servers are mutually exclusive and should not be consolidated together on the same destination server due to various reasons. For example, nodes of a RAC database should not be consolidated on the same destination server; nodes of an Oracle cluster should not be placed in the same failure group. Contractual obligations may also restrict certain applications from running on the same server.

10. Optionally click **Preview Effect of Constraints** to view a list of source servers that are not compatible based on the defined constraints.
11. Click **Next** to specify destination server candidates.
12. Choose destination server candidates using either of the following options:

- Click **Use New (Phantom) Servers** if you plan to use target servers that have yet to be provisioned or purchased. The Targets Planning region shows the minimum required target servers based on the resource requirements of all servers in the scenario.

- For a P2V project, select either of the following options:

**Use Oracle Engineered System** and click the search icon to select an appropriate configuration type.

**Use Generic Servers** and provide the estimated benchmark rate if available; otherwise, click the search icon next to the CPU capacity input field, then select a server configuration that most closely matches your needs.

Accept or change the estimated resource requirements.

In a virtual environment, you can also specify a quantity of the resource to be set aside (reserved) for use by supervisory software in the database machine. This quantity is subtracted from the total capacity of the destination before consolidating source servers into the remaining resource. For example, if your estimated memory requirement is 12 GB and you specify a reserve of 2 GB, only 10 GB is available for consolidation.

- For a P2P project, select either of the following options:

**Use Oracle Engineered System** and select either Exadata Database Machine or Exalogic Elastic Cloud. Click the search icon and select a configuration type appropriate to either choice.

**Use Generic Servers** and provide the estimated benchmark rate if available; otherwise, click the search icon next to the CPU capacity input field, then select a server configuration that most closely matches your needs.

Accept or change the estimated resource requirements.

Consolidation Planner will determine how many target servers are required as part of the consolidation results.

- Click **Use Existing Servers** to specify a set of existing managed servers to use as targets for both P2V and P2P projects.

These are the servers you specified when defining the scope for the consolidation project. Consolidation Planner will determine the available hardware resources based on collected usage data. For a P2P project, if you did not explicitly specify destination candidates, all source targets are potential destinations for consolidation.

By default, the consolidation process will try to use as few target servers as possible. If you prefer, choose to balance the source load across all destinations.

13. Enter percentages for **Maximum Allowed Resource Utilization on Destination Servers**. The defaults are 75 for CPU and 90 for memory, respectively. Contrast these allowances, which provide headroom on destination servers, with the scale factor, which provides headroom for individual source servers.
14. Click **Next** to map the source servers to the destination servers they will be consolidated to. The objective is to fit source server requirements with each destination server's available resources as tightly as possible.

It is recommended that you allow Consolidation Planner to perform this mapping automatically. This will allow the tool to maximize resource utilization of destination servers based on server resource capabilities and the various consolidation constraints specified.

If you prefer to map each source and destination server manually:

- a. Click a source server in the list.
- b. Click the flashlight icon to select the destination server to map to the source server. Note that you can map a single source server to a target server or multiple source servers to a target server, but there can be only one target server.

The resulting consolidation report will show any resource and/or constraint violations due to such manual mapping.

15. Click **Next** to review the consolidation scenario. Use the **Back** button if you need to make changes; otherwise, proceed as follows:
  - Optionally, you can save the scenario as a template, which can then be shared with other users. If you want to do this, click **Save Scenario as a Template**. In the dialog that opens, browse to a location in the local file system and save the template as an XML file.
  - Click **Submit**. A message appears confirming that a job has been submitted for further analysis of the scenario. Results appear at the bottom of the Consolidation Planner page when done.

### 23.3.4 Other Scenario Creation Options

You can also create a consolidation scenario based on an existing scenario. Select an applicable scenario under a consolidation project and then select **Create Like Scenario** from the **Actions** menu. Modify the scenario as desired and save under a different name.

If you saved a scenario as a template, you can subsequently import the scenario into another environment.

1. On the Consolidation Planner home page, select **Create Scenario from Template** from the **Actions** menu.
2. Browse to a saved template XML file in the local file system. Click **Open**.
3. Indicate the extent to which you want replicate the saved template; that is, in terms of the resources, constraints, and targets planning represented by the template. Click **Update** if you make any changes.
4. Click **OK** to import the saved template.

The imported scenario opens in the wizard where you can edit and save it in Consolidation Planner.

### 23.3.5 Evaluating Consolidation Scenarios

You can view details for your consolidation scenarios using the Consolidation Planner console. After evaluating the consolidation scenario results, you can define different plans as well as rerun existing scenarios to re-evaluate them based on the previously specified conditions with the latest available data. The results of the previous analysis will be over-written. You can also create a new scenario based on an existing scenario, where you tweak certain values to customize the new scenario. This iterative process helps you obtain the optimized consolidation scenario which is generated by compromising various factors and weighing different trade-offs.

Compare the consolidation scenarios you create to determine which consolidation strategy best meets your requirements.

Your objective is to:

- Match source server resource requirements with destination servers best able to meet those requirements.
- Fit source server requirements with each destination server's available resources as tightly as possible, so you can get maximum usage of destination server capacity.
- Provide room for growth on destination servers by allowing for headroom as a factor of resource requirements.
- Optionally balance the source server workload across all available destination servers.

Note that in the current release, consolidation scenarios can be created for planning purposes only. Consolidation scenarios cannot be executed using Consolidation Planner.

1. From the **Enterprise** menu, select **Consolidation Planner**.
2. First, examine the project containing the scenario you want to view.
  - The **Status** column indicates the status of data collection, based on the minimum and maximum collection days specified for the project.
  - Click the **Source Servers** tab to view CPU, memory and disk storage data and utilization rates collected for the source servers defined in the project.
  - Click the **Source Workload** tab to view a graph showing source server resource usage data collected. Data is shown for each 24 hour period.
  - Click the **Report** button above the table when the project is selected to view summarized information and more details.
3. Next, view the data for a specific scenario. For a completed analysis, click any metric in the row to view details. Clicking the **Sources** metric takes you to the **General** tab. Clicking the remaining metrics takes you to the respective tab, as follows:
  - **Destinations:** The list of destination servers to which the source servers will be consolidated. Resource configuration and calculated utilization are shown for each destination server.
  - **Ratio:** The ratio of source servers to destination servers. By default, Consolidation Planner will try to "fit" source servers into as few destination servers as possible.
  - **Mapping:** The destination servers to which specific source servers will be mapped. The analysis includes estimated CPU and memory requirements and

utilization, enhanced by suggested CPU and memory allocation figures to consider. These suggestions represent a reasonable compromise between requirements and destination server capacity.

- **Confidence:** The percentage of the data collected for source servers that meet the source server usage requirements defined in the scenario. This value is aggregated for all source servers defined with the project.
- **Violations:** The number of violations of technical or business constraints defined in the scenario.
- **Exclusions:** The number of source servers that do not have a qualified mapping to a destination server. These are source servers that exceed the capacity of available destination servers. This metric is applicable only if auto-mapping of source servers to destination servers is used.

A different set of constraints may result in a different optimal scenario. Modify the constraints to come up with different scenario results.

### Share Reports with Others

Select a report in the table and click **Publish Report** to make report contents public. This action integrates with BI Publisher, where you can:

- Save reports in a variety of formats (Excel, PowerPoint, HTML, PDF)
- Distribute generated reports to e-mail lists (users who do not have access to Enterprise Manager, for example) on a defined schedule

## 23.3.6 Managing Data Collections

Manage data collections by viewing the status of your projects.

1. On the Consolidation Planner home page select **View Data Collection** from the **Actions** menu.
2. The view lists source targets within a project where you can perform the following tasks:
  - View the latest collection status by project.
  - Select a target to see its collection history and troubleshoot potential problems with the collection.
  - Click the link under Data Collection Jobs to go to the job activity page where you can view and administer the latest data collection job.
  - Update the latest SPECint rates by following the instructions to download a CSV file with the latest rates. After downloading the file, click **Browse** to locate the file in the local file system and click **Load** to update the rates in Consolidation Planner.

# Part VII

---

## Using the Cloud APIs

This section describes the cloud application programming interfaces (APIs) and the RESTful (Representational State Transfer) API.

It contains the following chapters:

- [Chapter 24, "Introduction to Cloud APIs"](#)
- [Chapter 25, "Cloud Resource Models"](#)
- [Chapter 26, "Infrastructure as a Service APIs"](#)
- [Chapter 27, "Database as a Service Family APIs"](#)
- [Chapter 28, "Java as a Service APIs"](#)
- [Chapter 29, "Chargeback and Metering EMCLI Verbs"](#)
- [Chapter 30, "SSA Administration APIs"](#)
- [Chapter 31, "Introduction to Blueprints"](#)
- [Chapter 32, "Cloud Blueprints and Blueprint Reference"](#)





---

## Introduction to Cloud APIs

Cloud computing is a style of computing in which dynamically scalable and deployed resources are provided as a service over the network. Users need not have knowledge of, expertise in, or control over the underlying infrastructure in the cloud that supports the services rendered to the users. As enterprises (companies, governments, and other organizations) integrate their existing IT infrastructures and IT resources with the sharable cloud paradigm, it is imperative for cloud enablers to provide a uniform API that these enterprises can use to tailor the cloud to their business processes and economic models.

The Representational State Transfer (RESTful) API presented here focuses on the resource models and their attributes.

This chapter contains the following sections:

- [Introduction](#)
- [When to Use Cloud APIs](#)
- [Change History](#)
- [Common Behaviors](#)

### 24.1 Introduction

As IT deployments become more complex, an abstraction of the infrastructure resources becomes more relevant to address concerns of compliance and configuration. Furthermore, such abstractions enable consumers to self serve the exact service they need and operationally control these services without any significant administrator involvement.

The RESTful API enables an infrastructure provider to service their customers by allowing them to perform the following:

- Browse templates that contain definitions and metadata of a logical unit of service.
- Deploy a template into the cloud and form an IT topology on demand.
- Perform operations (such as ONLINE, OFFLINE) on the resources.

The API is based on the HTTP protocol and GET, POST, PUT, and DELETE requests are all used. Resource representations documented here are in JavaScript Object Notation (JSON).

The API presupposes no particular structure in the Uniform Resource Identifier (URI) space. The starting point is a URI, supplied by the cloud service provider, that identifies the cloud itself. The cloud's representation contains URIs for the other

resources in the cloud. Operations on the cloud resources are performed by making an HTTP request against the URI of the resource.

The specification of this Cloud API includes the following:

- Common behaviors that apply across all requests and responses, error messages, and common resource attributes.
- Resource models, which describe the JSON data structures used in requests and responses.
- Requests that may be sent to cloud resources and the responses expected.

## 24.2 When to Use Cloud APIs

The Cloud APIs can be used to integrate Enterprise Manager with custom-built or 3rd party self service consoles and service desks. A few example scenarios include:

- Deploying databases and Oracle VM assemblies from custom-built self-service consoles.
- Integration into a bigger orchestration flow, such as a provisioning preceded with approval workflows.
- Deploying databases and Oracle VM assemblies from service desks.
- As a part of public clouds where the tenant interface of the cloud is typically different and serves service procurement, tenant registration, and so on.

## 24.3 Change History

The following table describes the high level changes to Enterprise Manager Cloud Control's Cloud Service Portal Plug-in 12.1.0.4 to 12.1.0.5.

**Table 24–1** *Changes to the Resource Model (from 12.1.0.4 to 12.1.0.5)*

| Abstract                                       | Description                                                                       | Impacts                                                                                         |
|------------------------------------------------|-----------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------|
| Availability of QuotaUsage resource.           | Service family type may provide resource quota usage information through the API. | Additional attributes are added to the respective resource models to reference the quota usage. |
| Availability of ServiceTemplateFinds resource. | Service family type may provide supports to query for Service Template resources. | Cloud level interactions return a list of service templates satisfying the listing criteria.    |

The following table describes the high level changes to Enterprise Manager Cloud Control's Cloud Service Portal Plug-in 12.1.0.2 to 12.1.0.4.

**Table 24–2 Changes to the Resource Model (from 12.1.0.2 to 12.1.0.4)**

| Abstract                                            | Description                                                                                                                                                                                                                                                                                                                                                                                                                | Impacts                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Introduction of x-specification-version header.     | <p>x-specification-version is a new header introduced in this release. The purpose of this header is to inform the backend server to switch between various supported versions of the Cloud Resource Model.</p> <p>By default, if this header is not specified, the server will automatically use the latest version of the resource model.</p> <p>Each server response will also now include x-specification-version.</p> | <p>If a client wishes to interact with the web service using the previous (pre-12.1.0.4 Cloud Service Portal release) resource model in 12.1.0.4 Cloud Service Portal, the following header can be specified.</p> <p>x-specification-version: 10000</p> <p>For example, GET /em/cloud with x-specification-version: 10000 will return the pre 12.1.0.4 release CloudServicePortal Cloud resource.</p>         |
| Cloud media type is updated to be service oriented. | Before this release, top level Cloud media type is IaaS centric. However, with the additions of Middleware as a Service in this release, the Cloud media type is transformed into a generic service model where additional XaaS can be plugged into the system without further remodel of the media type.                                                                                                                  | The Cloud media type is no longer backward compatible. (However, using the x-specification-version header can still allow a client to switch back to the older media type).                                                                                                                                                                                                                                   |
| Consolidate the concept of VDC into Zone.           | Resource relationships that are supported by the IaaS centric VDC are now consolidated into the generic Zone resource. The VDC resource is currently not visible from the 12.1.0.4 Cloud Service Portal resource model.                                                                                                                                                                                                    | Interacting with the 12.1.0.4 Cloud Service Portal resource model would be centric around Zone instead of VDC.                                                                                                                                                                                                                                                                                                |
| "media_type" attribute for resources.               | Each resource will have a "media_type" attribute to support client introspections in the generic service model. In other words, a Collection attribute in a resource may contain list of resources that are of different media type extending the same abstract media type.                                                                                                                                                | <p>The interaction with the web service may be introspective. For example, the Cloud resource contains "service_templates" attribute that describes the list of service templates that are from various XaaS.</p> <p>Each service template item in the "service_templates" may be of different media types, but these media types should contain all the attributes of the generic ServiceTemplate model.</p> |
| "Accept" header may be optional.                    | If the intention of the interaction is clear, you do not need to specify the "Accept" header.                                                                                                                                                                                                                                                                                                                              | The web service will introspect the resource URI automatically to return the resource with the most specific media type.                                                                                                                                                                                                                                                                                      |

**Table 24–2 (Cont.) Changes to the Resource Model (from 12.1.0.2 to 12.1.0.4)**

| Abstract                          | Description                                                                                                                                                                                                                                                                                                                                                                                                                  | Impacts                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|-----------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Generic and Specific media types. | <p>This version introduces the concept of Generic and Specific media types to support inheritance. For example, <code>application/oracle.com.cloud.common.ServiceTemplate</code> is a Generic media type and <code>application/oracle.com.cloud.common.VMTemplate</code> is a specific type that extends the above generic type.</p> <p>A specific type can extend another specific type (just like object inheritance).</p> | <p>It is now possible to use the Generic type to inquire about the resource. For example, if an URI represents a resource of specific type Y that extends generic type X, then, GET URI shall also accept the "accept" header of value X.</p> <p>Furthermore, attributes in specific type Y are expected to include the attributes in generic type X, and therefore clients shall be able to interact with the resource on the generic type.</p> |

## 24.4 Common Behaviors

The following sections specify constraints that apply to all the requests and responses that occur in the RESTful APIs supported by the Oracle Cloud Computing Platform.

### 24.4.1 Transport Protocol

All of the platform APIs are based on the Hypertext Transfer Protocol (HTTP), version 1.1 (RFC 2616). Each request will be authenticated using HTTP Basic Authentication (RFC 2617) unless otherwise noted. Therefore, requests sent from the public Internet (and not on a secure channel such as a VPN) must use the HTTPS protocol.

### 24.4.2 URI Space

The resources in the system are identified by URIs. To begin operations, the URI for a resource must be known. Dereferencing the URI yields a representation of the resource containing resource attributes and links to associated resources.

Assumptions about the layout of the URIs or the structure of resource URIs should not be made.

### 24.4.3 Media Types

In this specification, resource representations and request bodies are encoded in JavaScript Object Notation (JSON), as specified in RFC 4627. Each type of resource has its own media-type, which matches the following pattern:

```
application/oracle.com.cloud.common.Xxxxx+json
```

where `Xxxxx` represents the portion of the identifier unique to a particular representation format for each resource. The identifier must be globally unique in the space of `vnd.com.oracle.cloud`, and the media type should be registered in accordance to RFC 4288.

The Platform must provide representations of all resources available in JSON. The Platform must accept requests encoded in JSON.

## 24.4.4 Request Headers

In requests made to services implementing Oracle Cloud Platform APIs, several specific HTTP headers are used as described in the following table:

**Table 24–3 Request Headers**

| Header                              | Supported Values                                           | Description                                                                                  | Required                                                           |
|-------------------------------------|------------------------------------------------------------|----------------------------------------------------------------------------------------------|--------------------------------------------------------------------|
| Accept                              | Comma-delimited list of media types or media type patterns | Indicates to the server what media types are acceptable.                                     | Recommended on requests that will produce a response message body. |
| Authorization                       | “Basic” plus user name and password (per RFC 2617).        | Identifies the user making this request.                                                     | Yes on most of the requests.                                       |
| Content-Length                      | Length (in bytes) of the request message body.             | Describes the size of the message body.                                                      | Yes on requests that contain a message body.                       |
| Content-Type                        | Media type describing the request message body.            | Describes the representation and syntax of the request message body.                         | Yes on requests that contain a message body.                       |
| Host                                | Identifies the host receiving the message.                 | Required to allow support of multiple origin hosts at a single IP address.                   | All requests.                                                      |
| X-YYYYYClient-Specification-Version | String containing a specification version number.          | Declares the specification version of the YYYYY API that this client was programmed against. | No.                                                                |

## 24.4.5 Response Headers

Specific HTTP headers included in the responses returned from the platform are described in the following table:

**Table 24–4 Response Headers**

| Header       | Supported Values                                                                          | Description                                                                 | Required                                                                                  |
|--------------|-------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------|-------------------------------------------------------------------------------------------|
| Content-Type | Media type describing the response message body.                                          | Describes the representation and syntax of the response message body.       | Yes, on responses that contain a message body.                                            |
| Location     | Canonical URI of the resource, either newly created, or the original referenced resource. | Returns a URI that can be used to request a representation of the resource. | Yes, on responses to requests that create a new resource, or change an existing resource. |

**Table 24–4 (Cont.) Response Headers**

| Header        | Supported Values                                            | Description                                                                 | Required                                                                                                                                                                                                                                                                                     |
|---------------|-------------------------------------------------------------|-----------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Cache-Control | Max-age, public, no-store, mustrevalidate, proxyrevalidate. | How the representation of the resource should be cached, and its freshness. | No. For public resources (such as a list of public assemblies or templates) that do not change frequently, allowing lenient cache-control to optimize the response.<br><br>This will never be returned on a privileged resource or a resource request that contains an authorization header. |

## 24.4.6 HTTP Status Codes

Oracle Cloud Computing Platform APIs return standard HTTP response codes as described in the following table, under the conditions listed in the description.

**Table 24–5 HTTP Response Codes**

| Header           | Description                                                                                                                                                                                                                                                                                                                                                                |
|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 100 Continue     | The client should continue with its request. This interim response is used to inform the client that the initial part of the request has been received and has not yet been rejected by the platform. The client should continue by sending the remainder of the request or, if the request has already been completed, ignore this response.                              |
| 200 OK           | The request was successfully completed. If this request created a new resource that is addressable with a URI, and a response body is returned containing a representation of the new resource, a 200 status will be returned with a location header containing the canonical URI for the newly created resource.                                                          |
| 201 Created      | A request that created a new resource was completed and no response body containing a representation of the new resource is being returned. A location header containing the canonical URI for the newly created resource will be returned.                                                                                                                                |
| 202 Accepted     | The request has been accepted for processing, but the processing has not been completed. Per the HTTP/1.1 specification, the returned entity (if any) should include an indication of the request's current status. A location header containing the canonical URI for the not-yet completed resource is returned along with the status attribute indicating its progress. |
| 400 Bad Request  | The request could not be processed because it contains missing or invalid information (such as validation error on an input field, a missing required value, and so on).                                                                                                                                                                                                   |
| 401 Unauthorized | The authentication credentials included with this request are missing or invalid.                                                                                                                                                                                                                                                                                          |
| 403 Forbidden    | The server recognized your credentials, but you do not possess authorization to perform this request.                                                                                                                                                                                                                                                                      |
| 404 Not Found    | The request specified a URI of a resource that does not exist.                                                                                                                                                                                                                                                                                                             |

**Table 24–5 (Cont.) HTTP Response Codes**

| Header                    | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|---------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 405 Method Not Allowed    | The HTTP verb specified in the request (DELETE, GET, HEAD, POST, PUT) is not supported for this request URI.                                                                                                                                                                                                                                                                                                                                                                                  |
| 406 Not Acceptable        | The resource identified by this request is not capable of generating a representation corresponding to one of the media types in the Accept header of the request.                                                                                                                                                                                                                                                                                                                            |
| 409 Conflict              | A creation or update request could not be completed because it would cause a conflict in the current state of the resources supported by the platform. For example, an attempt is made to create a new resource with a unique identifier that is already assigned to some existing resource or an attempt is made to modify a resource attribute which is not yet completed.                                                                                                                  |
| 410 Gone                  | The requested resource is no longer available at the server and no forwarding address is known. This condition is expected to be considered permanent. Clients with link editing capabilities should delete references to the Request-URI after user approval.<br><br>If the server does not know, or has no facility to determine, whether or not the condition is permanent, the status code 404 (Not Found) should be used instead. This response is cacheable unless indicated otherwise. |
| 412 Precondition Failed   | The precondition given in one or more of the request-header fields evaluated to + when it was tested on the server. This response code allows the client to place preconditions on the current resource meta-information (header field data) and thus prevent the requested method from being applied to a resource other than the one intended.                                                                                                                                              |
| 500 Internal Server Error | The server encountered an unexpected condition which prevented it from fulfilling the request.                                                                                                                                                                                                                                                                                                                                                                                                |
| 501 Not Implemented       | The server does not (currently) support the functionality required to fulfill the request.                                                                                                                                                                                                                                                                                                                                                                                                    |
| 503 Service Unavailable   | The server is currently unable to handle the request due to temporary overloading or maintenance of the server.                                                                                                                                                                                                                                                                                                                                                                               |

## 24.4.7 Common Resource Attributes

All the resource entities in this specification may contain the following common resource attributes.

### 24.4.7.1 Resource State

This attribute denotes the state of the resource describing the lifecycle of the resource. This differs from the status of the entity represented by the resource which has entity specific semantics.

The following table shows the Data Model of this attribute.

**Table 24–6 ResourceState Data Model**

| Field | Type   | Occurs | Description                                                                                                                                                                                                                                                            |
|-------|--------|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| state | String | 1      | The current state of the resource as last known. This is a label containing lifecycle state (for example, INITIATED, CREATING, CREATED, DESTROYING, DESTROYED, READY).<br><br>When there are vendor extensions, the vendor shall publish and document their semantics. |

**Table 24–6 (Cont.) ResourceState Data Model**

| Field    | Type                      | Occurs | Description                                                                                                                |
|----------|---------------------------|--------|----------------------------------------------------------------------------------------------------------------------------|
| progress | Integer between 0 and 100 | 0..1   | This indicates the progress made as an approximate percentage. Not all state labels assign semantic meaning to this field. |
| messages | Message[]                 | 0..1   | Include the message data model instances to denote noteworthy communications.                                              |

## 24.4.8 Collection

This attribute is a meta resource that represents a collection field in a resource. For example, a VDC contains a collection of VMs and the field that represents the list of VMs is implemented in this type.

In the resource model, a collection field is denoted as Collection<type>, for example, Collection<VM>.

**Table 24–7 Collection<type> Data Model**

| Field    | Type     | Occurs | Description                                                                                                                                                                                                                                                                                      |
|----------|----------|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| uri      | URI      | 1      | URI that represents the collection of entities.                                                                                                                                                                                                                                                  |
| type     | String   | 1      | Type of the entity that this collection contains.                                                                                                                                                                                                                                                |
| total    | Integer  | 0..1   | Total number of elements that can be safely assumed to be in the elements list.                                                                                                                                                                                                                  |
| elements | <TYPE>[] | 0..1   | List of entities in this collection. The URI of the entities must be populated by the platform. When dereferencing the URI, the client must use the type field in the Accept header (except in the case where type = URI).<br><br>If this is not returned, then the collection is an empty list. |

In addition to the resource type, the collection also supports Collection<URI> where the type field is “URI”. This basic type collection requires additional type casting where the URI can be dereferenced properly.

This basic type collection requires additional type casting where the URI can be dereferenced properly. It is also permissible to have a collection of type collection, for example, Collection<Collection<Server>>.

## 24.4.9 Error Response Message

Successful requests will generally return an HTTP status code of 200 (OK), 201 (Created), 202 (Accepted), or 204 (No Content) to indicate that the requested action has been successfully performed or submitted.

In addition, a response message body (with an appropriate media type) containing a representation of the requested information might also be included. However, it is possible for a number of things to go wrong.

The various underlying causes are described by various HTTP status codes in the range 400-499 (for client side errors) or 500-599 (for server side problems).



If a response is returned with an error status code (400-499 or 500-599), the server will also return a response message body containing a message data model, with zero or more message data models, describing what went wrong. The text values of such messages might be used, for example, to communicate with a human user of the client side application.

The entire list of messages included in a single error response is encapsulated in a messages data model. The media type shall be returned in the Content-Type header. The client shall not include the Messages media type in the Accept header.

**Table 24–8 Messages Data Model**

| Field    | Type    | Occurs | Description                                            |
|----------|---------|--------|--------------------------------------------------------|
| messages | Message | 0..n   | Zero or more message data for each individual message. |

An individual message contains the following fields:

**Table 24–9 Individual Message Data Model**

| Field       | Type   | Occurs | Description                                                                                                                                                                        |
|-------------|--------|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| code        | String | 0..1   | Zero or more message data for each individual message.                                                                                                                             |
| field       | String | 0..1   | Name of the field from the request data model that this message is associated with.                                                                                                |
| hint        | String | 0..1   | Localized text further describing the nature of the problem, possibly including potential workarounds that the client could try.                                                   |
| text        | String | 1      | Localized text describing the nature of the problem reported by this message.                                                                                                      |
| severity    | String | 0..1   | Label indicating the severity of the error condition represented by this message.<br>Vendor shall publish the enumerators that are associated with this field and their semantics. |
| stack_trace | String | 0..1   | Vendor specific stack trace associated with this message.                                                                                                                          |
| source      | String | 0..1   | Symbolic identifier of the service implementation component that triggered this message.                                                                                           |



## Cloud Resource Models

This chapter specifies the representations of the resources that the API operates on. It contains the following sections:

- [Enterprise Manager Cloud Resource Model 10001](#)
- [About Cloud Resources](#)
- [Resource Data Models](#)

### 25.1 Enterprise Manager Cloud Resource Model 10001

This version of the Cloud resource model is the culmination of the XaaS implementation for Cloud Service Portal Plug-in, where additional services can be used to support the creation and lifecycle management of service instances.

For this release of Enterprise Manager Cloud Control's Cloud Service Portal Plug-in, 10001 is the default version when /em/cloud is accessed by the client. Specifically, the client can specify x-specification-version: 10001 in the HTTP header to examine whether or not a resource supports version 10001. It is important to note that the resource type from different versions may not be inter-mixed in a request. For example, 10001 may reject a request where the content is shaped like the 10000 version of the resource.

The following table contains the top level description of the resources:

**Table 25–1 Top Level Resource Descriptions**

| Resource                            | Description                                                                                                                                                                                                                                                                                                                                                                                   |
|-------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| application/oracle.com.cloud.common | This media type namespace contains "common" resources that are not specific to resource types in XaaS implementation. It is customary that XaaS implementation may extend a common resource type by responding to the request where the accept/content-type is the base type. This is similar to Java's class typing where one can always typecast the base class to an implementation class. |
| Cloud                               | This is the top level media type of the entry point /em/cloud where the client can get a high level view of the Cloud as a whole.                                                                                                                                                                                                                                                             |
| ServiceTemplate                     | This is the common media type that represents a service template each XaaS can extend to create a service instance.                                                                                                                                                                                                                                                                           |
| ServiceInstance                     | This represents the common media type that represents a service instance. Each XaaS can extend this media type to represent the service instances that it supports.                                                                                                                                                                                                                           |
| ServiceFamilyType                   | This represents XaaS or a service family.                                                                                                                                                                                                                                                                                                                                                     |

**Table 25–1 (Cont.) Top Level Resource Descriptions**

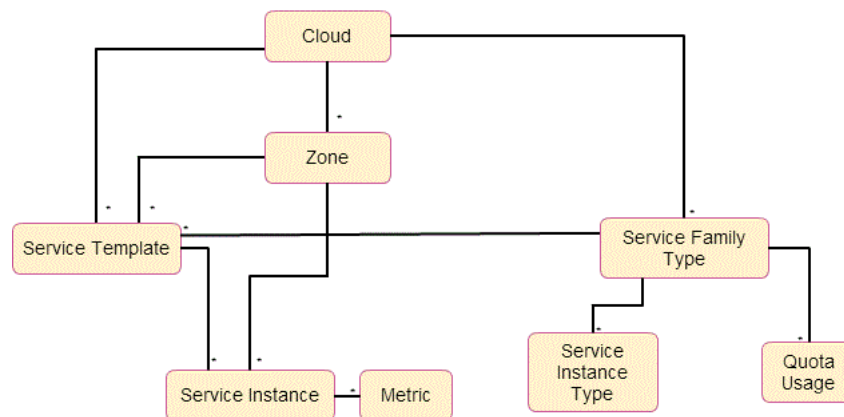
| Resource          | Description                                                                                                                                                                                                                                                                                            |
|-------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| InstanceType      | This represents a type of service instance. For example, VM is an extending media type of ServiceInstance and there is a corresponding InstanceType? that describes what a VM instance may contain.                                                                                                    |
| Metric            | This represents a metric of a resource. For example, performance metrics such as CPU utilization and network utilization are metrics of a VM instance. This does not need to be constrained to performance metrics only. Any measurable or observable time series data can be represented as a metric. |
| MultipartMetric   | This extends the Metric, but represents a multi-dimensional observation.                                                                                                                                                                                                                               |
| CloudInteractions | This represents the collection of interactions that are supported by the Cloud. The data of this resource would be protocol dependent. For example, in HTTP protocol, the data would consist of the supported HTTP request triple of (request method, accept type, and content type).                  |
| QuotaUsage        | This represents the quota information in the context of a service family.                                                                                                                                                                                                                              |

## 25.2 About Cloud Resources

Cloud resource representations are made up of fields, each with a name and value, encoded using a JavaScript Object Notation (JSON) dictionary. The values may be lists, dictionaries, or numeric or string literals, each of which is represented in JSON in accordance with RFC 4627.

Each type of cloud resource has its own Internet Media Type. The media type SHALL conform to the pattern `application/oracle.com.cloud.common.Xxxxxxxx+json`, and the specific media type for each resource model is included in square brackets in the corresponding section header.

Cloud resources are now organized by common service entities where specific services (for example, Infrastructure as a Service) provide resources extending common service entities where appropriate. The figure below shows the resource model relationships on the common resources and Infrastructure as a service specific resources.

**Figure 25–1 Cloud Resource Model**

## 25.3 Resource Data Models

The following sections provide details on the different cloud resource data models.

### 25.3.1 Cloud [application/oracle.com.cloud.common.Cloud+json]

A Cloud represents the user's starting view of all accessible resources. The following table describes the cloud data model.

**Table 25–2 Cloud Data Model**

| Field                | Type                          | Occurs | Description                                                                                                                                                                                                                                                                 | Since    |
|----------------------|-------------------------------|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| uri                  | URI                           | 1      | GET against this URI refreshes the client representation of the resources accessible to this user.                                                                                                                                                                          | 12.1.0.1 |
| name                 | String                        | 1      | Human readable name. It is a UNICODE string to support different languages.                                                                                                                                                                                                 | 12.1.0.1 |
| description          | String                        | 0..1   | Brief description. This is a UNICODE string to support different languages.                                                                                                                                                                                                 | 12.1.0.1 |
| service_templates    | Collection<ServiceTemplates>  | 0..1   | List of service templates that are accessible to the user.                                                                                                                                                                                                                  | 12.1.0.4 |
| service_family_types | Collection<ServiceFamilyType> | 0..1   | The list of service family types that are supported by the Cloud.                                                                                                                                                                                                           | 12.1.0.4 |
| zones                | Collection<Zone>              | 0..1   | List of zones that are supported by the cloud and accessible to the user.                                                                                                                                                                                                   | 12.1.0.2 |
| resource_state       | ResourceState                 | 0..1   | Cloud that is online and running would have READY as its state. If this field is not returned, the client can assume the cloud is READY. If the state of the returned field is not READY, the client cannot assume the viability of subsequent interactions into the cloud. | 12.1.0.1 |
| media_type           | String                        | 1      | Value of this media type with the payload format. For example, application/oracle.com.cloud.common.Cloud+json.                                                                                                                                                              | 12.1.0.4 |

### 25.3.2 ServiceTemplate [application/oracle.com.cloud.common.ServiceTemplate+json]

A ServiceTemplate represents the definition of the deployable service. For a user, a ServiceTemplate represents the definition of the deployable service. Users can create cloud resources by interacting with the URI of a ServiceTemplate. The cloud shall instantiate the resources and their configurations as specified in the definition of the ServiceTemplate.

VMTemplate is a subclass to the ServiceTemplate resource and therefore it is possible to get the ServiceTemplate from the URI of a VMTemplate. The following table shows the ServiceTemplate Data Model.

**Table 25–3 ServiceTemplate Data Model**

| Field                       | Type                          | Occurs | Description                                                                                                                                                                                                                       | Since    |
|-----------------------------|-------------------------------|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| uri                         | URI                           | 1      | GET against this URI refreshes the client representation of the ServiceTemplate definition to this user.                                                                                                                          | 12.1.0.4 |
| name                        | String                        | 1      | Human readable name given to the ServiceTemplate. It is a UNICODE string to support different languages.                                                                                                                          | 12.1.0.4 |
| description                 | String                        | 0..1   | A brief description given to the ServiceTemplate. It is a UNICODE string to support different languages.                                                                                                                          | 12.1.0.4 |
| type                        | String                        | 1      | The string that describes the type of the service template.                                                                                                                                                                       | 12.1.0.4 |
| created                     | Timestamp                     | 1      | Date and time, in ISO 8601 format, when the ServiceTemplate was created.                                                                                                                                                          | 12.1.0.4 |
| service_family_type         | String                        | 0..1   | The name of the service family type under which this service template is categorized.                                                                                                                                             | 12.1.0.4 |
| default_instance_media_type | String                        | 0..1   | The default media type of the service instance that can be created using this template.<br><br>For example, a VMTemplate would have VM for this attribute to indicate that VMTemplate can be deployed into a VM service instance. | 12.1.0.4 |
| resource_state              | ResourceState                 | 1      | Only a service template with a READY state can be deployed.                                                                                                                                                                       | 12.1.0.4 |
| deployment_params           | List of Deployment Parameters | 0..1   | Contains the list of data structure of Deployment Parameters to indicate the parameters that may be specified during the service instance deployment using this template                                                          | 12.1.0.4 |
| zones                       | List of Zones                 | 0..1   | Contains the list of Zone resources that this service template can be used to create service instances with                                                                                                                       | 12.1.0.4 |
| service_instances           | Collection<? >                | 0..1   | Contains the list of service instances that are created with this template.                                                                                                                                                       | 12.1.0.4 |

The following table describes the structure of the Deployment Parameters.

**Table 25–4 Deployment Parameter Structure**

| Field       | Type   | Occurs | Description                                                             |
|-------------|--------|--------|-------------------------------------------------------------------------|
| name        | String | 1      | The name of the parameter that needs to be specified during deployment. |
| description | String | 0..1   | A brief description of the deployment parameter.                        |

**Table 25–4 (Cont.) Deployment Parameter Structure**

| Field         | Type    | Occurs | Description                                                                                                            |
|---------------|---------|--------|------------------------------------------------------------------------------------------------------------------------|
| type          | String  | 1      | The type of the deployment parameter. It is an enumeration of the following values, STRING, INTEGER, NUMBER, and LIST. |
| default_value | String  | 0..1   | The default value for this parameter.                                                                                  |
| required      | Boolean | 1      | TRUE if this parameter is required, FALSE otherwise.                                                                   |
| sensitive     | Boolean | 1      | TRUE if this parameter denotes a password. FALSE otherwise.                                                            |

### 25.3.3 Zone [application/oracle.com.cloud.common.Zone+json]

A Zone represents a logical boundary where the resources may reside. A Zone can represent a particular geographically location such as Europe Zone, North America Zone, East Asia Zone. A zone can also represent characteristics, such as high network bandwidth or DMZ secured. Furthermore, a Zone can be organizational in nature, such as Financial Department Zone, Testing Zone, and Development Zone.

There should not be any assumption of exclusivity of underlying infrastructures in the Zones unless otherwise noted. For example, Zone A and Zone B can be on the same physical network serving two different departments, but their physical infrastructure setup is transparent to cloud users.

**Table 25–5 Zone Data Model**

| Field               | Type                         | Occurs | Description                                                                                                                           | Since    |
|---------------------|------------------------------|--------|---------------------------------------------------------------------------------------------------------------------------------------|----------|
| uri                 | URI                          | 1      | GET against this URI refreshes the client representation of the Zone definition to this user.                                         | 12.1.0.1 |
| name                | String                       | 1      | Name of the Zone.                                                                                                                     | 12.1.0.1 |
| description         | String                       | 0..1   | Human readable description of the Zone. It is a UNICODE string to support different languages.                                        | 12.1.0.1 |
| media_type          | String                       | 1      | The media type of the resource.                                                                                                       | 12.1.0.4 |
| resource_state      | ResourceState                | 0..1   | The resource state of the resource.                                                                                                   | 12.1.0.1 |
| service_family_type | String                       | 0..1   | The name of the service family type that this zone is associated with.<br><br>Each Zone is contextualized in the service family type. | 12.1.0.4 |
| service_templates   | Collection <ServiceTemplate> | 0..1   | Collection of the service templates that this zone supports and can be deployed into.                                                 | 12.1.0.4 |
| service_instances   | Collection <ServiceInstance> | 0..1   | Collection of the service instances that are in this zone.                                                                            | 12.1.0.4 |

### 25.3.4 Service Family Type

#### [application/oracle.com.cloud.common.ServiceFamilyType+json]

A service family type is a category of services that are offered by the cloud. For example, IaaS is a category that encapsulates infrastructure services. Mwaas is a category that encapsulates middleware services. These categories are predefined.

The following table describes the ServiceFamilyType Data Model

**Table 25–6 Service Family Type**

| Field             | Type                        | Occurs | Description                                                                                                                                                                                               | Since    |
|-------------------|-----------------------------|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| uri               | URI                         | 1      | A GET against this URI refreshes the client representation of the ServiceFamilyType definition to this user.                                                                                              | 12.1.0.4 |
| name              | String                      | 1      | Name of the ServiceFamilyType.                                                                                                                                                                            | 12.1.0.4 |
| description       | String                      | 0..1   | Human readable description of the Service Family Type. Shall be a UNICODE string to support different languages.                                                                                          | 12.1.0.4 |
| resource_state    | ResourceState               | 0..1   | The resource state of the resource.                                                                                                                                                                       | 12.1.0.4 |
| media_type        | String                      | 1      | The media type of the resource.                                                                                                                                                                           | 12.1.0.4 |
| service_templates | Collection<ServiceTemplate> | 0..1   | The collection of service templates that are of this service family type.                                                                                                                                 | 12.1.0.4 |
| zones             | Collection<Zone>            | 0..1   | The collection of zones that support service instances of this service family type.                                                                                                                       | 12.1.0.4 |
| instance_types    | Collection<InstanceType>    | 0..1   | The collection of instance types that are supported by this service family type.                                                                                                                          | 12.1.0.4 |
| quota_usages      | Collection<QuotaUsage>      | 0..1   | The collection of quota usages of the perspective family type of the authenticated user.                                                                                                                  | 12.1.0.5 |
| quota_def         | Json Object                 | 0..1   | The family type specific quota metadata information, this may include some permission information.<br><br>The structure of this object would be documented in the perspective family type resource model. | 12.1.0.5 |

### 25.3.5 Service Instance Type

#### [application/oracle.com.cloud.common.InstanceType+json]

A service instance type describes the common metadata about service instances of the type. This can be interpreted analogously as a Class where the service instance of the type is an instantiation of the class.

The following table describes the ServiceFamilyType Data Model.



**Table 25–7 Service Instance Type**

| Field               | Type               | Occurs | Description                                                                                                                                             | Since    |
|---------------------|--------------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| uri                 | URI                | 1      | A GET against this URI refreshes the client representation of the ServiceInstanceType definition to this user.                                          | 12.1.0.4 |
| name                | String             | 1      | Name of the ServiceInstanceType.                                                                                                                        | 12.1.0.4 |
| description         | String             | 0..1   | Human readable description of the Service Instance Type. Shall be a UNICODE string to support different languages.                                      | 12.1.0.4 |
| resource_state      | ResourceState      | 0..1   | The resource state of the resource.                                                                                                                     | 12.1.0.4 |
| media_type          | String             | 1      | The media type of the resource.                                                                                                                         | 12.1.0.4 |
| instance_media_type | String             | 0..1   | The media type of the instance of this type.                                                                                                            | 12.1.0.4 |
| metrics             | List<List<String>> | 0..1   | List of metrics that may be supported by the instance type. Each list element represents the triple of "name", "description", and "type" of the metric. | 12.1.0.4 |

### 25.3.6 Metric [application/oracle.com.cloud.common.Metric+json]

This resource represents a time series data that may be performance or configuration related. For example, CPU utilization could be a metric for a VM resource, Tablespace usages could be a metric for a Database resource.

The following table describes the Metric Data Model.

**Table 25–8 Metric Data Model**

| Field          | Type          | Occurs | Description                                                                                         | Since    |
|----------------|---------------|--------|-----------------------------------------------------------------------------------------------------|----------|
| uri            | URI           | 1      | A GET against this URI refreshes the client representation of the Metric definition to this user.   | 12.1.0.4 |
| name           | String        | 1      | Name of the Metric.                                                                                 | 12.1.0.4 |
| description    | String        | 0..1   | Human readable description of the Metric. SHALL be a UNICODE string to support different languages. | 12.1.0.4 |
| resource_state | ResourceState | 0..1   | The resource state of the resource.                                                                 | 12.1.0.4 |
| media_type     | String        | 1      | The media type of the resource.                                                                     | 12.1.0.4 |
| type           | String        | 1      | The type of the metric value. This is an enumeration of STRING, INTEGER, PERCENTAGE, NUMBER, MIX.   | 12.1.0.4 |
| current_value  | String        | 0..1   | The latest known value of the metric                                                                | 12.1.0.4 |

**Table 25–8 (Cont.) Metric Data Model**

| Field            | Type                     | Occurs | Description                                                                                                                                                                                                                                                                                                                              | Since    |
|------------------|--------------------------|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| current_time     | ISO 8601 Date            | 0..1   | The UTC time when the current_value was last observed                                                                                                                                                                                                                                                                                    | 12.1.0.4 |
| time_range       | List of 2 ISO 8601 Dates | 0..1   | The first one is begin date, and the second one is end date. This range determines the time series window of "values" attribute                                                                                                                                                                                                          | 12.1.0.4 |
| rollup_unit      | String                   | 0..1   | The rollup unit for the time series data. Could be one of DAY, DAILY, DAYLY, HOUR, HOURLY, RAW                                                                                                                                                                                                                                           | 12.1.0.4 |
| values           | List of objects          | 0..1   | For rollup_unit = RAW, this would be a list of object each containing "time_utc", and "value" attributes to indicate a timed observation.<br><br>For all other rollup_unit, this would be a list of object each containing "time_utc", "average", "min", "max", "std" attributes to indicate an aggregated observations where supported. | 12.1.0.4 |
| time_range_epoch | List of 2 Number         | 0..1   | Same as time_range, but in the form of epoch time.                                                                                                                                                                                                                                                                                       | 12.1.0.4 |

### 25.3.7 Service Instance [application/oracle.com.cloud.common.ServiceInstance+json]

A service instance describes the some common metadata about service instances. This is an abstract media type where the actual implementation would be provided by each XaaS services. It is expected that the extending resources respect GET request with this abstract media type.

The following table describes the Service Instance Data Model.

**Table 25–9 Service Instance Data Model**

| Field          | Type                | Occurs | Description                                                                                                   | Since    |
|----------------|---------------------|--------|---------------------------------------------------------------------------------------------------------------|----------|
| uri            | URI                 | 1      | A GET against this URI refreshes the client representation of the ServiceInstance definition to this user.    | 12.1.0.4 |
| name           | String              | 1      | Name of the ServiceInstance.                                                                                  | 12.1.0.4 |
| description    | String              | 0..1   | Human readable description of the Service Instance. Shall be a UNICODE string to support different languages. | 12.1.0.4 |
| resource_state | ResourceState       | 1      | The resource state of the resource.                                                                           | 12.1.0.4 |
| media_type     | String              | 1      | The media type of the resource.                                                                               | 12.1.0.4 |
| metrics        | Collection <Metric> | 0..1   | Collection of metrics that are observed on the resource.                                                      | 12.1.0.4 |

**Table 25–9 (Cont.) Service Instance Data Model**

| Field               | Type          | Occurs | Description                                                                                                                                                                                                                                                                      | Since    |
|---------------------|---------------|--------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| service_family_type | String        | 1      | The name of the ServiceFamilyType that this service instance is in context of.                                                                                                                                                                                                   | 12.1.0.4 |
| status              | String        | 0..1   | The status of the entity represented by the service instance. For example, for service instance VM, the value could be STARTED, STOPPED, or any other values that are appropriate for VM. Each extending resource should describe the enumeration of status that are applicable. | 12.1.0.4 |
| created             | ISO 8601 Date | 0..1   | The date of which the service instance is created.                                                                                                                                                                                                                               | 12.1.0.4 |

### 25.3.8 Quota Usage [application/oracle.com.cloud.common.QuotaUsage+json]

A quota usage describes the amount of some resource used by the authentication used in the context of a service family type. For example, in IaaS, there were quota defined for CPU, Memory, number of servers, where the usages of each is encapsulated into a Quota Usage resource.

The following table describes the Quota Usage Data Model.

**Table 25–10 QuotaUsage Data Model**

| Field       | Type   | Occurs | Description                                                                                                   | Since    |
|-------------|--------|--------|---------------------------------------------------------------------------------------------------------------|----------|
| uri         | URI    | 1      | A GET against this URI refreshes the client representation of the QuotaUsage                                  | 12.1.0.5 |
| name        | String | 1      | Name of the QuotaUsage.                                                                                       | 12.1.0.5 |
| description | String | 0..1   | Human readable description of the Service Instance. Shall be a UNICODE string to support different languages. | 12.1.0.5 |
| media_type  | String | 1      | The media type of the resource.                                                                               | 12.1.0.5 |
| id          | String | 1      | The identification representation of the resource.                                                            | 12.1.0.5 |
| type        | String | 1      | The service family type of the resource.                                                                      | 12.1.0.5 |
| used        | Number | 1      | How much of the resource is used. The type is represented by the "unit" attribute.                            | 12.1.0.5 |
| maxAllowed  | Number | 0..1   | What is the maximum amount of resource that may be used by the authenticated user.                            | 12.1.0.5 |
| unit        | String | 1      | The unit of the resource.                                                                                     | 12.1.0.5 |

## 25.3.9 Service Template Finds

### [application/oracle.com.cloud.common.ServiceTemplateFinds+json]

Service Template Finds is a resource that would search and filter service templates under the entire cloud given the search criteria. Each service family type shall document the scope of support (for example, list of filter attributes).

The following table describes the Service Template Finds Data Model.

**Table 25–11 ServiceTemplateFinds Data Model**

| Field       | Type              | Occurs | Description                                                                                                                                                                                                      | Since    |
|-------------|-------------------|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| uri         | URI               | 1      | A GET against this URI refreshes the client representation of the Service TemplateFinds.<br><br>If service templates were added/removed that matched the criteria, refreshing the link will reflect the changes. | 12.1.0.5 |
| name        | String            | 1      | Name of the ServiceTemplateFinds.                                                                                                                                                                                | 12.1.0.5 |
| description | String            | 0..1   | Human readable description of the Service Instance. Shall be a UNICODE string to support different languages.                                                                                                    | 12.1.0.5 |
| filters     | Json Object       | 1      | Map of key-value pair to indicate the filter criteria.                                                                                                                                                           | 12.1.0.5 |
| finds       | Collection<br><*> | 1      | .Collection of entities (could be any sub type of the service template) that satisfies the filter criteria.                                                                                                      | 12.1.0.5 |

The following table describes the attributes supported in the filters:

**Table 25–12 Filter Attributes**

| Attributes | Description                                                                                                                                          | Since                                                                 |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------|
| media_type | The specific service template media to search for. This will be an exact match of the media type.                                                    | 12.1.0.5 Cloud Service Portal Plug-in (IaaS service family type only) |
| name       | A service template will be returned if the name contains the specified string. This is case insensitive, so OEL and oel will return the same result. | 12.1.0.5 Cloud Service Portal Plug-in (IaaS service family type only) |

## 25.4 Cloud API Examples

The following sections provide examples of different interactions in Enterprise Manager Cloud Control 12.1.0.2 with the 12.1.0.5 Cloud Service Portal Plug-in. These examples illustrate the shape of the various resources. In this version, default x-specification-version is assumed at 10001.

### 25.4.1 Cloud Resource

Use the top level /em/cloud to introspect the Cloud resource. The following table describes the GET method features:

**Table 25–13 Cloud Resource GET Method**

| Feature | Description                                                                     |
|---------|---------------------------------------------------------------------------------|
| URL     | https://example.oracle.com/em/cloud                                             |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI<br>=X-specification-Version: 10001 |
| Method  | GET                                                                             |
| Since   | 12.1.0.4 Cloud Service Portal Plug-in                                           |

The following returns the Cloud resource:

```
{
 "uri" : "/em/cloud" ,
 "name" : "Oracle Cloud by Enterprise Manager" ,
 "description" : "This represents the Cloud resource of the Oracle Enterprise
Manager Cloud Management solution" ,
 "resource_state" : {
 "state" : "READY"
 } ,
 "media_type" : "application/oracle.com.cloud.common.Cloud+json" ,
 "zones" : {
 "media_type" : "application/oracle.com.cloud.common.Zone+json" ,
 "total" : "3" ,
 "elements" :
 [
 {
 "uri" : "/em/cloud/jaas/zone/450121204B4703131FC0FDD72E7BF494" ,
 "name" : "PaaS_Zone_Middleware" ,
 "media_type" : "application/oracle.com.cloud.jaas.Zone+json" ,
 "service_family_type" : "jaas"
 } ,
 {
 "uri" : "/em/cloud/opc/opczone" ,
 "name" : "OPC Zone" ,
 "media_type" : "application/oracle.com.cloud.opc.OpcZone+json" ,
 "service_family_type" : "opc" ,
 "type" : "opc"
 } ,
 {
 "uri" : "/em/cloud/iaas/zone/D73AF0D42C8459E11419862797D1F37D" ,
 "name" : "cloud_zone" ,
 "media_type" : "application/oracle.com.cloud.iaas.Zone+json" ,
 "service_family_type" : "iaas" ,
 "type" : "VM_ZONE"
 }
]
 } ,
 "service_templates" : {
 "media_type" : "application/oracle.com.cloud.common.ServiceTemplate+json" ,
 "total" : "7" ,
 "elements" :
 [
 {
 "uri" :
"/em/cloud/iaas/servicetemplate/vm/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1
%3Acmp%3AVirtualization%3ATemplate%3AC76CEB5563EA5E13E040578CDA817FAF%3A0.1" ,
 "name" : "template_sanity" ,
 "media_type" : "application/oracle.com.cloud.common.VMTemplate+json" ,
```

```
 "service_family_type" : "iaas" ,
 "type" : "Template"
 } ,
 {
 "uri" :
"/em/cloud/iaas/servicetemplate/assembly/oracle%3AdefaultService%3Aem%3Aprovisioni
ng%3A1%3Acmp%3AVirtualization%3AAssembly%3AC75E88B04D7FEDEDE040578CDA810E49%3A0.1"
,
 "name" : "sidb_assembly" ,
 "media_type" :
"application/oracle.com.cloud.common.AssemblyTemplate+json" ,
 "service_family_type" : "iaas" ,
 "type" : "Assembly"
 } ,
 {
 "uri" :
"/em/cloud/iaas/servicetemplate/assembly/oracle%3AdefaultService%3Aem%3Aprovisioni
ng%3A1%3Acmp%3AVirtualization%3AAssembly%3AC769B1F361529309E040578CDA813D57%3A0.1"
,
 "name" : "wls_assembly" ,
 "media_type" :
"application/oracle.com.cloud.common.AssemblyTemplate+json" ,
 "service_family_type" : "iaas" ,
 "type" : "Assembly"
 } ,
 {
 "uri" :
"/em/cloud/iaas/servicetemplate/assembly/oracle%3AdefaultService%3Aem%3Aprovisioni
ng%3A1%3Acmp%3AVirtualization%3AAssembly%3AC76F733BC7A41AF7E040578CDA812CDC%3A0.1"
,
 "name" : "fmw_venkat" ,
 "media_type" :
"application/oracle.com.cloud.common.AssemblyTemplate+json" ,
 "service_family_type" : "iaas" ,
 "type" : "Assembly"
 } ,
 {
 "uri" :
"/em/cloud/iaas/servicetemplate/assembly/oracle%3AdefaultService%3Aem%3Aprovisioni
ng%3A1%3Acmp%3AVirtualization%3AAssembly%3AC76C144A4A245B62E040578CDA8163B9%3A0.1"
,
 "name" : "fmw_abby" ,
 "media_type" :
"application/oracle.com.cloud.common.AssemblyTemplate+json" ,
 "service_family_type" : "iaas" ,
 "type" : "Assembly"
 } ,
 {
 "uri" :
"/em/cloud/iaas/servicetemplate/assembly/oracle%3AdefaultService%3Aem%3Aprovisioni
ng%3A1%3Acmp%3AVirtualization%3AAssembly%3AC77122B0A916D95CE040578CDA814854%3A0.1"
,
 "name" : "sidbasmA_abby" ,
 "media_type" :
"application/oracle.com.cloud.common.AssemblyTemplate+json" ,
 "service_family_type" : "iaas" ,
 "type" : "Assembly"
 } ,
 {
 "uri" :
```

```

"/em/cloud/iaas/servicetemplate/assembly/oracle%3AdefaultService%3Aem%3Aprovisioni
ng%3A1%3Acmp%3AVirtualization%3AAssembly%3AC76C8792DE2A0937E040578CDA81795E%3A0.1"
,
 "name" : "WLS_abby" ,
 "media_type" :
"application/oracle.com.cloud.common.AssemblyTemplate+json" ,
 "service_family_type" : "iaas" ,
 "type" : "Assembly"
}
]
} ,
"service_family_types" : {
 "media_type" : "application/oracle.com.cloud.common.ServiceFamilyType+json"
,
 "total" : "4" ,
 "elements" :
 [
 {
 "uri" : "/em/cloud/service_family_type/jaas" ,
 "name" : "jaas" ,
 "media_type" :
"application/oracle.com.cloud.common.ServiceFamilyType+json" ,
 "type" : "jaas"
 } ,
 {
 "uri" : "/em/cloud/service_family_type/dbaas" ,
 "name" : "dbaas" ,
 "media_type" :
"application/oracle.com.cloud.common.ServiceFamilyType+json" ,
 "type" : "dbaas"
 } ,
 {
 "uri" : "/em/cloud/service_family_type/opc" ,
 "name" : "opc" ,
 "media_type" :
"application/oracle.com.cloud.common.ServiceFamilyType+json" ,
 "type" : "opc"
 } ,
 {
 "uri" : "/em/cloud/service_family_type/iaas" ,
 "name" : "iaas" ,
 "media_type" :
"application/oracle.com.cloud.iaas.IaasServiceFamilyType+json" ,
 "type" : "iaas"
 }
]
}
}

```

---

**Note:** even though the Accept type was not specified, the Web service still unambiguously returns the Cloud resource because `/em/cloud` address uniquely identify the Cloud as the default resource to be returned.

---

The following table shows the Cloud resource picking up specific attributes

**Table 25–14 Cloud Resource GET Method**

| Feature | Description                                                   |
|---------|---------------------------------------------------------------|
| URL     | https://example.oracle.com/em/cloud?service_family_types,name |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=                 |
| Method  | GET                                                           |
| Since   | 12.1.0.4 Cloud Service Portal Plug-in                         |

Here are the selective attributes:

```
{
 "name" : "Oracle Cloud by Enterprise Manager" ,
 "service_family_types" : {
 "media_type" : "application/oracle.com.cloud.common.ServiceFamilyType+json"
 } ,
 "total" : "4" ,
 "elements" : [
 {
 "uri" : "/em/cloud/service_family_type/jaas" ,
 "name" : "jaas" ,
 "media_type" :
"application/oracle.com.cloud.common.ServiceFamilyType+json" ,
 "type" : "jaas"
 } ,
 {
 "uri" : "/em/cloud/service_family_type/dbaas" ,
 "name" : "dbaas" ,
 "media_type" :
"application/oracle.com.cloud.common.ServiceFamilyType+json" ,
 "type" : "dbaas"
 } ,
 {
 "uri" : "/em/cloud/service_family_type/opc" ,
 "name" : "opc" ,
 "media_type" :
"application/oracle.com.cloud.common.ServiceFamilyType+json" ,
 "type" : "opc"
 } ,
 {
 "uri" : "/em/cloud/service_family_type/iaas" ,
 "name" : "iaas" ,
 "media_type" :
"application/oracle.com.cloud.iaas.IaasServiceFamilyType+json" ,
 "type" : "iaas"
 }
]
}
```

This feature is supported on all the common resources where only selective attributes would be gathered and returned to the client.

## 25.4.2 Service Family Type Resource

The following table shows the features of the Service Family Type resource:



**Table 25–15 Service Family Type Resource**

| Feature | Description                                                  |
|---------|--------------------------------------------------------------|
| URL     | https://example.oracle.com/em/cloud/service_family_type/iaas |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=                |
| Method  | GET                                                          |
| Since   | 12.1.0.4 Cloud Service Portal Plug-in                        |

Here is the returned headers and content for this interaction.

```
X-specification-version: 10001
Content-Type: application/oracle.com.cloud.iaas.IaaSServiceFamilyType+json;
charset=ISO-8859-1
{
 "uri" : "/em/cloud/service_family_type/iaas" ,
 "name" : "iaas" ,
 "resource_state" : {
 "state" : "READY"
 } ,
 "media_type" : "application/oracle.com.cloud.iaas.IaaSServiceFamilyType+json" ,
 "type" : "iaas" ,
 "service_templates" : {
 "media_type" : "application/oracle.com.cloud.common.ServiceTemplate+json" ,
 "total" : "7" ,
 "elements" :
 [
 {
 "uri" :
"/em/cloud/iaas/servicetemplate/vm/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp%3AVirtualization%3ATemplate%3AC76CEB5563EA5E13E040578CDA817FAF%3A0.1" ,
 "name" : "template_sanity" ,
 "media_type" : "application/oracle.com.cloud.common.VMTemplate+json" ,
 "type" : "Template"
 } ,
 {
 "uri" :
"/em/cloud/iaas/servicetemplate/assembly/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp%3AVirtualization%3AAssembly%3AC75E88B04D7FEDEDE040578CDA810E49%3A0.1"
,
 "name" : "sidb_assembly" ,
 "media_type" :
"application/oracle.com.cloud.common.AssemblyTemplate+json" ,
 "type" : "Assembly"
 } ,
 {
 "uri" :
"/em/cloud/iaas/servicetemplate/assembly/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp%3AVirtualization%3AAssembly%3AC769B1F361529309E040578CDA813D57%3A0.1"
,
 "name" : "wls_assembly" ,
 "media_type" :
"application/oracle.com.cloud.common.AssemblyTemplate+json" ,
 "type" : "Assembly"
 } ,
 {
 "uri" :
"/em/cloud/iaas/servicetemplate/assembly/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp%3AVirtualization%3AAssembly%3AC76F733BC7A41AF7E040578CDA812CDC%3A0.1"

```

```
,
 "name" : "fmw_venkat" ,
 "media_type" :
"application/oracle.com.cloud.common.AssemblyTemplate+json" ,
 "type" : "Assembly"
} ,
{
 "uri" :
"/em/cloud/iaas/servicetemplate/assembly/oracle%3AdefaultService%3Aem%3Aprovisioni
ng%3A1%3Acmp%3AVirtualization%3AAssembly%3AC76C144A4A245B62E040578CDA8163B9%3A0.1"
,
 "name" : "fmw_abby" ,
 "media_type" :
"application/oracle.com.cloud.common.AssemblyTemplate+json" ,
 "type" : "Assembly"
} ,
{
 "uri" :
"/em/cloud/iaas/servicetemplate/assembly/oracle%3AdefaultService%3Aem%3Aprovisioni
ng%3A1%3Acmp%3AVirtualization%3AAssembly%3AC77122B0A916D95CE040578CDA814854%3A0.1"
,
 "name" : "sidbasmA_abby" ,
 "media_type" :
"application/oracle.com.cloud.common.AssemblyTemplate+json" ,
 "type" : "Assembly"
} ,
{
 "uri" :
"/em/cloud/iaas/servicetemplate/assembly/oracle%3AdefaultService%3Aem%3Aprovisioni
ng%3A1%3Acmp%3AVirtualization%3AAssembly%3AC76C8792DE2A0937E040578CDA81795E%3A0.1"
,
 "name" : "WLS_abby" ,
 "media_type" :
"application/oracle.com.cloud.common.AssemblyTemplate+json" ,
 "type" : "Assembly"
}
} ,
"zones" : {
 "media_type" : "application/oracle.com.cloud.common.Zone+json" ,
 "total" : "1" ,
 "elements" :
[
 {
 "uri" : "/em/cloud/iaas/zone/D73AF0D42C8459E11419862797D1F37D" ,
 "name" : "cloud_zone" ,
 "media_type" : "application/oracle.com.cloud.iaas.Zone+json" ,
 "service_family_type" : "iaas" ,
 "type" : "VM_ZONE"
 }
]
} ,
"instance_types" : {
 "media_type" : "application/oracle.com.cloud.common.InstanceType+json" ,
 "total" : "2" ,
 "elements" :
[
 {
 "uri" : "/em/cloud/instance_type/vm%40iaas" ,
 "name" : "vm" ,
```

```

 "media_type" : "application/oracle.com.cloud.common.InstanceType+json"
 ,
 "instance_media_type" : "application/oracle.com.cloud.common.VM+json"
} ,
{
 "uri" : "/em/cloud/instance_type/assembly%40iaas" ,
 "name" : "assembly" ,
 "media_type" : "application/oracle.com.cloud.common.InstanceType+json"
,
 "instance_media_type" :
"application/oracle.com.cloud.common.AssemblyInstance+json"
}
]
} ,
"quota_usages" : {
 "media_type" : "application/oracle.com.cloud.common.QuotaUsage+json" ,
 "total" : "6" ,
 "elements" :
 [
 {
 "uri" : "/em/cloud/quota_usage/iaas_3" ,
 "name" : "Local Storage Quota" ,
 "description" : "Local Storage Quota" ,
 "context_id" : "3" ,
 "media_type" : "application/oracle.com.cloud.common.QuotaUsage+json" ,
 "used" : "10240" ,
 "maxAllowed" : "2000" ,
 "unit" : "GB"
 } ,
 {
 "uri" : "/em/cloud/quota_usage/iaas_2" ,
 "name" : "RAM Quota" ,
 "description" : "RAM Quota" ,
 "context_id" : "2" ,
 "media_type" : "application/oracle.com.cloud.common.QuotaUsage+json" ,
 "used" : "512" ,
 "maxAllowed" : "600000" ,
 "unit" : "MB"
 } ,
 {
 "uri" : "/em/cloud/quota_usage/iaas_1" ,
 "name" : "CPU Quota" ,
 "description" : "CPU Quota" ,
 "context_id" : "1" ,
 "media_type" : "application/oracle.com.cloud.common.QuotaUsage+json" ,
 "used" : "1" ,
 "maxAllowed" : "100" ,
 "unit" : "COUNT"
 } ,
 {
 "uri" : "/em/cloud/quota_usage/iaas_6" ,
 "name" : "Servers Quota" ,
 "description" : "Servers Quota" ,
 "context_id" : "6" ,
 "media_type" : "application/oracle.com.cloud.common.QuotaUsage+json" ,
 "used" : "1" ,
 "maxAllowed" : "20" ,
 "unit" : "COUNT"
 }
]
}

```

```
 "uri" : "/em/cloud/quota_usage/iaas_5" ,
 "name" : "SWLIB Storage Quota" ,
 "description" : "SWLIB Storage Quota" ,
 "context_id" : "5" ,
 "media_type" : "application/oracle.com.cloud.common.QuotaUsage+json" ,
 "used" : "0" ,
 "maxAllowed" : "-1" ,
 "unit" : "GB"
 } ,
 {
 "uri" : "/em/cloud/quota_usage/iaas_4" ,
 "name" : "Extra Storage Quota" ,
 "description" : "Extra Storage Quota" ,
 "context_id" : "4" ,
 "media_type" : "application/oracle.com.cloud.common.QuotaUsage+json" ,
 "used" : "0" ,
 "maxAllowed" : "-1" ,
 "unit" : "GB"
 }
]
} ,
"quota_def" : {
 "uri" : "/em/cloud/iaas/quota" ,
 "name" : "iaas Quota" ,
 "description" : "iaas Quota" ,
 "media_type" : "application/oracle.com.cloud.iaas.Quota+json"
} ,
"instance_options" :
[
 {
 "name" : "Medium" ,
 "cpu" : "4" ,
 "memory" : "8192" ,
 "local_storage" : "512000" ,
 "id" : "2"
 } ,
 {
 "name" : "Large" ,
 "cpu" : "8" ,
 "memory" : "15360" ,
 "local_storage" : "1024000" ,
 "id" : "3"
 } ,
 {
 "name" : "Small" ,
 "cpu" : "2" ,
 "memory" : "4096" ,
 "local_storage" : "256000" ,
 "id" : "1"
 }
]
}
```

The `application/oracle.com.cloud.iaas.IaaSServiceFamilyType+json` media type is returned automatically as the most detailed resource. Similarly, you can interact with the system by accepting the common type:

**Table 25–16 Service Family Type Resource**

| Feature | Description                                                                                                              |
|---------|--------------------------------------------------------------------------------------------------------------------------|
| URL     | https://example.oracle.com/em/cloud/service_family_type/iaas                                                             |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=<br>Accept-Type: application/oracle.com.cloud.common.ServiceFamilyType+json |
| Method  | GET                                                                                                                      |

The Web service will return the content as shown in the previous interaction.

### 25.4.3 Quota Resource

The following describes the quota resource that is exposed through the IaaS service family type to introspect the quota definition for IaaS services.

**Table 25–17 Service Family Type Resource**

| Feature | Description                                    |
|---------|------------------------------------------------|
| URL     | https://example.oracle.com/em/cloud/iaas/quota |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=  |
| Method  | GET                                            |
| Since   | 12.1.0.5 Cloud Service Portal Plug-in          |

The following example shows the returned headers and content for this interaction:

```
X-specification-version: 10001
Content-Type: application/oracle.com.cloud.iaas.Quota+json; charset=ISO-8859-1
{
 "uri" : "/em/cloud/iaas/quota" ,
 "name" : "iaas Quota" ,
 "description" : "iaas Quota" ,
 "resource_state" : {
 "state" : "READY"
 } ,
 "media_type" : "application/oracle.com.cloud.iaas.Quota+json" ,
 "service_family_type" : "iaas" ,
 "allowedSaveToSwlib" : "1" ,
 "allowedSaveToSwlibOnExpiry" : "0" ,
 "resourceQuotas" : {
 "total" : "6" ,
 "elements" :
 [
 {
 "name" : "Local Storage Quota" ,
 "description" : "Local Storage Quota" ,
 "maxAllowed" : "2000" ,
 "unit" : "GB"
 } ,
 {
 "name" : "RAM Quota" ,
 "description" : "RAM Quota" ,
 "maxAllowed" : "600000" ,
 "unit" : "MB"
 }
]
 }
}
```

```

 "name" : "CPU Quota" ,
 "description" : "CPU Quota" ,
 "maxAllowed" : "100" ,
 "unit" : "COUNT"
 } ,
 {
 "name" : "Servers Quota" ,
 "description" : "Servers Quota" ,
 "maxAllowed" : "20" ,
 "unit" : "COUNT"
 } ,
 {
 "name" : "SWLIB Storage Quota" ,
 "description" : "SWLIB Storage Quota" ,
 "maxAllowed" : "-1" ,
 "unit" : "GB"
 } ,
 {
 "name" : "Extra Storage Quota" ,
 "description" : "Extra Storage Quota" ,
 "maxAllowed" : "-1" ,
 "unit" : "GB"
 }
]
}

```

## 25.4.4 Service Instance Type Resource

This is the resource that describes the type of services that are provided by a particular Service Family Type.

**Table 25–18 Service Instance Type Resource**

| Feature | Description                                                                                                                           |
|---------|---------------------------------------------------------------------------------------------------------------------------------------|
| URL     | <a href="https://example.oracle.com/em/cloud/instance_type/vm%40iaas">https://example.oracle.com/em/cloud/instance_type/vm%40iaas</a> |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=                                                                                         |
| Method  | GET                                                                                                                                   |
| Since   | 12.1.0.4 Cloud Service Portal Plug-in                                                                                                 |

The following is the Service Instance Type resource:

```

{
 "uri" : "/em/cloud/instance_type/vm%40iaas" ,
 "name" : "vm" ,
 "description" : "A Virtual Machine represents a computational unit that contains CPU, Memory, Network Instances, and Disks" ,
 "media_type" : "application/oracle.com.cloud.common.InstanceType+json" ,
 "instance_media_type" : "application/oracle.com.cloud.common.VM+json" ,
 "metrics" :
 [
 [
 "CPU_Utilization" ,
 "CPU Utilization of the Server" ,
 "GENERIC"
] ,
 [
 "Total_Network_Throughput" ,

```

```

 "Total Network Throughput of the Server" ,
 "GENERIC"
] ,
 [
 "Total_Disk_Throughput" ,
 "Total Disk Throughput of the Server" ,
 "GENERIC"
] ,
 [
 "Filesystem_Total_Used" ,
 "Summary of the total file system usage" ,
 "GENERIC"
]
] ,
"instance_options" :
[
 {
 "name" : "Medium" ,
 "cpu" : "4" ,
 "memory" : "8192" ,
 "local_storage" : "512000" ,
 "id" : "2"
 } ,
 {
 "name" : "Large" ,
 "cpu" : "8" ,
 "memory" : "15360" ,
 "local_storage" : "1024000" ,
 "id" : "3"
 } ,
 {
 "name" : "Small" ,
 "cpu" : "2" ,
 "memory" : "4096" ,
 "local_storage" : "256000" ,
 "id" : "1"
 }
]
}

```

## 25.4.5 Zone Resource

The following table shows the features of the Zone resource:

**Table 25–19 ZoneResource**

| Feature | Description                                                                                                                                                                 |
|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| URL     | <a href="https://example.oracle.com/em/cloud/iaas/zone/D73AF0D42C8459E11419862797D1F37D">https://example.oracle.com/em/cloud/iaas/zone/D73AF0D42C8459E11419862797D1F37D</a> |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=                                                                                                                               |
| Method  | GET                                                                                                                                                                         |
| Since   | 12.1.0.4 Cloud Service Portal Plug-in                                                                                                                                       |

The following is the Zone resource, automatically cast to the application/oracle.com.cloud.iaas.Zone+json media type

```

{
 "uri" : "/em/cloud/iaas/zone/D73AF0D42C8459E11419862797D1F37D" ,

```

```
"name" : "cloud_zone" ,
"resource_state" : {
 "state" : "READY"
} ,
"context_id" : "D73AF0D42C8459E11419862797D1F37D" ,
"media_type" : "application/oracle.com.cloud.iaas.Zone+json" ,
"service_family_type" : "iaas" ,
"type" : "VM_ZONE" ,
"service_templates" : {
 "media_type" : "application/oracle.com.cloud.common.ServiceTemplate+json" ,
 "total" : "7" ,
 "elements" :
 [
 {
 "uri" :
"/em/cloud/iaas/servicetemplate/vm/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp%3AVirtualization%3ATemplate%3AC76CEB5563EA5E13E040578CDA817FAF%3A0.1" ,
 "name" : "template_sanity" ,
 "media_type" : "application/oracle.com.cloud.common.VMTemplate+json" ,
 "type" : "Template"
 } ,
 {
 "uri" :
"/em/cloud/iaas/servicetemplate/assembly/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp%3AVirtualization%3AAssembly%3AC75E88B04D7FEDEDE040578CDA810E49%3A0.1"
 ,
 "name" : "sidb_assembly" ,
 "media_type" :
"application/oracle.com.cloud.common.AssemblyTemplate+json" ,
 "type" : "Assembly"
 } ,
 {
 "uri" :
"/em/cloud/iaas/servicetemplate/assembly/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp%3AVirtualization%3AAssembly%3AC77122B0A916D95CE040578CDA814854%3A0.1"
 ,
 "name" : "sidbasmA_abby" ,
 "media_type" :
"application/oracle.com.cloud.common.AssemblyTemplate+json" ,
 "type" : "Assembly"
 } ,
 {
 "uri" :
"/em/cloud/iaas/servicetemplate/assembly/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp%3AVirtualization%3AAssembly%3AC769B1F361529309E040578CDA813D57%3A0.1"
 ,
 "name" : "wls_assembly" ,
 "media_type" :
"application/oracle.com.cloud.common.AssemblyTemplate+json" ,
 "type" : "Assembly"
 } ,
 {
 "uri" :
"/em/cloud/iaas/servicetemplate/assembly/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp%3AVirtualization%3AAssembly%3AC76C144A4A245B62E040578CDA8163B9%3A0.1"
 ,
 "name" : "fmw_abby" ,
 "media_type" :
"application/oracle.com.cloud.common.AssemblyTemplate+json" ,
 "type" : "Assembly"
 }
]
}
```



```

 },
 {
 "uri" :
"/em/cloud/iaas/servicetemplate/assembly/oracle%3AdefaultService%3Aem%3Aprovisioni
ng%3A1%3Acmp%3AVirtualization%3AAssembly%3AC76C8792DE2A0937E040578CDA81795E%3A0.1"
,
 "name" : "WLS_abby" ,
 "media_type" :
"application/oracle.com.cloud.common.AssemblyTemplate+json" ,
 "type" : "Assembly"
 } ,
 {
 "uri" :
"/em/cloud/iaas/servicetemplate/assembly/oracle%3AdefaultService%3Aem%3Aprovisioni
ng%3A1%3Acmp%3AVirtualization%3AAssembly%3AC76F733BC7A41AF7E040578CDA812CDC%3A0.1"
,
 "name" : "fmw_venkat" ,
 "media_type" :
"application/oracle.com.cloud.common.AssemblyTemplate+json" ,
 "type" : "Assembly"
 }
]
},
"service_instances" : {
 "media_type" : "application/oracle.com.cloud.common.ServiceInstance+json" ,
 "total" : "0" ,
 "elements" :
[
]
} ,
"vnets" : {
 "media_type" : "application/oracle.com.cloud.common.VNet+json" ,
 "total" : "3" ,
 "elements" :
[
 {
 "uri" :
"/em/cloud/iaas/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3AnetConfig%3AC76CEADBBBE6B23FE040578CDA817FB1" ,
 "context_id" :
"oracle:defaultService:em:provisioning:1:netConfig:C76CEADBBBE6B23FE040578CDA817FB
1" ,
 "media_type" : "application/oracle.com.cloud.common.VNet+json" ,
 "id" :
"oracle:defaultService:em:provisioning:1:netConfig:C76CEADBBBE6B23FE040578CDA817FB
1"
 } ,
 {
 "uri" :
"/em/cloud/iaas/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3AnetConfig%3AC77076C8FDEC6BD7E040578CDA813B2B" ,
 "context_id" :
"oracle:defaultService:em:provisioning:1:netConfig:C77076C8FDEC6BD7E040578CDA813B2
B" ,
 "media_type" : "application/oracle.com.cloud.common.VNet+json" ,
 "id" :
"oracle:defaultService:em:provisioning:1:netConfig:C77076C8FDEC6BD7E040578CDA813B2
B"
 } ,
 {
 "uri" :

```

```

"/em/cloud/iaas/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3AnetConfig%3AC76F741AFD7EB760E040578CDA812CD8" ,
 "context_id" :
"/em/cloud/iaas/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3AnetConfig%3AC76F741AFD7EB760E040578CDA812CD8" ,
 "media_type" : "application/oracle.com.cloud.common.VNet+json" ,
 "id" :
"/em/cloud/iaas/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3AnetConfig%3AC76F741AFD7EB760E040578CDA812CD8"
}
]
} ,
"delegated_credentials" : {
 "media_type" :
"application/oracle.com.cloud.iaas.ExalogicZoneCredential+json" ,
 "total" : "0" ,
 "elements" :
[
]
}
}

```

## 25.4.6 Service Template Resource

As can be seen from the previous examples, the elements in the collection shown all have "media\_type" as an attribute to clearly indicate what its resource type is. The following table shows the features of the Service Template resource:

**Table 25–20 Service Template Resource**

| Feature | Description                                                                                                                                                                                         |
|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| URL     | <code>https://example.oracle.com/em/cloud/iaas/servicetemplate/vm/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp%3AVirtualization%3ATemplate%3AC76CEB5563EA5E13E040578CDA817FAF%3A0.1</code> |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=                                                                                                                                                       |
| Method  | GET                                                                                                                                                                                                 |
| Since   | 12.1.0.4 Cloud Service Portal Plug-in                                                                                                                                                               |

This is the resource returned as the complete VMTemplate resource.

```

{
 "uri" :
"/em/cloud/iaas/servicetemplate/vm/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp%3AVirtualization%3ATemplate%3AC76CEB5563EA5E13E040578CDA817FAF%3A0.1" ,
 "name" : "template_sanity" ,
 "description" : "test template" ,
 "resource_state" : {
 "state" : "READY"
 } ,
 "context_id" :
"/em/cloud/iaas/servicetemplate/vm/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp%3AVirtualization%3ATemplate%3AC76CEB5563EA5E13E040578CDA817FAF%3A0.1" ,
 "media_type" : "application/oracle.com.cloud.common.VMTemplate+json" ,
 "type" : "Template" ,
 "created" : "2012-08-17T11:53:26+0000" ,
 "default_instance_media_type" : "application/oracle.com.cloud.common.VM+json" ,
 "zones" : {
 "media_type" : "application/oracle.com.cloud.iaas.Zone+json" ,
 "total" : "1" ,

```

```

 "elements" :
 [
 {
 "uri" : "/em/cloud/iaas/zone/D73AF0D42C8459E11419862797D1F37D" ,
 "name" : "cloud_zone" ,
 "media_type" : "application/oracle.com.cloud.iaas.Zone+json" ,
 "service_family_type" : "iaas" ,
 "type" : "VM_ZONE"
 }
]
} ,
"service_instances" : {
 "media_type" : "application/oracle.com.cloud.common.ServiceInstance+json" ,
 "total" : "0" ,
 "elements" :
 [
]
} ,
"deployment_params" :
[
 {
 "name" : "zone" ,
 "description" : "'zone' attribute is the URI of the zone that the
'application/oracle.com.cloud.common.VM' resource is to be created in" ,
 "type" : "STRING" ,
 "require" : "false" ,
 "sensitive" : "false"
 } ,
 {
 "name" : "based_on" ,
 "description" : "'based_on' attribute is the URI of the service template
or format String of the originated source which the deployed
'application/oracle.com.cloud.common.VM' resource is to be followed" ,
 "type" : "STRING" ,
 "require" : "false" ,
 "sensitive" : "false"
 } ,
 {
 "name" : "cpu" ,
 "description" : "The information that determined how much 'CPU' is to be
allocated to the deployed resource" ,
 "type" : "LIST" ,
 "require" : "false" ,
 "sensitive" : "false"
 } ,
 {
 "name" : "memory" ,
 "description" : "The information that determined how much 'MEMORY' is to
be allocated to the deployed resource" ,
 "type" : "NUMBER" ,
 "require" : "false" ,
 "sensitive" : "false"
 } ,
 {
 "name" : "disks" ,
 "description" : "The list of disks to be included in the deployed
resource" ,
 "type" : "LIST" ,
 "require" : "false" ,
 "sensitive" : "false"
 } ,
]

```

```
{
 "name" : "params.domain_type" ,
 "description" : "The domain type of the Virtual Machine" ,
 "type" : "STRING" ,
 "require" : "false" ,
 "sensitive" : "false"
} ,
{
 "name" : "params.network_profile" ,
 "description" : "The identifier of the network profile to be used for the
network instances of the Virtual Machine" ,
 "type" : "STRING" ,
 "require" : "false" ,
 "sensitive" : "false"
} ,
{
 "name" : "params.server_prefix" ,
 "description" : "The server prefix to be specified for the Virtual
Machine" ,
 "type" : "STRING" ,
 "require" : "true" ,
 "sensitive" : "false"
} ,
{
 "name" : "params.server_size" ,
 "description" : "The name of the server instance size" ,
 "type" : "STRING" ,
 "require" : "false" ,
 "sensitive" : "false"
} ,
{
 "name" : "params.request_name" ,
 "description" : "The request name to be tracked for the Virtual Machine
creation" ,
 "type" : "STRING" ,
 "require" : "false" ,
 "sensitive" : "false"
} ,
{
 "name" : "params.request_description" ,
 "description" : "The request description to be tracked for the Virtual
Machine creation" ,
 "type" : "STRING" ,
 "require" : "false" ,
 "sensitive" : "false"
} ,
{
 "name" : "params.vnc_password" ,
 "description" : "The VNC password for the Virtual Machine" ,
 "type" : "STRING" ,
 "require" : "true" ,
 "sensitive" : "true"
} ,
{
 "name" : "params.root_password" ,
 "description" : "The Root password for the Virtual Machine" ,
 "type" : "STRING" ,
 "require" : "true" ,
 "sensitive" : "true"
} ,
}
```

```

 {
 "name" : "params.start_vm" ,
 "description" : "Whether or not the Virtual Machine should be started
after creation" ,
 "type" : "STRING" ,
 "defaultValue" : "YES" ,
 "require" : "false" ,
 "sensitive" : "false"
 } ,
 {
 "name" : "params.ha_enabled" ,
 "description" : "Whether or not the Virtual Machine should be HA enabled"
,
 "type" : "STRING" ,
 "defaultValue" : "NO" ,
 "require" : "false" ,
 "sensitive" : "false"
 } ,
 {
 "name" : "params.disks_overwrite" ,
 "description" : "Whether or not the default Virtual Machine disks should
be overwritten" ,
 "type" : "STRING" ,
 "defaultValue" : "NO" ,
 "require" : "false" ,
 "sensitive" : "false"
 } ,
 {
 "name" : "params.networks" ,
 "description" : "Whether or not the default Virtual Machine disks should
be overwritten" ,
 "type" : "LIST" ,
 "require" : "false" ,
 "sensitive" : "false"
 }
] ,
 "directory" : "COMP_Components" ,
 "updated" : "2012-08-17T11:53:26+0000" ,
 "creator" : "DMBHAT1" ,
 "last_modified_by" : "DMBHAT1" ,
 "componenttype" : "Virtualization" ,
 "oracle_owned" : "false"
}

```

## 25.4.7 Metric Resource

A metric contains time series information and depending on the type of the metric, rollup information can also be retrieved. For example,

**Table 25–21 Service Template Resource**

| Feature | Description                                                                                                                                                                                                                       |
|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| URL     | <a href="https://example.oracle.com/em/cloud/metric/iaas%3Avm%3ACPU_Utilization%3A523CAE80A305928C9C5BE8A67A4181FD">https://example.oracle.com/em/cloud/metric/iaas%3Avm%3ACPU_Utilization%3A523CAE80A305928C9C5BE8A67A4181FD</a> |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=                                                                                                                                                                                     |
| Method  | GET                                                                                                                                                                                                                               |
| Since   | 12.1.0.4 Cloud Service Portal Plug-in                                                                                                                                                                                             |

Here is the response from the web service:

```
{
 "uri" : "/em/cloud/metric/iaas%3Avm%3ACPU_
Utilization%3A523CAE80A305928C9C5BE8A67A4181FD" ,
 "name" : "CPU_Utilization" ,
 "media_type" : "application/oracle.com.cloud.common.Metric+json" ,
 "resource_state" : {
 "state" : "READY"
 } ,
 "type" : "NUMBER" ,
 "current_value" : "52.53868103027344" ,
 "current_time" : "2012-08-22T01:28:14+0000" ,
 "time_range" :
 [
 "2012-08-22T01:18:14+0000" ,
 "2012-08-22T01:28:14+0000"
] ,
 "time_range_epoch" :
 [
 "1345598294965" ,
 "1345598894965"
] ,
 "rollup_unit" : "RAW" ,
 "values" :
 [
 {
 "time_utc" : "2012-08-22T01:21:44+0000" ,
 "value" : "52.53868103027344"
 }
] ,
 "key" : "iaas:vm:CPU_Utilization:523CAE80A305928C9C5BE8A67A4181FD"
}
```

The default time range is 10 minutes. To modify the time range, query parameters can be used to specify the time range. For example, the following would get the metric data for the last 50 minutes:

**Table 25–22    Service Template Resource**

| Feature | Description                                                                                                                                            |
|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| URL     | https://example.oracle.com/em/cloud/metric/iaas%3Avm%3ACPU_Utilization%3A523CAE80A305928C9C5BE8A67A4181FD?time_range_epoch=1345595894965~1345598894965 |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=                                                                                                          |
| Method  | GET                                                                                                                                                    |
| Since   | 12.1.0.4 Cloud Service Portal Plug-in                                                                                                                  |

The query parameter "time\_range\_epoch" contains 2 numbers separated by "~" to indicate FROM time and TO time, inclusive. Here is the response:

```
{
 "uri" : "/em/cloud/metric/iaas%3Avm%3ACPU_
Utilization%3A523CAE80A305928C9C5BE8A67A4181FD" ,
 "name" : "CPU_Utilization" ,
 "media_type" : "application/oracle.com.cloud.common.Metric+json" ,
 "resource_state" : {
 "state" : "READY"
 }
}
```

```

 } ,
 "type" : "NUMBER" ,
 "current_value" : "52.53868103027344" ,
 "current_time" : "2012-08-22T01:29:19+0000" ,
 "time_range" :
 [
 "2012-08-22T00:38:14+0000" ,
 "2012-08-22T01:28:14+0000"
] ,
 "time_range_epoch" :
 [
 "1345595894965" ,
 "1345598894965"
] ,
 "rollup_unit" : "RAW" ,
 "values" :
 [
 {
 "time_utc" : "2012-08-22T01:21:44+0000" ,
 "value" : "52.53868103027344"
 } ,
 {
 "time_utc" : "2012-08-22T01:16:44+0000" ,
 "value" : "52.67461395263672"
 } ,
 {
 "time_utc" : "2012-08-22T01:11:44+0000" ,
 "value" : "51.93536376953125"
 } ,
 {
 "time_utc" : "2012-08-22T01:06:44+0000" ,
 "value" : "52.61101531982422"
 } ,
 {
 "time_utc" : "2012-08-22T01:01:44+0000" ,
 "value" : "52.30859375"
 } ,
 {
 "time_utc" : "2012-08-22T00:56:44+0000" ,
 "value" : "52.778690338134766"
 } ,
 {
 "time_utc" : "2012-08-22T00:51:44+0000" ,
 "value" : "52.08208084106445"
 } ,
 {
 "time_utc" : "2012-08-22T00:46:44+0000" ,
 "value" : "52.42387771606445"
 } ,
 {
 "time_utc" : "2012-08-22T00:41:44+0000" ,
 "value" : "52.2672004699707"
 }
] ,
 "key" : "iaas:vm:CPU_Utilization:523CAE80A305928C9C5BE8A67A4181FD"
}

```

Similarly, to specify a different rollup unit, the following interaction shows the HOURLY rollup for the last 7 hours:

**Table 25–23 Service Template Resource**

| Feature | Description                                                                                                                                                               |
|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| URL     | https://example.oracle.com/em/cloud/metric/iaas%3Avm%3ACPU_Utilization%3A523CAE80A305928C9C5BE8A67A4181FD?time_range_epoch=1345571178252~1345611178252&rollup_unit=HOURLY |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=                                                                                                                             |
| Method  | GET                                                                                                                                                                       |
| Since   | 12.1.0.4 Cloud Service Portal Plug-in                                                                                                                                     |

Here is the response from the web service for the aggregated information:

```
{
 "uri" : "/em/cloud/metric/iaas%3Avm%3ACPU_Utilization%3A523CAE80A305928C9C5BE8A67A4181FD" ,
 "name" : "CPU_Utilization" ,
 "media_type" : "application/oracle.com.cloud.common.Metric+json" ,
 "resource_state" : {
 "state" : "READY"
 } ,
 "type" : "NUMBER" ,
 "current_value" : "52.37758255004883" ,
 "current_time" : "2012-08-22T04:55:13+0000" ,
 "time_range" :
 [
 "2012-08-21T17:46:18+0000" ,
 "2012-08-22T04:52:58+0000"
] ,
 "time_range_epoch" :
 [
 "1345571178252" ,
 "1345611178252"
] ,
 "rollup_unit" : "HOUR" ,
 "values" :
 [
 {
 "time_utc" : "2012-08-21T23:00:00+0000" ,
 "average" : "52.6733585993448905" ,
 "min" : "52.130245208740234" ,
 "max" : "53.29022216796875" ,
 "std" : "3.55263697371071501307481053785512321532E-01"
 } ,
 {
 "time_utc" : "2012-08-21T22:00:00+0000" ,
 "average" : "5.24303614298502599166666666666666666666667E01" ,
 "min" : "51.621360778808594" ,
 "max" : "52.85725402832031" ,
 "std" : "3.63521555326025399658498002900543252046E-01"
 } ,
 {
 "time_utc" : "2012-08-21T21:00:00+0000" ,
 "average" : "52.6410032908121745" ,
 "min" : "52.04186248779297" ,
 "max" : "53.23821258544922" ,
 "std" : "3.21552680564547999623645642115362106482E-01"
 }
]
}
```



}



## Infrastructure as a Service APIs

This chapter describes the resource models for Infrastructure as a Service (IaaS) API and the REST request/response interactions. The following topics are covered:

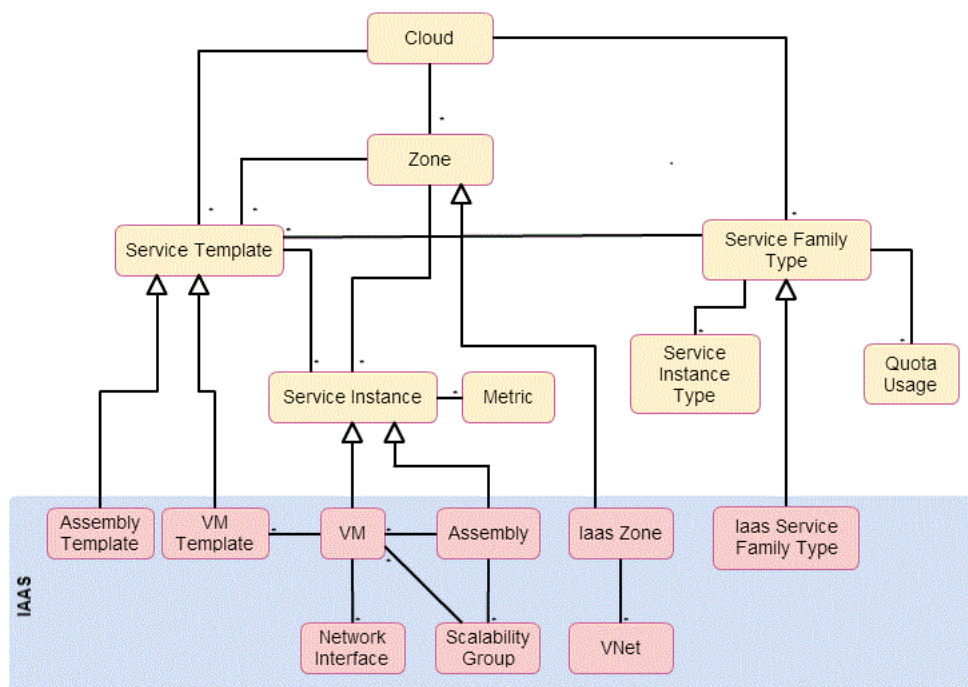
- [Resource Models for Infrastructure as a Service](#)
- [Supported Operations Examples](#)

### 26.1 Resource Models for Infrastructure as a Service

Infrastructure as a Service supports the compute, network, and storage resources that form the basics of IT components. The IaaS resource model supports common cloud interactions on these resources.

The following diagram illustrates the high level resource relationships and how they are related to the common resources:

**Figure 26–1 IaaS Resource Model**



The following describes the resource models that are supported by the IaaS family type.

### 26.1.1 Changes for the Infrastructure as a Service

The following describes the high level updates to the resources for Infrastructure as a Service:

**Table 26–1 Cloud Service Portal 12.1.0.5**

| Abstract                                                         | Description                                                                                                                       |
|------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------|
| Supports VMTemplate and AssemblyTemplate in ServiceTemplateFinds | Infrastructure as a service provides implementation for the filtering and searching of VMTemplate and AssemblyTemplate resources. |
| Supports cloning a VM                                            | An VM resource can be created from another VM directly.                                                                           |
| Supports saving a VM as a template                               | An VMTemplate resource can be created from a VM resource.                                                                         |
| Supports Infrastructure Quota                                    | Provides quota information and quota usage for Infrastructure as a service resources.                                             |

### 26.1.2 IaaS Zone [application/oracle.com.cloud.iaas.Zone]

This extends application/oracle.com.cloud.common.Zone. This resource represents the view of an IaaS Zone that supports IaaS resources.

**Table 26–2 IaaS Zone Data Model**

| Field               | Type                         | Occurs | Description                                                                                                                                                                          | Since    |
|---------------------|------------------------------|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| uri                 | URI                          | 1      | A GET against this URI refreshes the client representation of the Zone definition to this user.                                                                                      | 12.1.0.4 |
| name                | String                       | 1      | Name of the Zone. SHALL be a UNICODE string to support different languages                                                                                                           | 12.1.0.4 |
| description         | String                       | 0..1   | Human readable description of the Zone. SHALL be a UNICODE string to support different languages.<br><br>The media type of the resource.<br>"application/oracle.com.cloud.iaas.Zone" | 12.1.0.4 |
| resource_state      | ResourceState                | 0..1   | The resource state of the resource.                                                                                                                                                  | 12.1.0.4 |
| service_family_type | String                       | 0..1   | The name of the service family type that this zone is associated with.<br><br>It is "iaas" for this.                                                                                 | 12.1.0.4 |
| service_templates   | Collection <ServiceTemplate> | 0..1   | Collection of the service templates that this zone supports and can be deployed into. Currently, VMTemplate and AssemblyTemplate are types of the members.                           | 12.1.0.4 |
| service_instances   | Collection <ServiceInstance> | 0..1   | Collection of the service instances that are in this zone.                                                                                                                           | 12.1.0.4 |
| vnets               | Collection <VNet>            | 0..1   | Collection of the virtual networks that are associated with this zone.                                                                                                               | 12.1.0.4 |

**Table 26–2 (Cont.) IaaS Zone Data Model**

| Field      | Type   | Occurs | Description                                                                 | Since    |
|------------|--------|--------|-----------------------------------------------------------------------------|----------|
| media_type | String | 1      | The media type of the resource.<br>"application/oracle.com.cloud.iaas.Zone" | 12.1.0.4 |

### 26.1.3 AssemblyInstance

[application/oracle.com.cloud.common.AssemblyInstance+json]

Extends: application/oracle.com.cloud.common.ServiceInstance

An AssemblyInstance is a logical grouping of resources from a deployment request of an Assembly template. The lifecycle of the resources in an AssemblyInstance can be managed centrally through the assembly instance.

---

**Note:** A service template of Assembly type is deployed into an AssemblyInstance.

---

The following table describes the AssemblyInstance Data Model.

**Table 26–3 AssemblyInstance Data Model**

| Field              | Type                         | Occurs | Description                                                                                                                                                                                                                                                                                                                                                                     | Since    |
|--------------------|------------------------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| uri                | URI                          | 1      | GET against this URI refreshes the client representation of the Assembly Instance definition.                                                                                                                                                                                                                                                                                   | 12.1.0.2 |
| name               | String                       | 1      | Name of the Assembly Instance as given by the user. It is a UNICODE string to support different languages.                                                                                                                                                                                                                                                                      | 12.1.0.2 |
| description        | String                       | 0..1   | Human readable description of the Assembly Instance given by the user. It is a UNICODE string to support different languages.                                                                                                                                                                                                                                                   | 12.1.0.2 |
| based_on           | URI                          | 0..1   | URI of the service template of which this AssemblyInstance is based on.                                                                                                                                                                                                                                                                                                         | 12.1.0.2 |
| scalability_groups | Collection<ScalabilityGroup> | 0..1   | List of scalability groups that are included in this Assembly Instance.                                                                                                                                                                                                                                                                                                         | 12.1.0.2 |
| servers            | Collection<VM>               | 0..1   | List of VMs that are directly included in this Assembly Instance.                                                                                                                                                                                                                                                                                                               | 12.1.0.2 |
| status             | String                       | 0..1   | Indicates the status of the Assembly Instance. This field contains the semantics that the service provider implements. For example, a service provider may implement an ONLINE status to indicate that all the entities, recursively, are in an ONLINE status. Or a service provider may implement an ONLINE status to indicate that critical entities are in an ONLINE status. | 12.1.0.2 |
| resource_state     | ResourceState                | 1      | Validity of the other Assembly Instance fields on a GET should be guaranteed only when the resource state is READY. Otherwise, the client should not assume the validity of the fields.                                                                                                                                                                                         | 12.1.0.1 |

**Table 26–3 (Cont.) AssemblyInstance Data Model**

| Field           | Type      | Occurs | Description                                                                                                                                                                                                                                                                                                                           | Since    |
|-----------------|-----------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| created         | Timestamp | 0..1   | Date and time, in ISO 8601 format, when the Assembly Instance was created.                                                                                                                                                                                                                                                            | 12.1.0.2 |
| expiry          | Timestamp | 0..1   | Date and time, in ISO 8601 format, when the Assembly Instance should expire. If not specified, the Assembly Instance never expires.                                                                                                                                                                                                   | 12.1.0.2 |
| params          | {}        | 0..1   | Vendor specific configuration parameters for this deployment.                                                                                                                                                                                                                                                                         | 12.1.0.2 |
| deployment_plan | String    | 0..1   | Vendor specific configuration parameters for this deployment in the prescribed XML format from the corresponding AssemblyTemplate that this AssemblyInstance is based on. When this is specified in the POST request, the content of this attribute will be processed to substitute the parameters for deploying an AssemblyInstance. | 12.1.0.2 |
| media_type      | String    | 1      | The media type of the resource.                                                                                                                                                                                                                                                                                                       | 12.1.0.4 |

#### 26.1.4 Scalability Group [application/oracle.com.cloud.common.ScalabilityGroup+json]

Extends: application/oracle.com.cloud.common.ServiceInstance

The Scalability Group is a collection of servers and corresponding virtual networks. Scalability Group contains a collection of homogenous entities. The cloud service provider should enforce the semantics of “sameness”. Operations such as scale\_out and scale\_in of the scalability group may be supported by the provider in the homogenous scalability group.

The following table describes the ScalabilityGroup Data Model.

**Table 26–4 ScalabilityGroup Data Model**

| Field       | Type             | Occurs | Description                                                                                                                   | Since    |
|-------------|------------------|--------|-------------------------------------------------------------------------------------------------------------------------------|----------|
| uri         | URI              | 1      | GET against this URI refreshes the client representation of the scalability group definition.                                 | 12.1.0.2 |
| name        | String           | 1      | Name of the scalability group as given by the user. It is a UNICODE string to support different languages.                    | 12.1.0.2 |
| description | String           | 0..1   | Human readable description of the scalability group given by the user. It is a UNICODE string to support different languages. | 12.1.0.2 |
| nodes       | Collection< URI> | 1      | List of URIs that represent the entities making up this scalability group.                                                    | 12.1.0.2 |
| count       | Integer          | 1      | Count of the nodes that are in the scalability group.                                                                         | 12.1.0.2 |

**Table 26–4 (Cont.) ScalabilityGroup Data Model**

| Field          | Type           | Occurs | Description                                                                                                                                                                                                                                                                                                                                                                                       | Since    |
|----------------|----------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| status         | String         | 0..1   | Indicates the status of the scalability group. This field contains the vendor dependent semantics that the service provider implements. For example, a service provider may implement an ONLINE status to indicate that all the entities, recursively, are in an ONLINE status. Or a service provider may implement an ONLINE status to indicate that at least one entity is in an ONLINE status. | 12.1.0.2 |
| resource_state | ResourceState  | 1      | Validity of the other scalability group fields is guaranteed only when the resource state is READY.<br><br>Otherwise, the client shall not assume the validity of the fields.                                                                                                                                                                                                                     | 12.1.0.2 |
| created        | Timestamp      | 1      | Date and time, in ISO 8601 format, when the scalability group is created.                                                                                                                                                                                                                                                                                                                         | 12.1.0.2 |
| max            | Signed Integer | 0..1   | Maximum number of nodes this scalability group can hold. If not provided, the client should assume it is unlimited, which is the specified with as the value “-1”.                                                                                                                                                                                                                                | 12.1.0.2 |
| min            | Integer        | 0..1   | Minimal number of nodes this scalability group should hold to be considered a functional scalability group. If not specified, the client should assume it is 1.                                                                                                                                                                                                                                   | 12.1.0.2 |
| homogenous     | Boolean        | 1      | TRUE when the scalability group contains homogenous entities and FALSE otherwise.                                                                                                                                                                                                                                                                                                                 | 12.1.0.2 |

### 26.1.5 VM [application/oracle.com.cloud.common.VM+json]

Extends: application/oracle.com.cloud.common.ServiceInstance

A VM is a computing container providing a complete system platform that supports the execution of a complete OS stack. The following table describes the VM Data Model.

**Table 26–5 VM Data Model**

| Field       | Type   | Occurs | Description                                                                                                                                     | Since    |
|-------------|--------|--------|-------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| uri         | URI    | 1      | GET against this URI refreshes the client representation of the VM definition.                                                                  | 12.1.0.1 |
| name        | String | 1      | Name of the VM as given by the user or generated by the platform.                                                                               | 12.1.0.1 |
| description | String | 0..1   | Human readable description of the Server given by the user or generated by the platform. It is a UNICODE string to support different languages. | 12.1.0.1 |
| status      | String | 1      | Current running status of this Server. The service provider can overwrite the valid values for this field and may implement status operations.  | 12.1.0.1 |

**Table 26–5 (Cont.) VM Data Model**

| Field           | Type                         | Occurs | Description                                                                                                                                                         | Since    |
|-----------------|------------------------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| based_on        | URI                          | 0..1   | URI of the service template on which this Server is based.                                                                                                          | 12.1.0.1 |
| hostname        | String                       | 0..1   | Qualified host name of this Server when available.                                                                                                                  | 12.1.0.2 |
| cpu             | [Number,Number]              | 1      | Numeric sizing of the CPU where the first number indicates the count of the CPU cores and the second number indicates the CPU speed in MHz per core.                | 12.1.0.1 |
| memory          | Number                       | 1      | Numeric sizing of the RAM in MByte.                                                                                                                                 | 12.1.0.1 |
| disks           | {String,Number}[]            | 0..1   | Name and size in GB of local disks.                                                                                                                                 | 12.1.0.2 |
| interfaces      | Collection<NetworkInterface> | 1      | Network interfaces associated with this Server.                                                                                                                     | 12.1.0.2 |
| params          | { }                          | 0..1   | Vendor specific configuration parameters for this Server.                                                                                                           | 12.1.0.1 |
| resource_state  | ResourceState                | 1      | Validity of the other Server fields is guaranteed only when the resource state is READY.<br><br>Otherwise, the client should not assume the validity of the fields. | 12.1.0.1 |
| created         | Timestamp                    | 1      | Date and time, in ISO 8601 format, when the Server was created.                                                                                                     | 12.1.0.1 |
| type            | String                       | 1      | Type of server that represents this computational container (PHYSICAL, VIRTUAL).                                                                                    | 12.1.0.1 |
| os              | String                       | 0..1   | OS information of the Server where appropriate.                                                                                                                     | 12.1.0.1 |
| media_type      | String                       | 1      | The media type of the resource.                                                                                                                                     | 12.1.0.4 |
| saved_templates | Collection<VMTemplate>       | 0..1   | The collection of VMTemplate that were created based on this VM resource.                                                                                           | 12.1.0.5 |

The status field of the VM Data Model should contain the running status of the VM. It is expected that the service provider implements at least the following valid values:

- STOPPED
- STOPPING
- STARTING
- STARTED
- SUSPENDED
- SUSPENDING
- RESUMING
- RESTARTING
- DELETED



### 26.1.6 VNet [application/oracle.com.cloud.common.VNet+json]

A VNet is a service that is capable of providing network addresses, routing rules, security constraints, and access limits.

The following table describes the VNet Data Model.

**Table 26–6 VNet Data Model**

| Field       | Type                         | Occurs | Description                                                                                                                                   | Since    |
|-------------|------------------------------|--------|-----------------------------------------------------------------------------------------------------------------------------------------------|----------|
| uri         | URI                          | 1      | GET against this URI refreshes the client representation of the VNet definition.                                                              | 12.1.0.2 |
| name        | String                       | 1      | Name of the VNet as given by the user or generated by the platform. It is a UNICODE string to support different languages.                    | 12.1.0.2 |
| description | String                       | 0..1   | Human readable description of the VNet given by the user or generated by the platform. It is a UNICODE string to support different languages. | 12.1.0.2 |
| created     | Timestamp                    | 1      | Date and time, in ISO 8601 format, when this VNet was created.                                                                                | 12.1.0.2 |
| interfaces  | Collection<NetworkInterface> | 0..1   | List of NetworkInterface resources that are part of the VNet.                                                                                 | 12.1.0.2 |

### 26.1.7 NetworkInterface [application/oracle.com.cloud.common.NetworkInterface+json]

An instance of the network interface is identified by a network end point and consists of a complete address that can be interpreted by the underlying network infrastructure.

The following table describes the NetworkInterface Data Model.

**Table 26–7 NetworkInterface Data Model**

| Field          | Type          | Occurs | Description                                                                                                                                                                | Since    |
|----------------|---------------|--------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| uri            | URI           | 1      | GET against this URI refreshes the client representation of the NetworkInterface definition.                                                                               | 12.1.0.2 |
| name           | String        | 1      | Name of the NetworkInterface as given by the user or generated by the platform. It is a UNICODE string to support different languages.                                     | 12.1.0.2 |
| description    | String        | 0..1   | Human readable description of the NetworkInterface given by the user or generated by the platform. It is a UNICODE string to support different languages.                  | 12.1.0.2 |
| address        | String        | 1      | Address of this interface. If an IP based protocol, specified as a dotted notation IPv4 or IPv6 address.                                                                   | 12.1.0.2 |
| resource_state | ResourceState | 0..1   | The validity of the other NetworkInterface fields is guaranteed only when the resource state is READY. Otherwise, the client should not assume the validity of the fields. | 12.1.0.2 |

**Table 26–7 (Cont.) NetworkInterface Data Model**

| Field | Type   | Occurs | Description                                                 | Since    |
|-------|--------|--------|-------------------------------------------------------------|----------|
| mac   | String | 0..1   | The MAC address of the network interface where appropriate. | 12.1.0.2 |

## 26.1.8 VMTemplate [application/oracle.com.cloud.common.VMTemplate+json]

Extends: application/oracle.com.cloud.common.ServiceTemplate

VMTemplate is a class of ServiceTemplate. It is permissible to accept the ServiceTemplate media type on an instance of VMTemplate. VMTemplate is a preconfigured deployable entity that realizes a VM resource. The following table describes the VMTemplate Data Model.

**Table 26–8 VMTemplate Data Model**

| Field                       | Type                          | Occurs | Description                                                                                                                                                               | Since    |
|-----------------------------|-------------------------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| uri                         | URI                           | 1      | A GET against this URI refreshes the client representation of the VMTemplate definition to this user.                                                                     | 12.1.0.1 |
| name                        | String                        | 1      | A human readable name, given to the VMTemplate.                                                                                                                           | 12.1.0.1 |
| description                 | String                        | 0..1   | A brief description, given to the ServiceTemplate. SHALL be a UNICODE string to support different languages.                                                              | 12.1.0.1 |
| type                        | String                        | 1      | The String that describes the type of the VM Template.<br>This has value "Template".                                                                                      | 12.1.0.1 |
| created                     | Timestamp                     | 1      | Date and time, in ISO 8601 format, when the VMTemplate was created.                                                                                                       | 12.1.0.1 |
| service_family_type         | String                        | 0..1   | The name of the service family type that this service template is categorized under.<br>This has value "iaas".                                                            | 12.1.0.4 |
| default_instance_media_type | String                        | 0..1   | The default media type of the service instance that can be created using this template.<br>This has value "application/oracle.com.cloud.common.VM".                       | 12.1.0.4 |
| resource_state              | ResourceState                 | 0..1   | Only a service template with READY state can be deployed.                                                                                                                 | 12.1.0.1 |
| deployment_params           | List of Deployment Parameters | 0..1   | Contains the list of data structure of Deployment Parameters to indicate the parameters that may be specified during the service instance deployment using this template. | 12.1.0.4 |
| zones                       | List of Zones                 | 0..1   | Contains the list of Zone resources that this service template can be used to create service instances with.                                                              | 12.1.0.2 |
| service_instances           | Collection<?>                 | 0..1   | Contains the list of service instances that are created with this template.                                                                                               | 12.1.0.4 |

## 26.1.9 AssemblyTemplate

### [application/oracle.com.cloud.common.AssemblyTemplate+json]

Extends: application/oracle.com.cloud.common.ServiceTemplate+json

AssemblyTemplate is a class of ServiceTemplate. It is permissible to accept the ServiceTemplate media type on an instance of AssemblyTemplate.

AssemblyTemplate is a deployable entity that realizes an AssemblyInstance resource that may contain multiple resources that are interconnected. The following table describes the AssemblyTemplate Data Model.

**Table 26–9 AssemblyTemplate Data Model**

| Field                       | Type                          | Occurs | Description                                                                                                                                                               | Since    |
|-----------------------------|-------------------------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| uri                         | URI                           | 1      | A GET against this URI refreshes the client representation of the AssemblyTemplate definition to this user.                                                               | 12.1.0.2 |
| name                        | String                        | 1      | A human readable name, given to the AssemblyTemplate.                                                                                                                     | 12.1.0.2 |
| description                 | String                        | 0..1   | A brief description, given to the AssemblyTemplate. SHALL be a UNICODE string to support different languages.                                                             | 12.1.0.2 |
| type                        | String                        | 1      | The String that describes the type of the Assembly Template.<br>This has value "Assembly".                                                                                | 12.1.0.2 |
| created                     | Timestamp                     | 1      | Date and time, in ISO 8601 format, when the AssemblyTemplate was created.                                                                                                 | 12.1.0.2 |
| service_family_type         | String                        | 0..1   | The name of the service family type that this service template is categorized under.<br>This has value "iaas".                                                            | 12.1.0.4 |
| default_instance_media_type | String                        | 0..1   | The default media type of the service instance that can be created using this template.<br>This has value "application/oracle.com.cloud.common.AssemblyInstance".         | 12.1.0.4 |
| resource_state              | ResourceState                 | 0..1   | Only a service template with READY state can be deployed.                                                                                                                 | 12.1.0.1 |
| deployment_params           | List of Deployment Parameters | 0..1   | Contains the list of data structure of Deployment Parameters to indicate the parameters that may be specified during the service instance deployment using this template. | 12.1.0.4 |
| zones                       | List of Zones                 | 0..1   | Contains the list of Zone resources that this service template can be used to create service instances with.                                                              | 12.1.0.2 |
| service_instances           | Collection<?>                 | 0..1   | Contains the list of service instances that are created with this template.                                                                                               | 12.1.0.4 |

**Table 26–9 (Cont.) AssemblyTemplate Data Model**

| Field         | Type                     | Occurs | Description                                                                                                                                       | Since    |
|---------------|--------------------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| definition    | String                   | 0..1   | Contains the definition of the assembly template. This would be the OVF xml metadata representing the assembly template.                          | 12.1.0.1 |
| ovfdescriptor | String                   | 0..1   | Same as definition.                                                                                                                               | 12.1.0.1 |
| deploy_plans  | List of Deployment Plans | 0..1   | Contains the list of Deployment Plans that are associated with the assembly instance. The deployment plans are pre-filled configuration metadata. | 12.1.0.2 |

The following table describes the structure of the Deployment Plan.

**Table 26–10 Deployment Plan Structure**

| Field       | Type   | Occurs | Description                             |
|-------------|--------|--------|-----------------------------------------|
| name        | String | 1      | The name of the deployment plan.        |
| description | String | 0..1   | The description of the deployment plan. |
| xml         | String | 1      | The deployment plan XML.                |

### 26.1.10 IaasServiceFamilyType

[application/oracle.com.cloud.iaas.IaasServiceFamilyType+json]

Extends: application/oracle.com.cloud.common.ServiceFamilyType+json

The Infrastructure service family type describes the common service family type attributes and also the infrastructure specific attributes. The following table describes the IaaS Service Family Type Data Model.

**Table 26–11 AssemblyTemplate Data Model**

| Field             | Type                         | Occurs | Description                                                                                                        | Since    |
|-------------------|------------------------------|--------|--------------------------------------------------------------------------------------------------------------------|----------|
| uri               | URI                          | 1      | A GET against this URI refreshes the client representation of the ServiceFamilyType definition to this user.       | 12.1.0.4 |
| name              | String                       | 1      | A human readable name, given to the ServiceFamilyType.                                                             | 12.1.0.4 |
| description       | String                       | 0..1   | A human readable description of the Service Family Type. SHALL be a UNICODE string to support different languages. | 12.1.0.4 |
| resource_state    | ResourceState                | 0..1   | The resource state of the resource.                                                                                | 12.1.0.4 |
| media_type        | String                       | 1      | The media type of the resource.                                                                                    | 12.1.0.4 |
| service_templates | Collection <ServiceTemplate> | 0..1   | The collection of service templates that are of this service family type.                                          | 12.1.0.4 |
| zones             | Collection <Zone>            | 0..1   | The collection of zones that support service instances of this service family type.                                | 12.1.0.4 |

**Table 26–11 (Cont.) AssemblyTemplate Data Model**

| Field            | Type                         | Occurs | Description                                                                                                                                                                                               | Since    |
|------------------|------------------------------|--------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| instance_types   | Collection<br><InstanceType> | 0..1   | The collection of instance types that are supported by this service family type.                                                                                                                          | 12.1.0.4 |
| quota_usages     | Collection<br><QuotaUsage>   | 0..1   | The collection of quota usages of the perspective family type of the authenticated user.                                                                                                                  | 12.1.0.5 |
| quota_def        | Json Object                  | 0..1   | The family type specific quota metadata information, this may include some permission information.<br><br>The structure of this object would be documented in the perspective family type resource model. | 12.1.0.5 |
| instance_options | List of VM Size Options      | 0..1   | The list of VM size options that may be available to the authenticated user.                                                                                                                              | 12.1.0.4 |

The following table describes the structure of the VM Size option:

**Table 26–12 VM Size Option Structure**

| Field         | Type   | Occurs | Description                                                                           | Since    |
|---------------|--------|--------|---------------------------------------------------------------------------------------|----------|
| name          | String | 1      | The name of the VM Size. This may be used as the parameter for VM resource creations. | 12.1.0.4 |
| id            | String | 1      | The ID of the VM Size.                                                                | 12.1.0.4 |
| cpu           | Number | 1      | The amount of CPU.                                                                    | 12.1.0.4 |
| memory        | Number | 1      | The amount of Memory in MB.                                                           | 12.1.0.4 |
| local_storage | Number | 1      | The amount of local storage in GB.                                                    | 12.1.0.4 |

The IaaS supports the following Quota Usage.

**Table 26–13 Quota Usage for IaaS**

| Quota         | Description                                                       | Since    |
|---------------|-------------------------------------------------------------------|----------|
| CPU           | The number of CPUs in discrete units.                             | 12.1.0.5 |
| RAM           | Memory size in MB.                                                | 12.1.0.5 |
| Local Storage | The amount of storage local to the server in GB.                  | 12.1.0.5 |
| Extra Storage | The amount of storage external to the server in GB.               | 12.1.0.5 |
| SWLIB Storage | The amount of storage in the Enterprise Manager Software Library. | 12.1.0.5 |
| Server        | The number of servers in discrete units.                          | 12.1.0.5 |

The IaaS Quota definition is described in the media type `application/oracle.com.cloud.iaas.Quota`.

### 26.1.11 IaaS Quota Definition [application/oracle.com.cloud.iaas.Quota+json]

This resource describes the definition for the IaaS quota. The following table describes the IaaS Quota Definition Data Model

**Table 26–14 IaaS Quota Definition Data Model**

| Field                      | Type                        | Occurs | Description                                                                                                                                        | Since    |
|----------------------------|-----------------------------|--------|----------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| uri                        | URI                         | 1      | A GET against this URI refreshes the client representation of the IaaS Quota definition to this user.                                              | 12.1.0.5 |
| allowedSaveAsTemplate      | Boolean                     | 1      | Indicates whether the authenticated user has permission to save a VM as a template.                                                                | 12.1.0.5 |
| allowedSaveToSwlib         | Boolean                     | 1      | Indicates whether the authenticated user has permission to save entities to the Enterprise Manager software library.                               | 12.1.0.5 |
| allowedSaveToSwlibOnExpiry | Boolean                     | 1      | Indicates whether the authenticated user has permission to save VM as a template into the Enterprise Manager software library when the VM expires. | 12.1.0.5 |
| resourceQuotas             | Collection <Resource Quota> | 1      | The quotas of resources defined for the authenticated user.                                                                                        | 12.1.0.5 |

The following describes the structure of the Resource Quota

**Table 26–15 Structure of the Resource Quota**

| Field       | Type   | Occurs | Description                                                                                                                                                                          | Since    |
|-------------|--------|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------|
| name        | String | 1      | The name of the resource.                                                                                                                                                            | 12.1.0.5 |
| description | String | 0..1   | The description of the resource.                                                                                                                                                     | 12.1.0.5 |
| maxAllowed  | Number | 0..1   | The maximum number of units that the authenticated resource is allocated. If this value is negative or not defined, the user does not have any limit on the amount that can be used. | 12.1.0.5 |
| unit        | String | 1      | The measurement unit of the resource.                                                                                                                                                | 12.1.0.5 |

## 26.2 Supported Operations Examples

The following sections provide examples of the different operations supported by the IaaS APIs:

- [Creating a Service Instance](#)
- [Updating a Service Instance](#)
- [Adding a VM Disk](#)
- [Deleting a Service Instance](#)
- [Listing Service Instances](#)
- [Searching Service Templates](#)

## 26.2.1 Creating a Service Instance

This section describes how to create a service instance using the Cloud resources.

### 26.2.1.1 VM Creation

As VM is a service instance in the service family type "iaas", you need to identify the zone of the same family type in which to create the VM. There are two approaches to creating a VM:

1. POST to the "iaas" Zone of which VM is to reside
2. POST to the VMTemplate of which the VM is to be based on

Note the attributes in the submitted body for the POST. The pattern is the same for other service instance creations. Because the intention is clear, the "zone" attribute is not required while POSTing to the "iaas" Zone. Similarly, the "based\_on" attribute is not needed when POSTing to the VMTemplate.

#### Post to the "iaas" Zone

The following shows the configuration for POSTing to "iaas" Zone:

**Table 26–16 POST Request**

| Feature | Description                                                                                                                                                                                                                                                                                                                                      |
|---------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| URL     | <code>https://example.oracle.com/em/cloud/iaas/servicetemplate/vm/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp%3AVirtualization%3ATemplate%3AC76CEB5563EA5E13E040578CDA817FAF%3A0.1</code>                                                                                                                                              |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=<br>Content-Type: application/oracle.com.cloud.common.VM+json                                                                                                                                                                                                                                       |
| Body    | <pre>{   "based_on": "/em/cloud/iaas/servicetemplate/vm/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp%3AVirtualization%3ATemplate%3AC76CEB5563EA5E13E040578CDA817FAF%3A0.1",   "cpu" : [2,0],   "memory" : "512",   "params":{     "server_prefix":"ZONEPOST",     "vnc_password":"welcome1",     "root_password":"welcome1"   } }</pre> |
| Method  | POST                                                                                                                                                                                                                                                                                                                                             |
| Since   | 12.1.0.1 Cloud Service Portal plug-in                                                                                                                                                                                                                                                                                                            |

The following is the response from the above interaction:

```
{
 "uri" : "/em/cloud/iaas/server/byrequest/102" ,
 "name" : "VDOSI VM Creation 1345392541850" ,
 "resource_state" : {
 "state" : "INITIATED" ,
 "messages" :
 [
 {
 "text" : "The Request with ID '102' is scheduled with Job Id 'C79997609390CB9BE040578CDA817D96'" ,

```

```

 "date" : "2012-08-19T16:09:02+0000"
 }
]
 } ,
 "context_id" : "102" ,
 "media_type" : "application/oracle.com.cloud.common.VM+json" ,
 "service_family_type" : "iaas" ,
 "created" : "2012-08-19T16:09:02+0000"
}

```

### Post to the VMTemplate

The following describes the configurations for POSTing to the VMTemplate:

**Table 26–17 POST to VMTemplate**

| Feature | Description                                                                                                                                                                                                                                                 |
|---------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| URL     | https://example.oracle.com/em/cloud/iaas/servicetemplate/vm/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp%3AVirtualization%3ATemplate%3AC76CEB5563EA5E13E040578CDA817FAF%3A0.1                                                                      |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=<br>Content-Type: application/oracle.com.cloud.common.VM+json                                                                                                                                                  |
| Body    | {       "zone": "/em/cloud/iaas/zone/D73AF0D42C8459E11419862797D1F37D",       "cpu" : [2,0],       "memory" : "512",       "params":{         "server_prefix":"STPOST",         "vnc_password":"welcome1",         "root_password":"welcome1"       }     } |
| Method  | POST                                                                                                                                                                                                                                                        |
| Since   | 12.1.0.4 Cloud Service Portal plug-in                                                                                                                                                                                                                       |

The following is the response from the above interaction:

```

{
 "uri" : "/em/cloud/iaas/server/byrequest/101" ,
 "name" : "VDOSI VM Creation 1345391921407" ,
 "resource_state" : {
 "state" : "INITIATED" ,
 "messages" :
 [
 {
 "text" : "The Request with ID '101' is scheduled with Job Id 'C7999760937CCB9BE040578CDA817D96'" ,
 "date" : "2012-08-19T15:58:47+0000"
 }
]
 } ,
 "context_id" : "101" ,
 "media_type" : "application/oracle.com.cloud.common.VM+json" ,
 "service_family_type" : "iaas" ,
 "created" : "2012-08-19T15:58:42+0000"
}

```



This response is the same as the previous version. If the zone is not specified, then it is not possible for the Web service to determine where the VM should be deployed. The following interaction illustrates this:

**Table 26–18 POST Request**

| Feature | Description                                                                                                                                                                                         |
|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| URL     | <code>https://example.oracle.com/em/cloud/iaas/servicetemplate/vm/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp%3AVirtualization%3ATemplate%3AC76CEB5563EA5E13E040578CDA817FAF%3A0.1</code> |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=<br>Content-Type: application/oracle.com.cloud.common.VM+json                                                                                          |
| Body    | <pre>{   "cpu" : [2,0],   "memory" : "512",   "params":{     "server_prefix":"STPOST",     "vnc_password":"welcome1",     "root_password":"welcome1"   } }</pre>                                    |
| Method  | POST                                                                                                                                                                                                |

The following is the response from the Web service:

```
{
 "messages" :
 [
 {
 "hint" : "'zone' attribute needs to be specified to complete 'POST'
request for 'application/oracle.com.cloud.common.VM' resource" ,
 "stack_trace" :
"oracle.sysman.emInternalSDK.ssa.cloudapi.rest.RestServletException: 'zone'
attribute needs to be specified to complete 'POST' request for
'application/oracle.com.cloud.common.VM' resource\n\tat
oracle.sysman.ssa.cloudapi.iaas.VMRequest.GenerateVMFromDeployment (VMRequest.java:
149)\n\tat
oracle.sysman.ssa.cloudapi.iaas.VMTemplate.processRequest (VMTemplate.java:93)\n\tat
oracle.sysman.ssa.cloudapi.iaas.IaaSServiceProvider.processRequest (IaaSServiceProv
ider.java:582)\n\tat
oracle.sysman.emInternalSDK.ssa.cloudapi.EMCloudServlet.perform (EMCloudServlet.jav
a:236)\n\tat
oracle.sysman.emInternalSDK.ssa.cloudapi.EMCloudServlet.performPost (EMCloudServlet
.java:385)\n\tat
oracle.sysman.emInternalSDK.ssa.cloudapi.rest.AbstractRestServlet.doPost (AbstractR
estServlet.java:137)\n\tat
javax.servlet.http.HttpServlet.service (HttpServlet.java:727)\n\tat
javax.servlet.http.HttpServlet.service (HttpServlet.java:820)\n\tat
weblogic.servlet.internal.StubSecurityHelper$ServletServiceAction.run (StubSecurity
Helper.java:227)\n\tat web1" ,
 "text" : "Did not successfully execute 'POST' resource operation on
'application/oracle.com.cloud.common.Cloud' identified by
'/em/cloud/iaas/servicetemplate/vm/oracle:defaultService:em:provisioning:1:cmp:Vir
tualization:Template:C76CEB5563EA5E13E040578CDA817FAF:0.1' " ,
 "date" : "2012-08-19T16:02:51+0000"
 }
]
}
```

```
]
 }
```

### Polling the VM Creation

Once the VM creation has been POSTed, you can GET the resource identified by the return URI to keep track of the status.

**Table 26–19 Polling the VM Creation**

| Feature | Description                                                   |
|---------|---------------------------------------------------------------|
| URL     | https://example.oracle.com/em/cloud/iaas/server/byrequest/101 |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=                 |
| Method  | GET                                                           |

The following is the response:

```
{
 "uri" : "/em/cloud/iaas/server/byrequest/101" ,
 "name" : "VDOSI VM Creation 1345391921407" ,
 "resource_state" : {
 "state" : "CREATING" ,
 "messages" :
 [
 {
 "text" : "The Request with ID '101' is being processed, and has Job Id
'C7999760937CCB9BE040578CDA817D96'" ,
 "date" : "2012-08-19T16:14:11+0000"
 }
]
 } ,
 "context_id" : "101" ,
 "media_type" : "application/oracle.com.cloud.common.VM+json" ,
 "service_family_type" : "iaas" ,
 "created" : "2012-08-19T15:58:42+0000"
}
```

Note that the "resource\_state" attribute is showing "CREATING" state to indicate the VM is still being created. Once the VM is created, the same GET will return the VM resource with "READY" state, as shown below:

```
{
 "uri" : "/em/cloud/iaas/server/vm/C5CA59AF9386975E8AEA45B0F040F095" ,
 "name" : "STPOST" ,
 "resource_state" : {
 "state" : "READY"
 } ,
 "context_id" : "C5CA59AF9386975E8AEA45B0F040F095" ,
 "media_type" : "application/oracle.com.cloud.common.VM+json" ,
 "metrics" : {
 "media_type" : "application/oracle.com.cloud.common.Metric+json" ,
 "total" : "4" ,
 "elements" :
 [
 {
 "uri" : "/em/cloud/metric/iaas%3Avm%3ACPU_
Utilization%3AC5CA59AF9386975E8AEA45B0F040F095" ,
 "name" : "CPU_Utilization" ,
 "media_type" : "application/oracle.com.cloud.common.Metric+json" ,
```

```

 "type" : "NUMBER" ,
 "current_value" : "" ,
 "current_time" : "2012-08-19T16:17:02+0000"
 } ,
 {
 "uri" : "/em/cloud/metric/iaas%3Avm%3ATotal_Network_Throughput%3AC5CA59AF9386975E8AEA45B0F040F095" ,
 "name" : "Total_Network_Throughput" ,
 "media_type" : "application/oracle.com.cloud.common.Metric+json" ,
 "type" : "NUMBER" ,
 "current_value" : "" ,
 "current_time" : "2012-08-19T16:17:02+0000"
 } ,
 {
 "uri" : "/em/cloud/metric/iaas%3Avm%3ATotal_Disk_Throughput%3AC5CA59AF9386975E8AEA45B0F040F095" ,
 "name" : "Total_Disk_Throughput" ,
 "media_type" : "application/oracle.com.cloud.common.Metric+json" ,
 "type" : "NUMBER" ,
 "current_value" : "" ,
 "current_time" : "2012-08-19T16:17:02+0000"
 } ,
 {
 "uri" : "/em/cloud/metric/iaas%3Avm%3AFilesystem_Total_Used%3AC5CA59AF9386975E8AEA45B0F040F095" ,
 "name" : "Filesystem_Total_Used" ,
 "media_type" : "application/oracle.com.cloud.common.Metric+json" ,
 "type" : "NUMBER" ,
 "current_value" : "" ,
 "current_time" : "2012-08-19T16:17:02+0000"
 }
]
} ,
"service_family_type" : "iaas" ,
"status" : "STARTED" ,
"created" : "2012-08-19T16:15:21+0000" ,
"hostname" : "dadvfm0052.us.oracle.com" ,
"cpu" :
[
 "2" ,
 "0"
] ,
"memory" : "512" ,
"disks" :
[
 [
 "default_disk" ,
 "10240"
] ,
 [
 "7e8a5d25-7995-40c9-9b16-9274cc2d64f3.img (2)" ,
 "6918"
] ,
 [
 "fc4f377c-ed3e-4823-a9b3-6b4d2289e215.img (2)" ,
 "192"
] ,
 [
 "f5dba62c-eb49-40c6-af0b-06187b2b9856.img (2)" ,
 "5120"
]
]

```

```

]
],
 "interfaces" : {
 "media_type" : "application/oracle.com.cloud.common.NetworkInterface+json" ,
 "total" : "1" ,
 "elements" :
 [
 {
 "name" : "cloud_network" ,
 "address" : "10.229.136.185" ,
 "mac" : "00:21:f6:5f:73:33"
 }
]
 },
 "os" : "None" ,
 "type" : "VIRTUAL"
}

```

### 26.2.1.2 Assembly Instance Creation

Similar to VM creation, assembly instance creation can also be POSTed to both a Zone and an AssemblyTemplate. The following shows the service template details:

**Table 26–20 Assembly Instance Creation**

| Feature | Description                                                                                                                                                                                               |
|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| URL     | <code>https://example.oracle.com/em/cloud/iaas/servicetemplate/assembly/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp%3AVirtualization%3AAssembly%3AC75E88B04D7FEDEDE040578CDA810E49%3A0.1</code> |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=                                                                                                                                                             |
| Method  | GET                                                                                                                                                                                                       |
| Since   | 12.1.0.4 Cloud Service Portal plug-in                                                                                                                                                                     |

The following is the response from the Web service:

```

{
 "uri" :
"/em/cloud/iaas/servicetemplate/assembly/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp%3AVirtualization%3AAssembly%3AC75E88B04D7FEDEDE040578CDA810E49%3A0.1"
,
 "name" : "sidb_assembly" ,
 "resource_state" : {
 "state" : "READY"
 } ,
 "context_id" :
"oracle:defaultService:em:provisioning:1:cmp:Virtualization:Assembly:C75E88B04D7FEDEDE040578CDA810E49:0.1" ,
 "media_type" : "application/oracle.com.cloud.common.AssemblyTemplate+json" ,
 "type" : "Assembly" ,
 "created" : "2012-08-16T11:02:38+0000" ,
 "default_instance_media_type" :
"application/oracle.com.cloud.common.AssemblyInstance+json" ,
 "zones" : {
 "media_type" : "application/oracle.com.cloud.iaas.Zone+json" ,
 "total" : "1" ,
 "elements" :
 [
 {

```

```

 "uri" : "/em/cloud/iaas/zone/D73AF0D42C8459E11419862797D1F37D" ,
 "name" : "cloud_zone" ,
 "media_type" : "application/oracle.com.cloud.iaas.Zone+json" ,
 "service_family_type" : "iaas" ,
 "type" : "VM_ZONE"
 }
]
} ,
"service_instances" : {
 "media_type" : "application/oracle.com.cloud.common.ServiceInstance+json" ,
 "total" : "0" ,
 "elements" :
 [
]
} ,
"deployment_params" :
[
 {
 "name" : "zone" ,
 "description" : "'zone' attribute is the URI of the zone that the
'application/oracle.com.cloud.common.AssemblyInstance' resource is to be created
in" ,
 "type" : "STRING" ,
 "require" : "false" ,
 "sensitive" : "false"
 } ,
 {
 "name" : "based_on" ,
 "description" : "'based_on' attribute is the URI of the service template
or format String of the originated source which the deployed
'application/oracle.com.cloud.common.AssemblyInstance' resource is to be followed"
,
 "type" : "STRING" ,
 "require" : "false" ,
 "sensitive" : "false"
 } ,
 {
 "name" : "deployment_plan" ,
 "description" : "'deployment_plan' attribute describes the configuration
of which the deployed 'application/oracle.com.cloud.common.AssemblyInstance'
resource is to be used in the XML deployment plan format" ,
 "type" : "STRING" ,
 "require" : "false" ,
 "sensitive" : "false"
 }
] ,
"directory" : "COMP_Components" ,
"updated" : "2012-08-16T11:02:38+0000" ,
"creator" : "DMBHAT1" ,
"last_modified_by" : "DMBHAT1" ,
"componenttype" : "Virtualization" ,
"oracle_owned" : "false" ,
"deploy_plans" :
[
 {
 "name" : "cloud_zone" ,
 "description" : "Default Assembly Deployment Plan for Service Template
'sidb_assembly' on Zone 'cloud_zone'" ,
 "context_id" : "D73AF0D42C8459E11419862797D1F37D" ,
 "xml" : "<?xml version='1.0' encoding='UTF-8'
standalone='yes'?>\n<ns2:ConfigurationData

```

```

xmlns:ns2="http://www.oracle.com/sysman/vt/RequestConfigData">\n
<AssemblyDeployment assemblyInstanceName="DEFAULT_PLEASE_CHANGE">\n
<SourceBinary type="TYPE_SWLIB"
name="oracle:defaultService:em:provisioning:1:cmp:Virtualization:Assembly:C75E88B
04D7FEDEDE040578CDA810E49:0.1"/>\n
 <DeploymentTarget type="oracle_vm_
zone" name="D73AF0D42C8459E11419862797D1F37D"/>\n
<AcceptedAllEULAs>>false</AcceptedAllEULAs>\n
<InstallEMAgent>>false</InstallEMAgent>\n
<DisableCleanup>>false</DisableCleanup>\n
 <EMAgentConfig
installUserPassword=" " installUserName="oracle"/>\n
<AssemblyNetworkConfig>\n
 <AssemblyNetwork name="network-1"/>\n
</AssemblyNetworkConfig>\n
 <VirtualSystemCollectionConfig id="sidbasmA_
linuxx64_11203psu1">\n
 <ProductConfiguration>\n
<Product>\n
 <Properties/>\n
<EMAgent>>false</EMAgent>\n
 </Product>\n
</ProductConfiguration>\n
 <VirtualSystemConfig
agentPushInstallationEnabled="true" agentInstallationType="Push Agent
Installation" vmInstance="true" targetName="sidbasmA_linuxx64_
11203psu1/sidbasm:%assembly_instance_name%" id="sidbasmA_linuxx64_
11203psu1/sidbasm">\n
 <ProductConfiguration>\n
<Product>\n
 <Properties>\n
<Property
id="ocm.runConfiguration">\n<Name>ocm.runConfiguration</Name>\n<Required>>false</
Required>\n<Value>>false</Value>\n<Secret>>false</Secret>\n
</Property>\n
 <Property
id="ocm.metalinkCsiRegistration.CSI">\n<Name>ocm.metalinkCsiRegistration.CSI</Na
me>\n<Required>>false</Required>\n<Value></Value>\n<Secret>>false</Secret>\n
</Property>\n
 <Property
id="ocm.metalinkCsiRegistration.metalinkId">\n<Name>ocm.metalinkCsiRegistration.
metalinkId</Name>\n<Required>>false</Required>\n<Value></Value>\n<Secret>>false</Sec
ret>\n
 </Property>\n
 <Property
id="ocm.proxyPassword">\n<Name>ocm.proxyPassword</Name>\n<Required>>false</Requir
ed>\n<Value></Value>\n<Secret>>true</Secret>\n
 </Property>\n
 <Property
id="ocm.metalinkCsiRegistration.countryCode">\n<Name>ocm.metalinkCsiRegistration
.countryCode</Name>\n<Required>>false</Required>\n<Value></Value>\n<Secret>>false</S
ecret>\n
 </Property>\n
 <Property id="system-files|db_asm_disk_
7|size-units">\n<Name>system-files|db_asm_disk_
7|size-units</Name>\n<Required>>false</Required>\n<Value>MB</Value>\n<Secret>>false<
/Secret>\n
 </Property>\n
 <Property id="system-files|db_asm_disk_
1|size-units">\n<Name>system-files|db_asm_disk_
1|size-units</Name>\n<Required>>false</Required>\n<Value>MB</Value>\n<Secret>>false<
/Secret>\n
 </Property>\n
 <Property id="system-files|db_asm_disk_
3|size-units">\n<Name>system-files|db_asm_disk_
3|size-units</Name>\n<Required>>false</Required>\n<Value>MB</Value>\n<Secret>>false<
/Secret>\n
 </Property>\n
 <Property id="system-files|db_asm_disk_0|size">\n<Name>system-files|db_asm_
disk_
0|size</Name>\n<Required>>false</Required>\n<Value>2836</Value>\n<Secret>>false</Sec
ret>\n
 </Property>\n
 <Property id="system-files|db_asm_disk_1|size">\n<Name>system-files|db_asm_
disk_
1|size</Name>\n<Required>>false</Required>\n<Value>2836</Value>\n<Secret>>false</Sec
ret>\n
 </Property>\n
 <Property
id="asm-password">\n<Name>asm-password</Name>\n<Required>>true</Required>\n<Value

```

```

> </Value>\n<Secret>true</Secret>\n </Property>\n
<Property id=\"system-fileset|db_asm_disk_7|size\">\n<Name>system-fileset|db_asm_
disk_
7|size</Name>\n<Required>>false</Required>\n<Value>2836</Value>\n<Secret>>false</Sec
ret>\n </Property>\n
<Property
id=\"ocm.metalinkEmailRegistration.metalinkEmailId\">\n<Name>ocm.metalinkEmailRegi
stration.metalinkEmailId</Name>\n<Required>>false</Required>\n<Value></Value>\n<Sec
ret>>false</Secret>\n </Property>\n
<Property
id=\"input|listener-1|global-db-name\">\n<Name>input|listener-1|global-db-name</Na
me>\n<Required>>false</Required>\n<Value>orcl.us.oracle.com</Value>\n<Secret>>false<
/Secret>\n </Property>\n
<Property id=\"system-fileset|db_asm_disk_
4|size-units\">\n<Name>system-fileset|db_asm_disk_
4|size-units</Name>\n<Required>>false</Required>\n<Value>MB</Value>\n<Secret>>false<
/Secret>\n </Property>\n
<Property
id=\"ocm.proxyHost\">\n<Name>ocm.proxyHost</Name>\n<Required>>false</Required>\n<Va
lue></Value>\n<Secret>>false</Secret>\n </Property>\n
<Property
id=\"ocm.repeaterURI\">\n<Name>ocm.repeaterURI</Name>\n<Required>>false</Required>\n
<Value></Value>\n<Secret>>false</Secret>\n </Property>\n
<Property id=\"system-fileset|db_asm_
disk_2|size-units\">\n<Name>system-fileset|db_asm_disk_
2|size-units</Name>\n<Required>>false</Required>\n<Value>MB</Value>\n<Secret>>false<
/Secret>\n </Property>\n
<Property
id=\"input|listener-1|port\">\n<Name>input|listener-1|port</Name>\n<Required>>false
</Required>\n<Value>1521</Value>\n<Secret>>false</Secret>\n </Property>\n
<Property
id=\"ocm.anonymousEmailRegistration.emailId\">\n<Name>ocm.anonymousEmailRegistrati
on.emailId</Name>\n<Required>>false</Required>\n<Value></Value>\n<Secret>>false</Sec
ret>\n </Property>\n
<Property
id=\"db-account-password\">\n<Name>db-account-password</Name>\n<Required>>true</Req
uired>\n<Value> </Value>\n<Secret>true</Secret>\n </Property>\n
<Property id=\"system-fileset|db_asm_
disk_3|size\">\n<Name>system-fileset|db_asm_disk_
3|size</Name>\n<Required>>false</Required>\n<Value>2836</Value>\n<Secret>>false</Sec
ret>\n </Property>\n
<Property id=\"system-fileset|db_asm_disk_2|size\">\n<Name>system-fileset|db_asm_
disk_
2|size</Name>\n<Required>>false</Required>\n<Value>2836</Value>\n<Secret>>false</Sec
ret>\n </Property>\n
<Property
id=\"ocm.proxyPort\">\n<Name>ocm.proxyPort</Name>\n<Required>>false</Required>\n<Va
lue></Value>\n<Secret>>false</Secret>\n </Property>\n
<Property
id=\"ocm.proxyUsername\">\n<Name>ocm.proxyUsername</Name>\n<Required>>false</Requir
ed>\n<Value></Value>\n<Secret>>false</Secret>\n </Property>\n
<Property id=\"system-fileset|db_asm_
disk_6|size\">\n<Name>system-fileset|db_asm_disk_
6|size</Name>\n<Required>>false</Required>\n<Value>2836</Value>\n<Secret>>false</Sec
ret>\n </Property>\n
<Property id=\"system-fileset|db_asm_disk_5|size\">\n<Name>system-fileset|db_asm_
disk_
5|size</Name>\n<Required>>false</Required>\n<Value>2836</Value>\n<Secret>>false</Sec
ret>\n </Property>\n
<Property id=\"system-fileset|db_asm_disk_4|size\">\n<Name>system-fileset|db_asm_

```

```

disk_
4|size</Name>\n<Required>>false</Required>\n<Value>2836</Value>\n<Secret>>false</Sec
ret>\n </Property>\n
<Property
id=\"ocm.metalinkEmailRegistration.metalinkPassword\">\n<Name>ocm.metalinkEmailReg
istration.metalinkPassword</Name>\n<Required>>false</Required>\n<Value>
</Value>\n<Secret>>true</Secret>\n </Property>\n
<Property id=\"system-fileset|db_asm_disk_
6|size-units\">\n<Name>system-fileset|db_asm_disk_
6|size-units</Name>\n<Required>>false</Required>\n<Value>MB</Value>\n<Secret>>false<
/Secret>\n </Property>\n
<Property id=\"system-fileset|db_asm_disk_
5|size-units\">\n<Name>system-fileset|db_asm_disk_
5|size-units</Name>\n<Required>>false</Required>\n<Value>MB</Value>\n<Secret>>false<
/Secret>\n </Property>\n
<Property id=\"system-fileset|db_asm_disk_
0|size-units\">\n<Name>system-fileset|db_asm_disk_
0|size-units</Name>\n<Required>>false</Required>\n<Value>MB</Value>\n<Secret>>false<
/Secret>\n </Property>\n
</Properties>\n <EMAgent>>false</EMAgent>\n
</Product>\n </ProductConfiguration>\n
<HardwareConfiguration>\n <Memory>2048</Memory>\n
<VCPUs>1</VCPUs>\n <HaEnabled>>false</HaEnabled>\n
<StartAfterCreation>true</StartAfterCreation>\n
<CPUSchedulingPriority>50</CPUSchedulingPriority>\n
<CPUCap>100</CPUCap>\n <LocalDisks>\n
<disk fromDefinition=\"true\" name=\"2xZNMmSH5oZYW_System\">\n
<Size>5122</Size>\n <Mode>Read-Write</Mode>\n
<CreatedByVmSize>>false</CreatedByVmSize>\n
 </disk>\n <disk
fromDefinition=\"true\" name=\"AB\">\n <Size>0</Size>\n
<Mode>Read-Write</Mode>\n
<CreatedByVmSize>>false</CreatedByVmSize>\n </disk>\n
<disk fromDefinition=\"true\" name=\"2xZNMmSH5oZYW_sys-asm_base\">\n
<Size>502</Size>\n <Mode>Read-Write</Mode>\n
<CreatedByVmSize>>false</CreatedByVmSize>\n </disk>\n
<disk fromDefinition=\"true\" name=\"2xZNMmSH5oZYW_sys-asm_home\">\n
<Size>10307</Size>\n <Mode>Read-Write</Mode>\n
<CreatedByVmSize>>false</CreatedByVmSize>\n </disk>\n
<disk fromDefinition=\"true\" name=\"2xZNMmSH5oZYW_sys-db_base\">\n
<Size>502</Size>\n <Mode>Read-Write</Mode>\n
<CreatedByVmSize>>false</CreatedByVmSize>\n </disk>\n
<disk fromDefinition=\"true\" name=\"2xZNMmSH5oZYW_sys-db_home\">\n
<Size>9068</Size>\n <Mode>Read-Write</Mode>\n
<CreatedByVmSize>>false</CreatedByVmSize>\n </disk>\n
</LocalDisks>\n <SharedDisks>\n
<SharedDisk fromDefinition=\"true\" name=\"2xZNMmSH5oZYW_db_asm_disk_0\"/>\n
<SharedDisk fromDefinition=\"true\" name=\"2xZNMmSH5oZYW_db_asm_disk_1\"/>\n
<SharedDisk fromDefinition=\"true\" name=\"2xZNMmSH5oZYW_db_asm_disk_2\"/>\n
<SharedDisk fromDefinition=\"true\" name=\"2xZNMmSH5oZYW_db_asm_disk_3\"/>\n
<SharedDisk fromDefinition=\"true\" name=\"2xZNMmSH5oZYW_db_asm_disk_4\"/>\n
<SharedDisk fromDefinition=\"true\" name=\"2xZNMmSH5oZYW_db_asm_disk_5\"/>\n
<SharedDisk fromDefinition=\"true\" name=\"2xZNMmSH5oZYW_db_asm_disk_6\"/>\n
<SharedDisk fromDefinition=\"true\" name=\"2xZNMmSH5oZYW_db_asm_disk_7\"/>\n
</SharedDisks>\n <Nics>\n
<NetworkInterface fromDefinition=\"true\" name=\"eth0\">\n
<IPAssignmentMode>Dhcp</IPAssignmentMode>\n <QoS>Any_
Network_QoS_Type</QoS>\n <AssemblyNetwork
name=\"network-1\"/>\n </NetworkInterface>\n
</Nics>\n <RootPassword> </RootPassword>\n

```



```

<LogLocation>/assemblybuilder/logs</LogLocation>\n
<VmSize>Custom</VmSize>\n
<NetworkConfigurationTimeout>60</NetworkConfigurationTimeout>\n
<ProductConfigurationTimeout>60</ProductConfigurationTimeout>\n
</HardwareConfiguration>\n </VirtualSystemConfig>\n
</VirtualSystemCollectionConfig>\n
</AssemblyDeployment>\n</ns2:ConfigurationData>\n"
 } ,
 {
 "uri" :
"/em/cloud/iaas/servicetemplate/deploymentplan/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp%3AVirtualization%3ADeployment+Plan%3AC7790791F0831736E040578CDA815673%3A0.1" ,
 "name" : "sidb_deployplan" ,
 "media_type" : "application/oracle.com.cloud.iaas.DeploymentPlan+json" ,
 "type" : "Deployment Plan"
 }
]
}

```

The "deploy\_plans" attribute contains both the default deployment plan for each zone and the preconfigured deployment plans that are accessible by the user for this service template. To create an assembly instance, the deployment plan should be used.

The configuration specified in the "params" is the primary configuration. In addition, deployment plans can be used to describe detailed product and resource configurations. There are some product configurations where "<Required>true</Required>", and they should also be updated to capture the product configuration values.

The following shows an example of how to create an assembly instance:

**Table 26–21 POST Request for Assembly Instance Creation**

| Feature | Description                                                                                                                                                                                  |
|---------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| URL     | https://example.oracle.com/em/cloud/iaas/servicetemplate/assembly/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp%3AVirtualization%3AAssembly%3AC75E88B04D7FEDEDE040578CDA810E49%3A0.1 |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=<br>Content-Type:<br>application/oracle.com.cloud.common.AssemblyInstance+json                                                                  |
| Method  | POST                                                                                                                                                                                         |

The following shows the body:

```

{
 "zone": "/em/cloud/iaas/zone/D73AF0D42C8459E11419862797D1F37D",
 "params": {
 "vnc_password": "welcome1",
 "root_password": "welcome1",
 "assembly_instance_name": "AITEST"
 },
 "deployment_plan": "<?xml version='1.0' encoding='UTF-8'>
standalone='yes'><?xml version='1.0' encoding='UTF-8'>
xmlns:ns2='http://www.oracle.com/sysman/vt/RequestConfigData'>
<AssemblyDeployment assemblyInstanceName='DEFAULT_PLEASE_CHANGE'>
<SourceBinary type='TYPE_SWLIB'
name='oracle:defaultService:em:provisioning:1:cmp:Virtualization:Assembly:C75E88B04D7FEDEDE040578CDA810E49:0.1'>
<DeploymentTarget type='oracle_vm_

```

```
zone\" name=\"D73AF0D42C8459E11419862797D1F37D\"/>\n<AcceptedAllEULAs>>false</AcceptedAllEULAs>\n<InstallEMAgent>>false</InstallEMAgent>\n<DisableCleanup>>false</DisableCleanup>\n<EMAgentConfig installUserPassword=\" \" installUserName=\"oracle\"/>\n<AssemblyNetworkConfig>\n<AssemblyNetwork name=\"network-1\"/>\n</AssemblyNetworkConfig>\n<VirtualSystemCollectionConfig id=\"sidbasmA_linuxx64_11203psu1\">\n<ProductConfiguration>\n<Product>\n<Properties/>\n<EMAgent>>false</EMAgent>\n</Product>\n</ProductConfiguration>\n<VirtualSystemConfig agentPushInstallationEnabled=\"true\"\nagentInstallationType=\"Push Agent Installation\" vmInstance=\"true\"\ntargetName=\"sidbasmA_linuxx64_11203psu1/sidbasm:%assembly_instance_name%\"\nid=\"sidbasmA_linuxx64_11203psu1/sidbasm\">\n<ProductConfiguration>\n<Product>\n<Properties>\n<Property\nid=\"ocm.runConfiguration\">\n<Name>ocm.runConfiguration</Name>\n<Required>>false</\nRequired>\n<Value>>false</Value>\n<Secret>>false</Secret>\n</Property>\n<Property\nid=\"ocm.metalinkCsiRegistration.CSI\">\n<Name>ocm.metalinkCsiRegistration.CSI</Na\nme>\n<Required>>false</Required>\n<Value></Value>\n<Secret>>false</Secret>\n</Property>\n<Property\nid=\"ocm.metalinkCsiRegistration.metalinkId\">\n<Name>ocm.metalinkCsiRegistration.\nmetalinkId</Name>\n<Required>>false</Required>\n<Value></Value>\n<Secret>>false</Sec\nret>\n</Property>\n<Property\nid=\"ocm.proxyPassword\">\n<Name>ocm.proxyPassword</Name>\n<Required>>false</Requir\ned>\n<Value>\n</Value>\n<Secret>>true</Secret>\n</Property>\n<Property\nid=\"ocm.metalinkCsiRegistration.countryCode\">\n<Name>ocm.metalinkCsiRegistration\n.countryCode</Name>\n<Required>>false</Required>\n<Value></Value>\n<Secret>>false</S\necret>\n</Property>\n<Property id=\"system-files|db_asm_disk_\n7|size-units\">\n<Name>system-files|db_asm_disk_\n7|size-units</Name>\n<Required>>false</Required>\n<Value>MB</Value>\n<Secret>>false<\n/Secret>\n</Property>\n<Property id=\"system-files|db_asm_disk_\n1|size-units\">\n<Name>system-files|db_asm_disk_\n1|size-units</Name>\n<Required>>false</Required>\n<Value>MB</Value>\n<Secret>>false<\n/Secret>\n</Property>\n<Property id=\"system-files|db_asm_disk_\n3|size-units\">\n<Name>system-files|db_asm_disk_\n3|size-units</Name>\n<Required>>false</Required>\n<Value>MB</Value>\n<Secret>>false<\n/Secret>\n</Property>\n<Property id=\"system-files|db_asm_disk_0|size\">\n<Name>system-files|db_asm_
```

```

disk_
0|size</Name>\n<Required>>false</Required>\n<Value>2836</Value>\n<Secret>>false</Sec
ret>\n
</Property>\n
<Property id=\"system-fileset|db_asm_disk_1|size\">\n<Name>system-fileset|db_asm_
disk_
1|size</Name>\n<Required>>false</Required>\n<Value>2836</Value>\n<Secret>>false</Sec
ret>\n
</Property>\n
<Property
id=\"asm-password\">\n<Name>asm-password</Name>\n<Required>>true</Required>\n<Value
>welcome1</Value>\n<Secret>>true</Secret>\n
</Property>\n
<Property id=\"system-fileset|db_asm_disk_7|size\">\n<Name>system-fileset|db_asm_
disk_
7|size</Name>\n<Required>>false</Required>\n<Value>2836</Value>\n<Secret>>false</Sec
ret>\n
</Property>\n
<Property
id=\"ocm.metalinkEmailRegistration.metalinkEmailId\">\n<Name>ocm.metalinkEmailRegi
stration.metalinkEmailId</Name>\n<Required>>false</Required>\n<Value></Value>\n<Sec
ret>>false</Secret>\n
</Property>\n
<Property
id=\"input|listener-1|global-db-name\">\n<Name>input|listener-1|global-db-name</Na
me>\n<Required>>false</Required>\n<Value>orcl.us.oracle.com</Value>\n<Secret>>false<
/Secret>\n
</Property>\n
<Property id=\"system-fileset|db_asm_disk_
4|size-units\">\n<Name>system-fileset|db_asm_disk_
4|size-units</Name>\n<Required>>false</Required>\n<Value>MB</Value>\n<Secret>>false<
/Secret>\n
</Property>\n
<Property
id=\"ocm.proxyHost\">\n<Name>ocm.proxyHost</Name>\n<Required>>false</Required>\n<Va
lue></Value>\n<Secret>>false</Secret>\n
</Property>\n
<Property
id=\"ocm.repeaterURI\">\n<Name>ocm.repeaterURI</Name>\n<Required>>false</Required>\n
<Value></Value>\n<Secret>>false</Secret>\n
</Property>\n
<Property id=\"system-fileset|db_asm_disk_
2|size-units\">\n<Name>system-fileset|db_asm_disk_
2|size-units</Name>\n<Required>>false</Required>\n<Value>MB</Value>\n<Secret>>false<
/Secret>\n
</Property>\n
<Property
id=\"input|listener-1|port\">\n<Name>input|listener-1|port</Name>\n<Required>>false
</Required>\n<Value>1521</Value>\n<Secret>>false</Secret>\n
</Property>\n
<Property
id=\"ocm.anonymousEmailRegistration.emailId\">\n<Name>ocm.anonymousEmailRegistrati
on.emailId</Name>\n<Required>>false</Required>\n<Value></Value>\n<Secret>>false</Sec
ret>\n
</Property>\n
<Property
id=\"db-account-password\">\n<Name>db-account-password</Name>\n<Required>>true</Req
uired>\n<Value>welcome1</Value>\n<Secret>>true</Secret>\n
</Property>\n
<Property id=\"system-fileset|db_asm_disk_3|size\">\n<Name>system-fileset|db_asm_
disk_
3|size</Name>\n<Required>>false</Required>\n<Value>2836</Value>\n<Secret>>false</Sec
ret>\n
</Property>\n
<Property id=\"system-fileset|db_asm_disk_2|size\">\n<Name>system-fileset|db_asm_

```

```

disk_
2|size</Name>\n<Required>>false</Required>\n<Value>2836</Value>\n<Secret>>false</Sec
ret>\n
</Property>\n
<Property
id=\"ocm.proxyPort\">\n<Name>ocm.proxyPort</Name>\n<Required>>false</Required>\n<Va
lue></Value>\n<Secret>>false</Secret>\n
</Property>\n
<Property
id=\"ocm.proxyUsername\">\n<Name>ocm.proxyUsername</Name>\n<Required>>false</Requir
ed>\n<Value></Value>\n<Secret>>false</Secret>\n
</Property>\n
<Property id=\"system-fileset|db_asm_disk_6|size\">\n<Name>system-fileset|db_asm_
disk_
6|size</Name>\n<Required>>false</Required>\n<Value>2836</Value>\n<Secret>>false</Sec
ret>\n
</Property>\n
<Property id=\"system-fileset|db_asm_disk_5|size\">\n<Name>system-fileset|db_asm_
disk_
5|size</Name>\n<Required>>false</Required>\n<Value>2836</Value>\n<Secret>>false</Sec
ret>\n
</Property>\n
<Property id=\"system-fileset|db_asm_disk_4|size\">\n<Name>system-fileset|db_asm_
disk_
4|size</Name>\n<Required>>false</Required>\n<Value>2836</Value>\n<Secret>>false</Sec
ret>\n
</Property>\n
<Property
id=\"ocm.metalinkEmailRegistration.metalinkPassword\">\n<Name>ocm.metalinkEmailReg
istration.metalinkPassword</Name>\n<Required>>false</Required>\n<Value>
</Value>\n<Secret>>true</Secret>\n
</Property>\n
<Property id=\"system-fileset|db_asm_disk_
6|size-units\">\n<Name>system-fileset|db_asm_disk_
6|size-units</Name>\n<Required>>false</Required>\n<Value>MB</Value>\n<Secret>>false<
/Secret>\n
</Property>\n
<Property id=\"system-fileset|db_asm_disk_
5|size-units\">\n<Name>system-fileset|db_asm_disk_
5|size-units</Name>\n<Required>>false</Required>\n<Value>MB</Value>\n<Secret>>false<
/Secret>\n
</Property>\n
<Property id=\"system-fileset|db_asm_disk_
0|size-units\">\n<Name>system-fileset|db_asm_disk_
0|size-units</Name>\n<Required>>false</Required>\n<Value>MB</Value>\n<Secret>>false<
/Secret>\n
</Property>\n
</Properties>\n
<EMAgent>>false</EMAgent>\n
</Product>\n
</ProductConfiguration>\n
<HardwareConfiguration>\n
<Memory>2048</Memory>\n
<VCPU>1</VCPU>\n
<HaEnabled>>false</HaEnabled>\n
<StartAfterCreation>>true</StartAfterCreation>\n
<CPUSchedulingPriority>50</CPUSchedulingPriority>\n
<CPUCap>100</CPUCap>\n
<LocalDisks>\n
<disk fromDefinition=\"true\" name=\"2xZNMmSH5oZYW_System\">\n
<Size>5122</Size>\n
<Mode>Read-Write</Mode>\n

```

```

<CreatedByVmSize>>false</CreatedByVmSize>\n
</disk>\n
<disk fromDefinition=\"true\" name=\"AB\">\n <Size>0</Size>\n
<Mode>Read-Write</Mode>\n <CreatedByVmSize>>false</CreatedByVmSize>\n
</disk>\n <disk fromDefinition=\"true\" name=\"2xZNmSH5oZYW_sys-asm_base\">\n
<Size>502</Size>\n
<Mode>Read-Write</Mode>\n <CreatedByVmSize>>false</CreatedByVmSize>\n
</disk>\n <disk fromDefinition=\"true\" name=\"2xZNmSH5oZYW_sys-asm_home\">\n
<Size>10307</Size>\n
<Mode>Read-Write</Mode>\n <CreatedByVmSize>>false</CreatedByVmSize>\n
</disk>\n <disk fromDefinition=\"true\" name=\"2xZNmSH5oZYW_sys-db_base\">\n
<Size>502</Size>\n
<Mode>Read-Write</Mode>\n <CreatedByVmSize>>false</CreatedByVmSize>\n
</disk>\n <disk fromDefinition=\"true\" name=\"2xZNmSH5oZYW_sys-db_home\">\n
<Size>9068</Size>\n
<Mode>Read-Write</Mode>\n <CreatedByVmSize>>false</CreatedByVmSize>\n
</disk>\n </LocalDisks>\n
<SharedDisks>\n
<SharedDisk fromDefinition=\"true\" name=\"2xZNmSH5oZYW_db_asm_disk_0\"/>\n
<SharedDisk fromDefinition=\"true\" name=\"2xZNmSH5oZYW_db_asm_disk_1\"/>\n
<SharedDisk fromDefinition=\"true\" name=\"2xZNmSH5oZYW_db_asm_disk_2\"/>\n
<SharedDisk fromDefinition=\"true\" name=\"2xZNmSH5oZYW_db_asm_disk_3\"/>\n
<SharedDisk fromDefinition=\"true\" name=\"2xZNmSH5oZYW_db_asm_disk_4\"/>\n
<SharedDisk fromDefinition=\"true\" name=\"2xZNmSH5oZYW_db_asm_disk_5\"/>\n
<SharedDisk fromDefinition=\"true\" name=\"2xZNmSH5oZYW_db_asm_disk_6\"/>\n
<SharedDisk fromDefinition=\"true\" name=\"2xZNmSH5oZYW_db_asm_disk_7\"/>\n
</SharedDisks>\n
<Nics>\n
<NetworkInterface fromDefinition=\"true\" name=\"eth0\">\n
<IPAssignmentMode>Dhcp</IPAssignmentMode>\n
<QoS>Any_Network_QoS_Type</QoS>\n
<AssemblyNetwork name=\"network-1\"/>\n
</NetworkInterface>\n
</Nics>\n
<RootPassword>\n
</RootPassword>\n
<LogLocation>/assemblybuilder/logs</LogLocation>\n
<VmSize>Custom</VmSize>\n
<NetworkConfigurationTimeout>60</NetworkConfigurationTimeout>\n
<ProductConfigurationTimeout>60</ProductConfigurationTimeout>\n
</HardwareConfiguration>\n
</VirtualSystemConfig>\n
</VirtualSystemCollectionConfig>\n
</AssemblyDeployment>\n</ns2:ConfigurationData>\n"
}

```

The following is the Web service response after the request was successfully submitted:

```

{
 "uri" : "/em/cloud/iaas/assemblyrequest/223" ,
 "name" : "VDOSI Assembly Creation 1345575147280" ,
 "resource_state" : {
 "state" : "INITIATED" ,
 "messages" :
 [
 {
 "text" : "The Request with ID '223' is scheduled with Job Id 'C7CB0EBE83C22D70E040578CDA812E66'" ,
 "date" : "2012-08-21T18:52:31+0000"
 }
]
 }
}

```

```

]
 } ,
 "context_id" : "223" ,
 "media_type" : "application/oracle.com.cloud.common.AssemblyInstance+json" ,
 "service_family_type" : "iaas" ,
 "status" : "SCHEDULED" ,
 "created" : "2012-08-21T18:52:28+0000"
}

```

Similarly, tracking the process by getting the URI returns:

**Table 26–22 GET Request**

| Feature | Description                                                  |
|---------|--------------------------------------------------------------|
| URL     | https://example.oracle.com/em/cloud/iaas/assemblyrequest/223 |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=                |
| Method  | GET                                                          |
| Since   | 12.1.0.4 Cloud Service Portal plug-in                        |

The following is the response of the Web service response:

```

{
 "uri" : "/em/cloud/iaas/assemblyrequest/223" ,
 "name" : "VDOSI Assembly Creation 1345575147280" ,
 "resource_state" : {
 "state" : "CREATING" ,
 "messages" :
 [
 {
 "text" : "The Request with ID '223' is being processed, and has Job Id
'C7CB0EBE83C22D70E040578CDA812E66'" ,
 "date" : "2012-08-21T18:54:25+0000"
 }
]
 } ,
 "context_id" : "223" ,
 "media_type" : "application/oracle.com.cloud.common.AssemblyInstance+json" ,
 "service_family_type" : "iaas" ,
 "status" : "EXECUTING" ,
 "created" : "2012-08-21T18:52:28+0000"
}

```

Soon after the assembly instance is created successfully.

**Table 26–23 GET Request**

| Feature | Description                                                  |
|---------|--------------------------------------------------------------|
| URL     | https://example.oracle.com/em/cloud/iaas/assemblyrequest/223 |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=                |
| Method  | GET                                                          |
| Since   | 12.1.0.4 Cloud Service Portal plug-in                        |

The following is the response of the Web service response:

```

{
 "uri" : "/em/cloud/iaas/assembly/C0AACCD8234BBAFBA7AF6CD3865B1FD3" ,

```

```

"name" : "AITEST" ,
"resource_state" : {
 "state" : "READY"
} ,
"context_id" : "223" ,
"media_type" : "application/oracle.com.cloud.common.AssemblyInstance+json" ,
"service_family_type" : "iaas" ,
"based_on" :
"/em/cloud/iaas/servicetemplate/assembly/oracle%3AdefaultService%3Aem%3Aprovisioni
ng%3A1%3Acmp%3AVirtualization%3AAssembly%3AC75E88B04D7FEDEDE040578CDA810E49%3A0.1"
,
"reference" : {
 "ovf:id" : "sidbasmA_linuxx64_11203psu1"
} ,
"scalability_groups" : {
 "media_type" : "application/oracle.com.cloud.common.ScalabilityGroup+json" ,
 "total" : "0" ,
 "elements" :
[
]
} ,
"servers" : {
 "media_type" : "application/oracle.com.cloud.common.VM+json" ,
 "total" : "1" ,
 "elements" :
[
 {
 "uri" : "/em/cloud/iaas/server/vm/B28680F24303A42D6909FFC0F79F306C" ,
 "name" : "sidbasmA_linuxx64_11203psu1/sidbasM:DEPPLAN_1345575147280" ,
 "media_type" : "application/oracle.com.cloud.common.VM+json" ,
 "status" : "STARTED"
 }
]
} ,
"deployment_plan" : "<?xml version='1.0' encoding='UTF-8'
standalone='yes'?>\n<ns2:ConfigurationData
xmlns:ns2='http://www.oracle.com/sysman/vt/RequestConfigData'>\n
<AssemblyDeployment assemblyInstanceName='AITEST'>\n
 <SourceBinary
type='TYPE_SWLIB'
name='oracle:defaultService:em:provisioning:1:cmp:Virtualization:Assembly:C75E88B
04D7FEDEDE040578CDA810E49:0.1'/>\n
 <DeploymentTarget type='oracle_vm_
zone' name='D73AF0D42C8459E11419862797D1F37D'/>\n
<AcceptedAllEULAs>false</AcceptedAllEULAs>\n
<InstallEMAgent>false</InstallEMAgent>\n
<DisableCleanup>false</DisableCleanup>\n
 <EMAgentConfig
installUserPassword=' ' installUserName='oracle'/>\n
<AssemblyNetworkConfig>\n
 <AssemblyNetwork name='network-1'/>\n
</AssemblyNetworkConfig>\n
 <VirtualSystemCollectionConfig id='sidbasmA_
linuxx64_11203psu1'>\n
 <ProductConfiguration>\n
 <Product>\n
 <Properties/>\n
 </Product>\n
 </ProductConfiguration>\n
 <VirtualSystemConfig
agentPushInstallationEnabled='true' agentInstallationType='Push Agent
Installation' vmInstance='true' targetName='sidbasmA_linuxx64_
11203psu1/sidbasM:%assembly_instance_name%' id='sidbasmA_linuxx64_
11203psu1/sidbasM'>\n
 <ProductConfiguration>\n
 <Product>\n
 <Properties>\n
 <Property
id='ocm.runConfiguration'>\n<Name>ocm.runConfiguration</Name>\n<Required>false</
Required>\n<Value>false</Value>\n<Secret>false</Secret>\n
 </Property>

```

```

id=\ "ocm.metalinkCsiRegistration.CSI\">\n<Name>ocm.metalinkCsiRegistration.CSI</Name>
<Required>>false</Required>\n<Value></Value>\n<Secret>>false</Secret>\n
</Property>\n
<Property
id=\ "ocm.metalinkCsiRegistration.metalinkId\">\n<Name>ocm.metalinkCsiRegistration.
metalinkId</Name>\n<Required>>false</Required>\n<Value></Value>\n<Secret>>false</Sec
ret>\n
</Property>\n
<Property
id=\ "ocm.proxyPassword\">\n<Name>ocm.proxyPassword</Name>\n<Required>>false</Requir
ed>\n<Value> </Value>\n<Secret>>true</Secret>\n
</Property>\n
<Property
id=\ "ocm.metalinkCsiRegistration.countryCode\">\n<Name>ocm.metalinkCsiRegistration
.countryCode</Name>\n<Required>>false</Required>\n<Value></Value>\n<Secret>>false</S
ecret>\n
</Property>\n
<Property id=\ "system-fileset|db_asm_disk_
7|size-units\">\n<Name>system-fileset|db_asm_disk_
7|size-units</Name>\n<Required>>false</Required>\n<Value>MB</Value>\n<Secret>>false<
/Secret>\n
</Property>\n
<Property id=\ "system-fileset|db_asm_disk_
1|size-units\">\n<Name>system-fileset|db_asm_disk_
1|size-units</Name>\n<Required>>false</Required>\n<Value>MB</Value>\n<Secret>>false<
/Secret>\n
</Property>\n
<Property id=\ "system-fileset|db_asm_disk_
3|size-units\">\n<Name>system-fileset|db_asm_disk_
3|size-units</Name>\n<Required>>false</Required>\n<Value>MB</Value>\n<Secret>>false<
/Secret>\n
</Property>\n
<Property id=\ "system-fileset|db_asm_disk_0|size\">\n<Name>system-fileset|db_asm_
disk_
0|size</Name>\n<Required>>false</Required>\n<Value>2836</Value>\n<Secret>>false</Sec
ret>\n
</Property>\n
<Property id=\ "system-fileset|db_asm_disk_1|size\">\n<Name>system-fileset|db_asm_
disk_
1|size</Name>\n<Required>>false</Required>\n<Value>2836</Value>\n<Secret>>false</Sec
ret>\n
</Property>\n
<Property
id=\ "asm-password\">\n<Name>asm-password</Name>\n<Required>>true</Required>\n<Value
>
</Value>\n<ValueGuid>C7CC61AC07DC6E58E040578CDA8145F9</ValueGuid>\n<Secret>>true</S
ecret>\n
</Property>\n
<Property id=\ "system-fileset|db_asm_disk_7|size\">\n<Name>system-fileset|db_asm_
disk_
7|size</Name>\n<Required>>false</Required>\n<Value>2836</Value>\n<Secret>>false</Sec
ret>\n
</Property>\n
<Property
id=\ "ocm.metalinkEmailRegistration.metalinkEmailId\">\n<Name>ocm.metalinkEmailRegi
stration.metalinkEmailId</Name>\n<Required>>false</Required>\n<Value></Value>\n<Sec
ret>>false</Secret>\n
</Property>\n
<Property
id=\ "input|listener-1|global-db-name\">\n<Name>input|listener-1|global-db-name</Na
me>\n<Required>>false</Required>\n<Value>orcl.us.oracle.com</Value>\n<Secret>>false<
/Secret>\n
</Property>\n
<Property id=\ "system-fileset|db_asm_disk_
4|size-units\">\n<Name>system-fileset|db_asm_disk_
4|size-units</Name>\n<Required>>false</Required>\n<Value>MB</Value>\n<Secret>>false<
/Secret>\n
</Property>\n
<Property
id=\ "ocm.proxyHost\">\n<Name>ocm.proxyHost</Name>\n<Required>>false</Required>\n<Va
lue></Value>\n<Secret>>false</Secret>\n
</Property>\n
<Property
id=\ "ocm.repeaterURI\">\n<Name>ocm.repeaterURI</Name>\n<Required>>false</Required>\n
<Value></Value>\n<Secret>>false</Secret>\n

```



```

</Property>\n
<Property id=\"system-fileset|db_asm_
disk_2|size-units\">\n<Name>system-fileset|db_asm_disk_
2|size-units</Name>\n<Required>>false</Required>\n<Value>MB</Value>\n<Secret>>false<
/Secret>\n
</Property>\n
<Property
id=\"input|listener-1|port\">\n<Name>input|listener-1|port</Name>\n<Required>>false
</Required>\n<Value>1521</Value>\n<Secret>>false</Secret>\n
</Property>\n
<Property
id=\"ocm.anonymousEmailRegistration.emailId\">\n<Name>ocm.anonymousEmailRegistrati
on.emailId</Name>\n<Required>>false</Required>\n<Value></Value>\n<Secret>>false</Sec
ret>\n
</Property>\n
<Property
id=\"db-account-password\">\n<Name>db-account-password</Name>\n<Required>>true</Req
uired>\n<Value>
</Value>\n<ValueGuid>C7CC61AC07DF6E58E040578CDA8145F9</ValueGuid>\n<Secret>>true</S
ecret>\n
</Property>\n
<Property id=\"system-fileset|db_asm_disk_3|size\">\n<Name>system-fileset|db_asm_
disk_
3|size</Name>\n<Required>>false</Required>\n<Value>2836</Value>\n<Secret>>false</Sec
ret>\n
</Property>\n
<Property id=\"system-fileset|db_asm_disk_2|size\">\n<Name>system-fileset|db_asm_
disk_
2|size</Name>\n<Required>>false</Required>\n<Value>2836</Value>\n<Secret>>false</Sec
ret>\n
</Property>\n
<Property
id=\"ocm.proxyPort\">\n<Name>ocm.proxyPort</Name>\n<Required>>false</Required>\n<N<Va
lue></Value>\n<Secret>>false</Secret>\n
</Property>\n
<Property
id=\"ocm.proxyUsername\">\n<Name>ocm.proxyUsername</Name>\n<Required>>false</Requir
ed>\n<Value></Value>\n<Secret>>false</Secret>\n
</Property>\n
<Property id=\"system-fileset|db_asm_
disk_6|size\">\n<Name>system-fileset|db_asm_disk_
6|size</Name>\n<Required>>false</Required>\n<Value>2836</Value>\n<Secret>>false</Sec
ret>\n
</Property>\n
<Property id=\"system-fileset|db_asm_disk_5|size\">\n<Name>system-fileset|db_asm_
disk_
5|size</Name>\n<Required>>false</Required>\n<Value>2836</Value>\n<Secret>>false</Sec
ret>\n
</Property>\n
<Property id=\"system-fileset|db_asm_disk_4|size\">\n<Name>system-fileset|db_asm_
disk_
4|size</Name>\n<Required>>false</Required>\n<Value>2836</Value>\n<Secret>>false</Sec
ret>\n
</Property>\n
<Property
id=\"ocm.metalinkEmailRegistration.metalinkPassword\">\n<Name>ocm.metalinkEmailReg
istration.metalinkPassword</Name>\n<Required>>false</Required>\n<Value>
</Value>\n<Secret>true</Secret>\n
</Property>\n
<Property id=\"system-fileset|db_asm_disk_
6|size-units\">\n<Name>system-fileset|db_asm_disk_
6|size-units</Name>\n<Required>>false</Required>\n<Value>MB</Value>\n<Secret>>false<
/Secret>\n
</Property>\n
<Property id=\"system-fileset|db_asm_disk_
5|size-units\">\n<Name>system-fileset|db_asm_disk_
5|size-units</Name>\n<Required>>false</Required>\n<Value>MB</Value>\n<Secret>>false<
/Secret>\n
</Property>\n
<Property id=\"system-fileset|db_asm_disk_
0|size-units\">\n<Name>system-fileset|db_asm_disk_
0|size-units</Name>\n<Required>>false</Required>\n<Value>MB</Value>\n<Secret>>false<
/Secret>\n
</Property>\n
</Properties>\n
<EMAgent>>false</EMAgent>\n
</Product>\n
</ProductConfiguration>\n

```

```

<HardwareConfiguration>\n
 <Memory>2048</Memory>\n
 <VCPUs>1</VCPUs>\n
 <HaEnabled>>false</HaEnabled>\n
 <StartAfterCreation>>true</StartAfterCreation>\n
 <CPUSchedulingPriority>50</CPUSchedulingPriority>\n
 <CPUCap>100</CPUCap>\n
 <LocalDisks>\n
 <disk fromDefinition="true" name="2xZNMmSH5oZYW_System">\n
 <Size>5122</Size>\n
 <Mode>Read-Write</Mode>\n
 <CreatedByVmSize>>false</CreatedByVmSize>\n
 </disk>\n
 <disk fromDefinition="true" name="AB">\n
 <Size>0</Size>\n
 <Mode>Read-Write</Mode>\n
 <CreatedByVmSize>>false</CreatedByVmSize>\n
 </disk>\n
 <disk fromDefinition="true" name="2xZNMmSH5oZYW_sys-asm_base">\n
 <Size>502</Size>\n
 <Mode>Read-Write</Mode>\n
 <CreatedByVmSize>>false</CreatedByVmSize>\n
 </disk>\n
 <disk fromDefinition="true" name="2xZNMmSH5oZYW_sys-asm_home">\n
 <Size>10307</Size>\n
 <Mode>Read-Write</Mode>\n
 <CreatedByVmSize>>false</CreatedByVmSize>\n
 </disk>\n
 <disk fromDefinition="true" name="2xZNMmSH5oZYW_sys-db_base">\n
 <Size>502</Size>\n
 <Mode>Read-Write</Mode>\n
 <CreatedByVmSize>>false</CreatedByVmSize>\n
 </disk>\n
 <disk fromDefinition="true" name="2xZNMmSH5oZYW_sys-db_home">\n
 <Size>9068</Size>\n
 <Mode>Read-Write</Mode>\n
 <CreatedByVmSize>>false</CreatedByVmSize>\n
 </disk>\n
 </LocalDisks>\n
 <SharedDisks>\n
 <SharedDisk fromDefinition="true" name="2xZNMmSH5oZYW_db_asm_disk_0">\n
 <SharedDisk fromDefinition="true" name="2xZNMmSH5oZYW_db_asm_disk_1">\n
 <SharedDisk fromDefinition="true" name="2xZNMmSH5oZYW_db_asm_disk_2">\n
 <SharedDisk fromDefinition="true" name="2xZNMmSH5oZYW_db_asm_disk_3">\n
 <SharedDisk fromDefinition="true" name="2xZNMmSH5oZYW_db_asm_disk_4">\n
 <SharedDisk fromDefinition="true" name="2xZNMmSH5oZYW_db_asm_disk_5">\n
 <SharedDisk fromDefinition="true" name="2xZNMmSH5oZYW_db_asm_disk_6">\n
 <SharedDisk fromDefinition="true" name="2xZNMmSH5oZYW_db_asm_disk_7">\n
 </SharedDisks>\n
 <Nics>\n
 <NetworkInterface fromDefinition="true" name="eth0">\n
 <IPAssignmentMode>Dhcp</IPAssignmentMode>\n
 <QoS>Any_
 <Network_QoS_Type</QoS>\n
 <AssemblyNetwork
 name="network-1">\n
 </NetworkInterface>\n
 </Nics>\n
 <RootPassword>
 </RootPassword>\n
 <RootPasswordGuid>C7CC61AC07D96E58E040578CDA8145F9</RootPasswordGuid>\n
 <LogLocation>/assemblybuilder/logs</LogLocation>\n
 <VmSize>Custom</VmSize>\n
 <NetworkConfigurationTimeout>60</NetworkConfigurationTimeout>\n
 <ProductConfigurationTimeout>60</ProductConfigurationTimeout>\n
 </HardwareConfiguration>\n
 </VirtualSystemConfig>\n
 </VirtualSystemCollectionConfig>\n
 </AssemblyDeployment>\n</ns2:ConfigurationData>\n"
}

```

## 26.2.2 Updating a Service Instance

A service instance may be resized, its status updated, or additional resources added.

### 26.2.2.1 Updating a VM Resource

The VM resource supports the following changes:

- The "status" can be updated (for example, from STARTED to STOPPED or from STOPPED to STARTED).
- The "cpu" and "memory" can be updated (for example to resize the VM).

For the 12.1.0.4 Cloud Self Service Portal release, VM is no longer required to be STOPPED before modifying the "cpu" and "memory".

- Any updates supported on the VM resource from 10000 version are also supported on the 10001 version of the resource.

### Change VM Status

The following shows the GET of the VM's status:

**Table 26–24 GET Request for VM Status**

| Feature | Description                                                              |
|---------|--------------------------------------------------------------------------|
| URL     | https://example.oracle.com/em/cloud/iaas/server/byrequest/101?status,uri |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=                            |
| Method  | GET                                                                      |
| Since   | 12.1.0.1 Cloud Service Portal plug-in                                    |

The following is the response:

```
{
 "uri" : "/em/cloud/iaas/server/vm/C5CA59AF9386975E8AEA45B0F040F095" ,
 "status" : "STARTED"
}
```

To update the status, use the PUT method to the resource "uri". Note that the "uri" attribute is the one that directly references to the resource.

**Table 26–25 PUT Request for VM Status**

| Feature | Description                                                                                                |
|---------|------------------------------------------------------------------------------------------------------------|
| URL     | https://example.oracle.com/em/cloud/iaas/server/vm/C5CA59AF9386975E8AEA45B0F040F095                        |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=<br>Content-Type: application/oracle.com.cloud.common.VM+json |
| Body    | {<br>"status" : "STOPPED"<br>}                                                                             |
| Method  | PUT                                                                                                        |
| Since   | 12.1.0.1 Cloud Service Portal plug-in                                                                      |

The Web service response will return the 200 response code, along with the current VM attributes.

The following shows the GET request on the VM:

**Table 26–26 GET Request on VM**

| Feature | Description                                                              |
|---------|--------------------------------------------------------------------------|
| URL     | https://example.oracle.com/em/cloud/iaas/server/byrequest/101?status,uri |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=                            |
| Method  | GET                                                                      |

**Table 26–26 (Cont.) GET Request on VM**

| Feature | Description                           |
|---------|---------------------------------------|
| Since   | 12.1.0.1 Cloud Service Portal plug-in |

The Web service returns the following:

```
{
 "uri" : "/em/cloud/iaas/server/vm/C5CA59AF9386975E8AEA45B0F040F095" ,
 "status" : "STOPPED"
}
```

Attempting to submit a PUT transient status value would result in the following:

**Table 26–27 PUT Transient Status Value**

| Feature | Description                                                                                                |
|---------|------------------------------------------------------------------------------------------------------------|
| URL     | https://example.oracle.com/em/cloud/iaas/server/vm/C5CA59AF9386975E8AEA45B0F040F095                        |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=<br>Content-Type: application/oracle.com.cloud.common.VM+json |
| Body    | {       "status" : "STOPPING"     }                                                                        |
| Method  | PUT                                                                                                        |
| Since   | 12.1.0.1 Cloud Service Portal plug-in                                                                      |

The following shows the Web service response:

```
{
 "messages" :
 [
 {
 "date" : "2012-08-19T17:22:04+0000" ,
 "stack_trace" :
 "oracle.sysman.emInternalSDK.ssa.cloudapi.rest.RestControllerException: Status
 'STOPPING' specified by the PUT request is transient, and is not a valid
 status\n\tat
 oracle.sysman.ssa.cloudapi.iaas.VM.processVMStatusChange(VM.java:331)\n\tat
 oracle.sysman.ssa.cloudapi.iaas.VM.processVMChanges(VM.java:253)\n\tat
 oracle.sysman.ssa.cloudapi.iaas.VM.processRequest(VM.java:227)\n\tat
 oracle.sysman.ssa.cloudapi.iaas.IaaSServiceProvider.processRequest(IaaSServiceProv
 ider.java:582)\n\tat
 oracle.sysman.emInternalSDK.ssa.cloudapi.EMCloudServlet.perform(EMCloudServlet.jav
 a:236)\n\tat
 oracle.sysman.emInternalSDK.ssa.cloudapi.EMCloudServlet.performPut(EMCloudServlet.
 java:351)\n\tat
 oracle.sysman.emInternalSDK.ssa.cloudapi.rest.AbstractRestController.doPut(AbstractRe
 stServlet.java:163)\n\tat
 javax.servlet.http.HttpServlet.service(HttpServlet.java:730)\n\tat
 javax.servlet.http.HttpServlet.service(HttpServlet.java:820)\n\tat
 weblogic.servlet.internal.StubSecurityHelper$ServletServiceAction.run(StubSecurity
 Helper.java:227)\n\tat we" ,
 "hint" : "Status 'STOPPING' specified by the PUT request is transient, and
 is not a valid status" ,
 "text" : "Did not successfully execute 'PUT' resource operation on
```

```
'application/oracle.com.cloud.common.Cloud' identified by
'/em/cloud/iaas/server/vm/C5CA59AF9386975E8AEA45B0F040F095' "
 }
]
}
```

Attempting to submit the same PUT status to the VM will result in the following interaction:

**Table 26–28 PUT Status to the VM**

| Feature | Description                                                                                                |
|---------|------------------------------------------------------------------------------------------------------------|
| URL     | https://example.oracle.com/em/cloud/iaas/server/vm/C5CA59AF9386975E8AEA45B0F040F095?resource_state         |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=<br>Content-Type: application/oracle.com.cloud.common.VM+json |
| Body    | {<br>"status" : "STOPPED"<br>}                                                                             |
| Method  | PUT                                                                                                        |
| Since   | 12.1.0.1 Cloud Service Portal plug-in                                                                      |

The web service will response with:

```
{
 "resource_state" : {
 "state" : "READY" ,
 "messages" :
 [
 {
 "date" : "2012-08-19T17:23:49+0000" ,
 "text" : "Current Virtual Machine Status 'STOPPED' is the same as the
 requested status. Status change is not needed"
 }
]
 }
}
```

When modifying a service instance, the messages communicated will be encapsulated in the resource\_state attribute. For example, to START the VM again, the following PUT request can be issued:

**Table 26–29 PUT Request**

| Feature | Description                                                                                                |
|---------|------------------------------------------------------------------------------------------------------------|
| URL     | https://example.oracle.com/em/cloud/iaas/server/vm/C5CA59AF9386975E8AEA45B0F040F095?resource_state         |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=<br>Content-Type: application/oracle.com.cloud.common.VM+json |
| Body    | {<br>"status" : "STARTED"<br>}                                                                             |
| Method  | PUT                                                                                                        |

**Table 26–29 (Cont.) PUT Request**

| Feature | Description                           |
|---------|---------------------------------------|
| Since   | 12.1.0.1 Cloud Service Portal plug-in |

The web service will respond with the following:

```
{
 "resource_state" : {
 "state" : "READY" ,
 "messages" :
 [
 {
 "date" : "2012-08-19T17:25:29+0000" ,
 "text" : "'START' resource operation on
'application/oracle.com.cloud.common.VM' identified by
'C5CA59AF9386975E8AEA45B0F040F095' is successfully submitted with reference
'C7A2713B9A4D29C1E040578CDA817561'"
 }
]
 }
}
```

### Changing VM Sizes

The CPU and Memory of a VM can be adjusted to increase or decrease the capacity. For example, to adjust the VM's CPU to 1 vCPU and increase the memory from 512 to 1024, the following PUT request can be issued:

**Table 26–30 Changing VM Sizes**

| Feature | Description                                                                                                                                                                                                         |
|---------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| URL     | <a href="https://example.oracle.com/em/cloud/iaas/server/vm/C5CA59AF9386975E8AEA45B0F040F095?resource_state">https://example.oracle.com/em/cloud/iaas/server/vm/C5CA59AF9386975E8AEA45B0F040F095?resource_state</a> |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=<br>Content-Type: application/oracle.com.cloud.common.VM+json                                                                                                          |
| Body    | {       "cpu": [1,0],       "memory": "1024"     }                                                                                                                                                                  |
| Method  | PUT                                                                                                                                                                                                                 |
| Since   | 12.1.0.1 Cloud Service Portal plug-in                                                                                                                                                                               |

The following is returned by the Web service:

```
{
 "resource_state" : {
 "state" : "READY" ,
 "messages" :
 [
 {
 "text" : "'PUT' resource operation on
'application/oracle.com.cloud.common.VM' identified by
'C5CA59AF9386975E8AEA45B0F040F095' is successfully submitted with reference '106'"
 },
 {
 "date" : "2012-08-20T01:15:28+0000"
 }
]
 }
}
```

```

 }
]
}

```

After some period of time, the new capacity is reflected in the VM resource through the GET request:

**Table 26–31 GET Request**

| Feature | Description                                                                                                          |
|---------|----------------------------------------------------------------------------------------------------------------------|
| URL     | https://example.oracle.com/em/cloud/iaas/server/vm/C5CA59AF9386975E8AEA45B0F040F095?resource_state,cpu,memory,status |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=                                                                        |
| Method  | GET                                                                                                                  |

The Web service returns the following:

```

{
 "resource_state" : {
 "state" : "READY"
 } ,
 "status" : "STARTED" ,
 "cpu" :
 [
 "1" ,
 "0"
] ,
 "memory" : "1024"
}

```

## 26.2.3 Adding a VM Disk

An additional disk may be added to the VM resource. For example, to add two additional disks, the following PUT request can be issued:

**Table 26–32 Add VM Disk**

| Feature | Description                                                                                                |
|---------|------------------------------------------------------------------------------------------------------------|
| URL     | https://example.oracle.com/em/cloud/iaas/server/vm/C5CA59AF9386975E8AEA45B0F040F095?resource_state         |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=<br>Content-Type: application/oracle.com.cloud.common.VM+json |
| Body    | {<br>"disks": [<br>["additional_1", "5120"],<br>["additional_2", "10240"]<br>]<br>}                        |
| Method  | PUT                                                                                                        |
| Since   | 12.1.0.5 Cloud Service Portal plug-in                                                                      |

If an entry in the "disks" attribute is new, then a new disk would be added. If an entry in the "disks" attribute exists already, then, it would be ignored. Here is the web service response:

```
{
 "resource_state" : {
 "state" : "READY" ,
 "messages" :
 [
 {
 "text" : "'PUT' resource operation on
'application/oracle.com.cloud.common.VM' identified by '
C5CA59AF9386975E8AEA45B0F040F095' is successfully submitted with reference '2' " ,
 "date" : "2013-02-26T19:53:15+0000"
 }
]
 }
}
```

After some periods of time, the new disks are reflected in the VM resource using the GET request.

**Table 26–33    GET Request**

| Feature | Description                                                                                              |
|---------|----------------------------------------------------------------------------------------------------------|
| URL     | https://example.oracle.com/em/cloud/iaas/server/vm/C5CA59AF9386975E8AEA45B0F040F095?resource_state,disks |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=                                                            |
| Method  | GET                                                                                                      |

Here is the Web service response:

```
{
 "resource_state" : {
 "state" : "READY"
 } ,
 "disks" :
 [
 [
 "system.img (2)" ,
 "1400"
] ,
 [
 "default_disk" ,
 "10240"
] ,
 [
 "additional_1" ,
 "5120"
] ,
 [
 "additional_2" ,
 "10240"
]
]
}
```

**26.2.4 Deleting a Service Instance**

The DELETE request on the URI of the resource can be issued to delete a Service Instance. For example, the following DELETE request is issued:



**Table 26–34 Delete Request**

| Feature | Description                                                                         |
|---------|-------------------------------------------------------------------------------------|
| URL     | https://example.oracle.com/em/cloud/iaas/server/vm/C5CA59AF9386975E8AEA45B0F040F095 |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=                                       |
| Method  | DELETE                                                                              |
| Since   | 12.1.0.1 Cloud Service Portal plug-in                                               |

The following Web service response with code 200 is returned:

```
{
 "uri" : "/em/cloud/iaas/server/vm/C5CA59AF9386975E8AEA45B0F040F095" ,
 "resource_state" : {
 "state" : "READY" ,
 "messages" :
 [
 {
 "text" : "'DELETE' resource operation on
'application/oracle.com.cloud.common.VM' identified by
'C5CA59AF9386975E8AEA45B0F040F095' is successfully submitted with reference
'C7A83335CB63DA7BE040578CDA814DDC'" ,
 "date" : "2012-08-20T01:24:38+0000"
 }
]
 } ,
 "context_id" : "C5CA59AF9386975E8AEA45B0F040F095" ,
 "media_type" : "application/oracle.com.cloud.common.VM+json" ,
 "service_family_type" : "iaas"
}
```

After some time, the GET request on the VM URI will result in the following interaction:

**Table 26–35 GET Request on the VM URI**

| Feature | Description                                                                         |
|---------|-------------------------------------------------------------------------------------|
| URL     | https://example.oracle.com/em/cloud/iaas/server/vm/C5CA59AF9386975E8AEA45B0F040F095 |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=                                       |
| Method  | GET                                                                                 |

The Web service returns the following:

```
{
 "uri" : "/em/cloud/iaas/server/vm/C5CA59AF9386975E8AEA45B0F040F095" ,
 "resource_state" : {
 "state" : "READY"
 } ,
 "context_id" : "C5CA59AF9386975E8AEA45B0F040F095" ,
 "media_type" : "application/oracle.com.cloud.common.VM+json" ,
 "service_family_type" : "iaas" ,
 "status" : "DELETED"
}
```

Note the "status" of the VM has the "DELETED" value.

## 26.2.5 Listing Service Instances

There are various references to service instances in the resource model, and they are all in context of the resource themselves. For example, "service\_instances" in the Zone resource would list the elements that are in that Zone, while "service\_instances" in the ServiceTemplate resource would list the elements that are based on that service template.

For example, the following list the service instances in a zone.

**Table 26–36 GET Request**

| Feature | Description                                                                                                   |
|---------|---------------------------------------------------------------------------------------------------------------|
| URL     | <code>https://example.oracle.com/em/cloud/iaas/zone/D73AF0D42C8459E11419862797D1F37D?service_instances</code> |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=                                                                 |
| Method  | GET                                                                                                           |
| Since   | 12.1.0.4 Cloud Service Portal plug-in                                                                         |

The following response is returned:

```
{
 "service_instances" : {
 "media_type" : "application/oracle.com.cloud.common.ServiceInstance+json" ,
 "total" : "2" ,
 "elements" :
 [
 {
 "uri" : "/em/cloud/iaas/server/vm/523CAE80A305928C9C5BE8A67A4181FD" ,
 "name" : "ZONEPOST" ,
 "media_type" : "application/oracle.com.cloud.common.VM+json" ,
 "status" : "STARTED"
 } ,
 {
 "uri" : "/em/cloud/iaas/server/vm/C5CA59AF9386975E8AEA45B0F040F095" ,
 "name" : "STPOST" ,
 "media_type" : "application/oracle.com.cloud.common.VM+json" ,
 "status" : "STARTED"
 }
]
 }
}
```

Similarly, you can list the service instances from the service templates:

**Table 26–37 GET Request**

| Feature | Description                                                                                                                                                                                                           |
|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| URL     | <code>https://example.oracle.com/em/cloud/iaas/servicetemplate/vm/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp%3AVirtualization%3ATemplate%3AC76CEB5563EA5E13E040578CDA817FAF%3A0.1?service_instances</code> |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=                                                                                                                                                                         |
| Method  | GET                                                                                                                                                                                                                   |
| Since   | 12.1.0.4 Cloud Service Portal plug-in                                                                                                                                                                                 |

The following response is returned:

```
{
 "service_instances" : {
 "media_type" : "application/oracle.com.cloud.common.ServiceInstance+json" ,
 "total" : "2" ,
 "elements" :
 [
 {
 "uri" : "/em/cloud/iaas/server/vm/C5CA59AF9386975E8AEA45B0F040F095" ,
 "name" : "STPOST" ,
 "media_type" : "application/oracle.com.cloud.common.VM+json" ,
 "status" : "STARTED"
 } ,
 {
 "uri" : "/em/cloud/iaas/server/vm/523CAE80A305928C9C5BE8A67A4181FD" ,
 "name" : "ZONEPOST" ,
 "media_type" : "application/oracle.com.cloud.common.VM+json" ,
 "status" : "STARTED"
 }
]
 }
}
```

## 26.2.6 Searching Service Templates

Searching for a service template through cloud can be performed by constructing ServiceTemplateFinds resource. ServiceTemplateFinds is synonymous to a "report" that contains elements satisfying the condition specified.

For example, the following will return service template with the name OEL in it

**Table 26–38 POST Request**

| Feature | Description                                                                                                                                                                                               |
|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| URL     | https://example.oracle.com/em/cloud                                                                                                                                                                       |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=<br>Content-Type:<br>application/oracle.com.cloud.common.ServiceTemplateFinds+json<br>Accept : application/oracle.com.cloud.common.ServiceTemplateFinds+json |
| Body    | {           "filters":{             "name":"OEL"           }         }                                                                                                                                    |
| Method  | POST                                                                                                                                                                                                      |
| Since   | 12.1.0.5 Cloud Service Portal plug-in                                                                                                                                                                     |

Here is the response:

```
{
 "uri" : "/em/cloud/finds/service_
template/bWVkaWFfdHlwZSUyNTNEYXBwbGljYXRpb241MjUyRm9yYWNSZS5jb20uY2xvdWQuY29tbW9uL
lNl%0AcnZpY2VUZWlwbGF0ZSUyRm5hbWU1MjUzRE9FTA%3D%3D" ,
 "name" : "Service Template Finds" ,
 "media_type" : "application/oracle.com.cloud.common.ServiceTemplateFinds+json" ,
 "description" : "This is a Resource Finds for 'Service Template Finds' with
```

```

media type 'application/oracle.com.cloud.common.ServiceTemplateFinds+json' ,
 "resource_state" : {
 "state" : "READY"
 } ,
 "filters" : {
 "media_type" : "application/oracle.com.cloud.common.ServiceTemplate" ,
 "name" : "OEL"
 } ,
 "finds" : {
 "media_type" : "application/oracle.com.cloud.common.ServiceTemplate" ,
 "total" : "1" ,
 "elements" :
 [
 {
 "uri" :
"/em/cloud/iaas/servicetemplate/vm/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp%3AVirtualization%3ATemplate%3AD6A66B2DCAEAE177E040F20AB0527449%3A0.1" ,
 "name" : "OEL4.0 For Testing" ,
 "media_type" : "application/oracle.com.cloud.common.VMTemplate+json" ,
 "type" : "Template"
 }
]
 }
}

```

Currently, the following attributes are supported by the "filters":

**Table 26–39 Filters**

| Attributes | Description                                                                                                                                                           | Since                                                                |
|------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------|
| media_type | The specific service template media to search for. This would be an exact match of the media type.                                                                    | 12.1.0.5 Cloud Service Portal Plugin (IaaS service family type only) |
| name       | A service template would be returned if the name contains the specified string, case insensitive (so, OEL and oel in the example above would return the same result). | 12.1.0.5 Cloud Service Portal Plugin (IaaS service family type only) |

If an attribute specified is not recognized by the system, it will be ignored. For example, the following will return service templates containing the name DUMMY:

**Table 26–40 Unrecognized Attribute**

| Feature | Description                                                                                                                                                                                               |
|---------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| URL     | https://example.oracle.com/em/cloud                                                                                                                                                                       |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=<br>Content-Type:<br>application/oracle.com.cloud.common.ServiceTemplateFinds+json<br>Accept : application/oracle.com.cloud.common.ServiceTemplateFinds+json |
| Body    | {       "filters":{         "name": "DUMMY"       }     }                                                                                                                                                 |
| Method  | POST                                                                                                                                                                                                      |

**Table 26–40 (Cont.) Unrecognized Attribute**

| Feature | Description                           |
|---------|---------------------------------------|
| Since   | 12.1.0.5 Cloud Service Portal plug-in |

Here is the response:

```
{
 "uri" : "/em/cloud/finds/service_
template/bWVkaWFfdHlwZSUyNTNEYXBwbGljYXRpb241MjUyRm9yYWNSZS5jb20uY2xvdWQuY29tbW9uL
lNl%0AcnZpY2VUZW1wbGF0ZSUyRm5hbWU1MjUzRERVTU1Z" ,
 "name" : "Service Template Finds" ,
 "media_type" : "application/oracle.com.cloud.common.ServiceTemplateFinds+json" ,
 "description" : "This is a Resource Finds for 'Service Template Finds' with
media type 'application/oracle.com.cloud.common.ServiceTemplateFinds+json'" ,
 "resource_state" : {
 "state" : "READY"
 } ,
 "filters" : {
 "media_type" : "application/oracle.com.cloud.common.ServiceTemplate" ,
 "name" : "DUMMY"
 } ,
 "finds" : {
 "media_type" : "application/oracle.com.cloud.common.ServiceTemplate" ,
 "total" : "0" ,
 "elements" :
 [
]
 }
}
```

Since there were no service templates with "DUMMY" as part of its name, no service templates are returned.

Please note, once the resource ServiceTemplateFinds is created, its URI can be used for performing the same query by issuing a GET request. If new templates were added that also satisfy the conditions, it would be returned as well.

For example, let's use the DUMMY example's URI, and this time, a service template with "DUMMY" in its name was created.

**Table 26–41 Unrecognized Attribute**

| Feature | Description                                                                                                                                                                                    |
|---------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| URL     | https://slc03rqn.us.oracle.com:4473/em/cloud/finds/service_ template/bWVkaWFfdHlwZSUyNTNEYXBwbGljYXRpb241MjUyRm9yYW NsZS5jb20uY2xvdWQuY29tbW9uLlNl%0AcnZpY2VUZW1wbGF0ZSUyRm5 hbWU1MjUzRERVTU1Z |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=                                                                                                                                                  |
| Method  | GET                                                                                                                                                                                            |
| Since   | 12.1.0.5 Cloud Service Portal plug-in                                                                                                                                                          |

Here is the response of the GET request after the service template was created and published by administrators:

```
{
 "uri" : "/em/cloud/finds/service_
template/bWVkaWFfdHlwZSUyNTNEYXBwbGljYXRpb241MjUyRm9yYWNSZS5jb20uY2xvdWQuY29tbW9uL
lNl%0AcnZpY2VUZW1wbGF0ZSUyRm5hbWU1MjUzRERVTU1Z" ,
```

```

 "name" : "Service Template Finds" ,
 "media_type" : "application/oracle.com.cloud.common.ServiceTemplateFinds+json" ,
 "description" : "This is a Resource Finds for 'Service Template Finds' with
media type 'application/oracle.com.cloud.common.ServiceTemplateFinds+json' " ,
 "resource_state" : {
 "state" : "READY"
 } ,
 "filters" : {
 "media_type" : "application/oracle.com.cloud.common.ServiceTemplate" ,
 "name" : "DUMMY"
 } ,
 "finds" : {
 "media_type" : "application/oracle.com.cloud.common.ServiceTemplate" ,
 "total" : "1" ,
 "elements" :
 [
 {
 "uri" :
"/em/cloud/iaas/servicetemplate/assembly/oracle%3AdefaultService%3Aem%3Aprovisioni
ng%3A1%3Acmp%3AVirtualization%3AAssembly%3AD6A7D3D84B605D11E040F20AB0524E1C%3A0.1"
,
 "name" : "Dummy OEL Assembly" ,
 "media_type" :
"application/oracle.com.cloud.common.AssemblyTemplate+json" ,
 "type" : "Assembly"
 }
]
 }
}

```

## 26.3 Support for Version 1000

As the document previous noted, this 12.1.0.4 Cloud Service Portal Plugin supports the specification version 10000. To start, issue the following HTTP request on the Enterprise Manager:

**Table 26–42 HTTP Request**

| Feature | Description                                                                                                                               |
|---------|-------------------------------------------------------------------------------------------------------------------------------------------|
| URL     | https://example.oracle.com/em/cloud                                                                                                       |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=<br>X-Specification-Version: 10000<br>Accept: application/oracle.com.cloud.common.Cloud+json |
| Method  | GET                                                                                                                                       |

The following returns the resource as in 10000 version:

```

{
 "uri" : "/em/cloud" ,
 "name" : "Infrastructure Cloud" ,
 "description" : "Infrastructure Cloud with Self Service Portal and
Virtualization Managements" , "service_templates" : {
 "uri" : "/em/cloud/collection/servicetemplates" ,
 "type" : "application/oracle.com.cloud.common.ServiceTemplate+json" ,
 "total" : "7" ,
 "elements" :
 [

```

```

 {
 "uri" :
 "/em/cloud/servicetemplate/oracle:defaultService:em:provisioning:1:cmp:Virtualizat
ion:Assembly:C75E88B04D7FEDEDE040578CDA810E49:0.1" ,
 "name" : "sidb_assembly" ,
 "type" : "Assembly" ,
 "version" : "0.1"
 } ,
 {
 "uri" :
 "/em/cloud/servicetemplate/oracle:defaultService:em:provisioning:1:cmp:Virtualizat
ion:Assembly:C769B1F361529309E040578CDA813D57:0.1" ,
 "name" : "wls_assembly" ,
 "type" : "Assembly" ,
 "version" : "0.1"
 } ,
 {
 "uri" :
 "/em/cloud/servicetemplate/oracle:defaultService:em:provisioning:1:cmp:Virtualizat
ion:Template:C76CEB5563EA5E13E040578CDA817FAF:0.1" ,
 "name" : "template_sanity" ,
 "type" : "Template" ,
 "version" : "0.1"
 } ,
 {
 "uri" :
 "/em/cloud/servicetemplate/oracle:defaultService:em:provisioning:1:cmp:Virtualizat
ion:Assembly:C76F733BC7A41AF7E040578CDA812CDC:0.1" ,
 "name" : "fmw_venkat" ,
 "type" : "Assembly" ,
 "version" : "0.1"
 } ,
 {
 "uri" :
 "/em/cloud/servicetemplate/oracle:defaultService:em:provisioning:1:cmp:Virtualizat
ion:Assembly:C76C144A4A245B62E040578CDA8163B9:0.1" ,
 "name" : "fmw_abby" ,
 "type" : "Assembly" ,
 "version" : "0.1"
 } ,
 {
 "uri" :
 "/em/cloud/servicetemplate/oracle:defaultService:em:provisioning:1:cmp:Virtualizat
ion:Assembly:C77122B0A916D95CE040578CDA814854:0.1" ,
 "name" : "sidbasmA_abby" ,
 "type" : "Assembly" ,
 "version" : "0.1"
 } ,
 {
 "uri" :
 "/em/cloud/servicetemplate/oracle:defaultService:em:provisioning:1:cmp:Virtualizat
ion:Assembly:C76C8792DE2A0937E040578CDA81795E:0.1" ,
 "name" : "WLS_abby" ,
 "type" : "Assembly" ,
 "version" : "0.1"
 }
]
} ,
"zones" : {
 "uri" : "/em/cloud/collection/zones" ,

```

```
"type" : "application/oracle.com.cloud.common.Zone+json" ,
"total" : "1" ,
"elements" :
[
 {
 "uri" : "/em/cloud/zone/D73AF0D42C8459E11419862797D1F37D" ,
 "name" : "cloud_zone" ,
 "type" : "oracle_vm_zone"
 }
]
} ,
"vdcs" : {
 "uri" : "/em/cloud/collection/vdcs" ,
 "type" : "application/oracle.com.cloud.common.VDC+json" ,
 "total" : "1" ,
 "elements" :
 [
 {
 "uri" : "/em/cloud/vdc/default/D73AF0D42C8459E11419862797D1F37D" ,
 "name" : "Oracle Enterprise Manager Virtual Data Center on Zone
cloud_zone"
 }
]
} ,
"instance_options" : {
 "uri" : "/em/cloud/collection/instanceoptions" ,
 "type" : "application/oracle.com.cloud.common.InstanceOption+json" ,
 "total" : "3" ,
 "elements" :
 [
 {
 "uri" : "/em/cloud/instanceoption/virtual_machine_size/2" ,
 "name" : "Medium" ,
 "type" : "virtual_machine_size" ,
 "memory" : "8192" ,
 "cpu" : "4" ,
 "local_storage" : "512000"
 } ,
 {
 "uri" : "/em/cloud/instanceoption/virtual_machine_size/3" ,
 "name" : "Large" ,
 "type" : "virtual_machine_size" ,
 "memory" : "15360" ,
 "cpu" : "8" ,
 "local_storage" : "1024000"
 } ,
 {
 "uri" : "/em/cloud/instanceoption/virtual_machine_size/1" ,
 "name" : "Small" ,
 "type" : "virtual_machine_size" ,
 "memory" : "4096" ,
 "cpu" : "2" ,
 "local_storage" : "256000"
 }
]
} ,
"resource_state" : {
 "state" : "READY"
}
}
```



Similarly, to return the list of resources in the VDC, the following can be issued (note that X-specification-version must be included in the header whenever working with 10000 version resources):

**Table 26–43 Request to Return VDC Resources**

| Feature | Description                                                                                                                             |
|---------|-----------------------------------------------------------------------------------------------------------------------------------------|
| URL     | <code>https://example.oracle.com/em/cloud/vdc/default/D73AF0D42C8459E11419862797D1F37D</code>                                           |
| Headers | Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=<br>X-Specification-Version: 10000<br>Accept: application/oracle.com.cloud.common.VDC+json |
| Method  | GET                                                                                                                                     |

The following VDC resource is returned:

```
{
 "uri" : "/em/cloud/vdc/default/D73AF0D42C8459E11419862797D1F37D" ,
 "name" : "Oracle Enterprise Manager Virtual Data Center on Zone cloud_zone" ,
 "description" : "Virtual Data Center is a logical grouping of virtualized assets
in a particular zone" ,
 "scalability_groups" : {
 "uri" : "/em/cloud/vdc/default/D73AF0D42C8459E11419862797D1F37D?scalability_
groups" ,
 "type" : "application/oracle.com.cloud.common.ScalabilityGroup+json" ,
 "total" : "0" ,
 "elements" :
 [
]
 } ,
 "servers" : {
 "uri" : "/em/cloud/vdc/default/D73AF0D42C8459E11419862797D1F37D?servers" ,
 "type" : "application/oracle.com.cloud.common.Server+json" ,
 "total" : "0" ,
 "elements" :
 [
]
 } ,
 "vnets" : {
 "uri" : "/em/cloud/vdc/default/D73AF0D42C8459E11419862797D1F37D?vnets" ,
 "type" : "application/oracle.com.cloud.common.VNet+json" ,
 "total" : "3" ,
 "elements" :
 [
 {
 "uri" :
"/em/cloud/vnet/D73AF0D42C8459E11419862797D1F37D/oracle:defaultService:em:provisio
ning:1:netConfig:C76CEADBBBE6B23FE040578CDA817FB1" ,
 "name" : "venkat_profile" ,
 "id" :
"oracle:defaultService:em:provisioning:1:netConfig:C76CEADBBBE6B23FE040578CDA817FB
1"
 } ,
 {
 "uri" :
"/em/cloud/vnet/D73AF0D42C8459E11419862797D1F37D/oracle:defaultService:em:provisio
ning:1:netConfig:C77076C8FDEC6BD7E040578CDA813B2B" ,
 "name" : "nwprofile_abby" ,
 "id" :
"oracle:defaultService:em:provisioning:1:netConfig:C77076C8FDEC6BD7E040578CDA813B2

```

```

B"
 },
 {
 "uri" :
"/em/cloud/vnet/D73AF0D42C8459E11419862797D1F37D/oracle:defaultService:em:provisio
ning:1:netConfig:C76F741AFD7EB760E040578CDA812CD8" ,
 "name" : "bmp_profile" ,
 "id" :
"/em/cloud/vnet/D73AF0D42C8459E11419862797D1F37D/oracle:defaultService:em:provisioning:1:netConfig:C76F741AFD7EB760E040578CDA812CD8"
 }
]
},
"zone" : "/em/cloud/zone/D73AF0D42C8459E11419862797D1F37D" ,
"resource_state" : {
 "state" : "READY"
},
"assembly_instances" : {
 "uri" : "/em/cloud/vdc/default/D73AF0D42C8459E11419862797D1F37D?assemblies"
},
"type" : "application/oracle.com.cloud.common.AssemblyInstance+json" ,
"total" : "11" ,
"elements" :
[
 {
 "uri" : "/em/cloud/assembly/byrequest/64" ,
 "name" : "WLS_abby" ,
 "contained_in" :
"/em/cloud/vdc/default/D73AF0D42C8459E11419862797D1F37D" ,
 "guid" : "FD4F3945CDD6BD6EEFFAB064735ECDD"
 },
 {
 "uri" : "/em/cloud/assembly/byrequest/81" ,
 "name" : "WLS_abby_0" ,
 "contained_in" :
"/em/cloud/vdc/default/D73AF0D42C8459E11419862797D1F37D" ,
 "guid" : "79006B5C36F6ED90FC82EB58BB9580F2"
 },
 {
 "uri" : "/em/cloud/assembly/byrequest/82" ,
 "name" : "fmw_abby" ,
 "contained_in" :
"/em/cloud/vdc/default/D73AF0D42C8459E11419862797D1F37D" ,
 "guid" : "C7852B2D7740F4FAE040578CDA811EF3"
 },
 {
 "uri" : "/em/cloud/assembly/byrequest/26" ,
 "name" : "latest_attempt_wls_withagent_dinesh" ,
 "contained_in" :
"/em/cloud/vdc/default/D73AF0D42C8459E11419862797D1F37D" ,
 "guid" : "4DE0D2D8C81D0C09E665484BA9DBE668"
 },
 {
 "uri" : "/em/cloud/assembly/byrequest/61" ,
 "name" : "sidb_assembly" ,
 "contained_in" :
"/em/cloud/vdc/default/D73AF0D42C8459E11419862797D1F37D" ,
 "guid" : "235D95A93ABA2B50AFF5486540F48690"
 }
]

```

```

 "uri" : "/em/cloud/assembly/byrequest/62" ,
 "name" : "sidb_assembly_0" ,
 "contained_in" :
"/em/cloud/vdc/default/D73AF0D42C8459E11419862797D1F37D" ,
 "guid" : "BFA65EC9A5FB97AD2B971EE8FB5E91AC"
 } ,
 {
 "uri" : "/em/cloud/assembly/byrequest/27" ,
 "name" : "venkat_fmww" ,
 "contained_in" :
"/em/cloud/vdc/default/D73AF0D42C8459E11419862797D1F37D" ,
 "guid" : "3DBE707B981DEDA5380ED6DBD8BF399C"
 } ,
 {
 "uri" : "/em/cloud/assembly/byrequest/23" ,
 "name" : "venkat_nw_profile" ,
 "contained_in" :
"/em/cloud/vdc/default/D73AF0D42C8459E11419862797D1F37D" ,
 "guid" : "05C1F96EE6FBD049BFE1D57AA89CB0D0"
 } ,
 {
 "uri" : "/em/cloud/assembly/byrequest/1" ,
 "name" : "venkat_sidb" ,
 "contained_in" :
"/em/cloud/vdc/default/D73AF0D42C8459E11419862797D1F37D" ,
 "guid" : "3A5C3917D18E276BDC7D08949F1851D1"
 } ,
 {
 "uri" : "/em/cloud/assembly/byrequest/22" ,
 "name" : "wls_dinesh_32bit_agent" ,
 "contained_in" :
"/em/cloud/vdc/default/D73AF0D42C8459E11419862797D1F37D" ,
 "guid" : "896B27C0963AF3B5C5C69D62614D63ED"
 } ,
 {
 "uri" : "/em/cloud/assembly/byrequest/21" ,
 "name" : "wls_vt_dinesh" ,
 "contained_in" :
"/em/cloud/vdc/default/D73AF0D42C8459E11419862797D1F37D" ,
 "guid" : "E54B51C4FCFF6C56AE393C3B1EEA8464"
 }
]
} ,
"dbplatform_instances" : {
 "uri" : "/em/cloud/vdc/default/D73AF0D42C8459E11419862797D1F37D?dbplatform_
instances" ,
 "type" : "application/oracle.com.cloud.common.DbPlatformInstance+json" ,
 "total" : "0" ,
 "elements" :
 [
]
}
}

```

If the x-specification-Version header was not specified, the following would be returned:

```

{
 "messages" :
 [
 {
 "text" : "Media Type 'application/oracle.com.cloud.common.VDC' specified

```

```
in the request cannot be supported by the Cloud Resource. Cloud Resource supports
Media Type 'application/oracle.com.cloud.common.Cloud+json' ,
 "date" : "2012-08-19T15:17:08+0000"
 }
]
}
```

---

---

**Note:** Only Infrastructure Service resources are supported on 10000 version. Features supported are the same as the previous version.

---

---

---

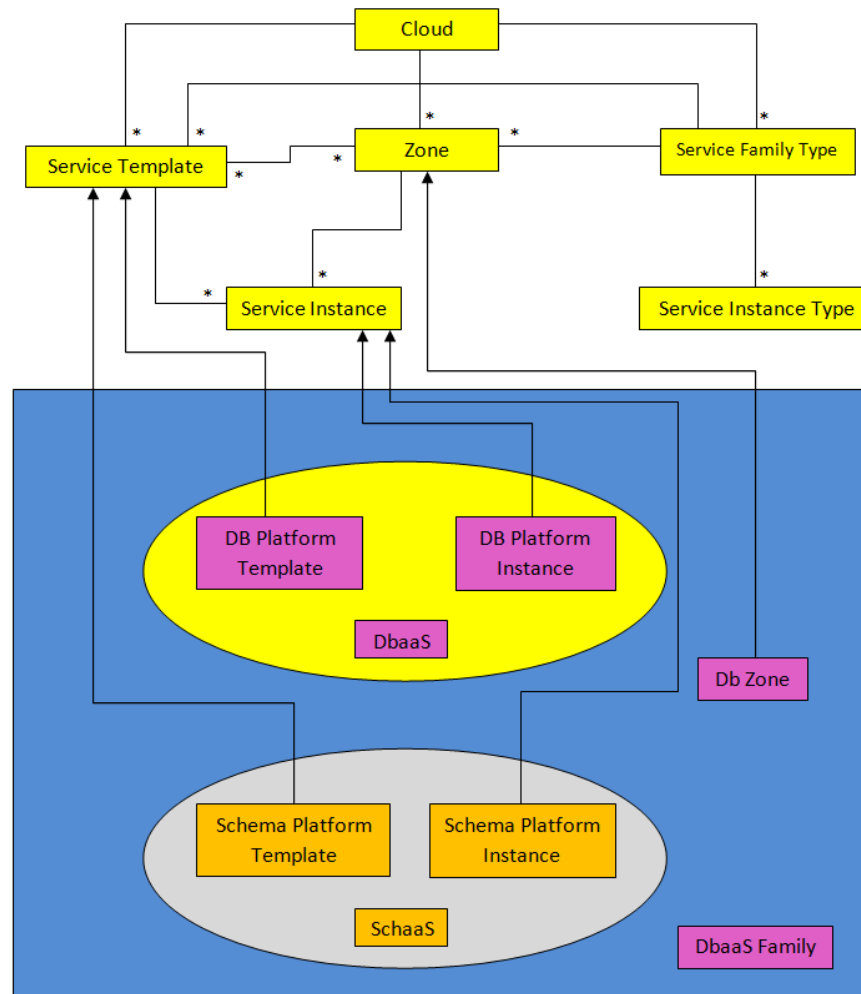
## Database as a Service Family APIs

This chapter describes the resource models for Database as a Service (DBaaS) API, Schema as a Service APIs, and the REST request/response interactions. The following topics are covered:

- [Resource Model for Database as a Service](#)
- [Supported Operations for DBaaS Resources](#)
- [Database as a Service API Examples](#)
- [Resource Model for Schema as a Service](#)
- [Supported Operations for Schema as a Service Resources](#)
- [Schema as a Service API Examples](#)
- [Using EMCLI to Create Database Profiles](#)

### 27.1 Resource Model for Database as a Service

Cloud resources are organized by common service entities and specific services (for example, Database as a Service) provide resources by extending these common service entities. The figure below shows the DBaaS resource model:

**Figure 27–1 Database as a Service Resource Model**

The following sections describe the resources that are supported by the Database as a Service family type and provide examples of supported operations. The following resource models are described.

- [DB Zone](#)
- [DBPlatformTemplate](#)
- [DBPlatformInstance](#)

### 27.1.1 DB Zone

This extends `application/oracle.com.cloud.common.Zone`. This resource represents the view of a PaaS Infrastructure Zone that supports DBaaS resources. DB Zone has the following media type and payload:

- **Media Type:** `application/oracle.com.cloud.common.DbZone+json`
- **Supported Payload:** json
- **URI Format:** `/em/cloud/dbaas/zone/<zone id>`

The following table describes the DB Zone Data Model.

**Table 27–1 DB Zone Data Model**

| Field               | Type                         | Occurs | Description                                                                                                                                     |
|---------------------|------------------------------|--------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| uri                 | URI                          | 1      | A GET against this URI refreshes the client representation of the Zone definition to this user.                                                 |
| name                | String                       | 1      | Name of the Zone. SHALL be a UNICODE string to support different languages.                                                                     |
| description         | String                       | 0..1   | Human readable description of the Zone. SHALL be a UNICODE string to support different languages.                                               |
| resource_state      | ResourceState                | 1      | The resource state of the resource.                                                                                                             |
| context_id          | String                       | 1      | Zone GUID                                                                                                                                       |
| resource_state      | ResourceState                | 1      | The resource state of the resource.                                                                                                             |
| service_family_type | String                       | 1      | The name of the service family type that this zone is associated with.<br>It is "dbaas" for this.                                               |
| templates           | Collection <ServiceTemplate> | 1      | Collection of the service templates that this zone supports and can be deployed into. Currently, DBPlatformTemplate is the only supported type. |
| service_instances   | Collection <ServiceInstance> | 1      | Collection of the service instances that are created in this zone.                                                                              |
| media_type          | String                       | 1      | The media type of the resource.                                                                                                                 |
| canonicalLink       | URI                          | 1      | Can be used to perform cloud interactions like GET and POST.                                                                                    |

### 27.1.2 DBPlatformTemplate

The DBPlatformTemplate extends the ServiceTemplate resource. This resource represents service templates created for creating databases. The databases may be created using profiles created using database templates, rman backups, and SnapClone. It is permissible to accept the ServiceTemplate media type on an instance of DBPlatformTemplate. It is a preconfigured deployable service that realizes a DBPlatformInstance resource. DBPlatformTemplate has the following media type and payload:

- **Media Type:** application/oracle.com.cloud.common.DbPlatformTemplate+json
- **Supported Payload:** json
- **URI Format:** /em/cloud/dbaas/dbplatformtemplate/<template id>

The following table describes the DBPlatformTemplate Data Model.

**Table 27–2 DBPlatformTemplate Data Model**

| Field       | Type   | Occurs | Description                                                                                          |
|-------------|--------|--------|------------------------------------------------------------------------------------------------------|
| uri         | URI    | 1      | A GET against this URI refreshes the client representation of the resources accessible to this user. |
| name        | String | 1      | A human readable name given to the template.                                                         |
| description | String | 0..1   | A brief description given to the template                                                            |

**Table 27–2 (Cont.) DBPlatformTemplate Data Model**

| Field                        | Type                           | Occurs | Description                                                                                                                                                         |
|------------------------------|--------------------------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| resource_state               | ResourceState                  | 1      | The validity of the fields on a GET should be guaranteed only when the resource state is READY. Otherwise, the client should not assume the validity of the fields. |
| service_family_type          | String                         | 1      | Denotes the type of Service Template. For example, "dbaas".                                                                                                         |
| service_instances            | Collection<DBPlatformInstance> | 1      | All service instances created based on this DBPlatformTemplate.                                                                                                     |
| media_type                   | String                         | 1      | The media type of the service template that this type represents.                                                                                                   |
| created                      | Timestamp                      | 1      | Date and time, in ISO 8601 format, when the template was created/last modified.                                                                                     |
| zones                        | Collection<DbZone>             | 1      | List of all zones on which this service template is published                                                                                                       |
| db_type                      | String                         | 1      | Type of the Database when provisioned. It can have one of the two values: "oracle_database" and "rac_database".                                                     |
| instance_configurable_params | String[]                       | 1      | List of parameters required to provision a DbPlatformInstance from this template. Username and Password in this case.                                               |
| type                         | String                         | 1      | The type of the service template which is "dbaas" by default.                                                                                                       |
| node_count                   |                                | 0..1   | Number of Nodes. Implicit 1 for SI DB and explicit count for the RAC Db types.                                                                                      |
| canonicalLink                | URI                            | 1      | Can be used to perform cloud interactions like GET and POST                                                                                                         |
| deployment_params            | List<Object>                   | 1      | Description of the parameters (username and password) which are required while creating a service instance.                                                         |
| subtype                      | String                         | 1      | The sub type of the instance created based on this template.<br>It has a value "db" in this case.                                                                   |
| context_id                   | URI                            | 1      | The unique id of the template.                                                                                                                                      |

### 27.1.3 DBPlatformInstance

The DBPlatformInstance extends the ServiceInstance resource. It is linked to an Oracle Database Instance or Cluster Database target in Enterprise Manager Cloud Control. DBPlatformInstance has the following media type and payload:

This section contains the following sections:

- **Media Type:** application/oracle.com.cloud.common.DbPlatformInstance+json
- **Supported Payload:** json
- **URI Format:** /em/cloud/dbaas/dbplatforminstance/byrequest/<request id>

The following table describes the DBPlatformInstance Data Model.



**Table 27–3 DBPlatformInstance Data Model**

| Field            | Type               | Occurs | Description                                                                                                                                                         |
|------------------|--------------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| uri              | URI                | 1      | A GET against this URI refreshes the client representation of the resources accessible to this user.                                                                |
| name             | String             | 1      | A human readable name given to the instance.<br>[POST]                                                                                                              |
| destination_zone | Zone               | 1      | The Zone to which the instance is deployed.<br>[POST]                                                                                                               |
| media_type       | String             | 1      | Indicates the additional media type that clients can use to perform a GET.                                                                                          |
| available_space  | Float              | 0..1   | Available space for this database instance.                                                                                                                         |
| resource_state   | ResourceState      | 1      | The validity of the fields on a GET should be guaranteed only when the resource state is READY. Otherwise, the client should not assume the validity of the fields. |
| based_on         | DbPlatformTemplate | 1      | The db platform template on which this instance is published.                                                                                                       |
| connect_string   | String             | 0..1   | Connect String required to establish a connection to this database instance.                                                                                        |
| created          | String             | 1      | Time of creation.                                                                                                                                                   |
| type             | String             | 0..1   | Type of the database : "oracle_database" for SIDB and "rac_database" for RAC                                                                                        |
| status           | String             | 1      | Status of the Database instance.                                                                                                                                    |
| db_version       | String             | 0..1   | Version of the Database.                                                                                                                                            |
| last_backup      | Timestamp          | 0..1   | Time at which the last backup of this database was taken.                                                                                                           |
| load             | Float              | 0..1   | Current average active sessions of this instance.                                                                                                                   |
| master_username  | String             | 1      | Name of the master user of this database.                                                                                                                           |
| total_sessions   | Integer            | 0..1   | Total number of sessions for this database                                                                                                                          |
| total_sga        | Float              | 0..1   | Total SGA for this database.                                                                                                                                        |
| up_time          | String             | 1      | Uptime in hours for the Service Instance.                                                                                                                           |
| context_id       | String             | 1      | The request id of this instance.                                                                                                                                    |
| canonicalLink    | URI                | 1      | Can be used to perform cloud interactions like GET and DELETE for this instance.                                                                                    |
| instances        | List<Object>       | 0..1   | List all the node level details of the database instance (exists only when the type is "rac_database")                                                              |
| total_memory*    | String             | 0..1   | Total memory used by the instance (node).                                                                                                                           |
| active_sessions* | Float              | 0..1   | Current average active sessions for this instance (node).                                                                                                           |
| cpu_utilization* | Percentage         | 0..1   | CPU utilization of this instance (node).                                                                                                                            |
| db_time*         | Float              | 0..1   | Database time of this instance (node).                                                                                                                              |
| cpu_count*       | Integer            | 0..1   | Number of CPU cores on the host on which the instance (node) is deployed.                                                                                           |

**Table 27–3 (Cont.) DBPlatformInstance Data Model**

| Field               | Type    | Occurs | Description                                                             |
|---------------------|---------|--------|-------------------------------------------------------------------------|
| allocated_sessions* | Integer | 0..1   | Number of sessions currently allocated.                                 |
| name*               | String  | 0..1   | Name of the node (exists as an field in the list of "instances" field). |

**Note:** \* In the case of database instance type "oracle\_database", these fields exist as it is in the DBPlatformInstance Data model. In the case of type "rac\_database", these fields exist as part of "instances" field and occur for each node of the database instance.

## 27.2 Supported Operations for DBaaS Resources

The following table lists all the operations that are supported on the various Database as a Service resources.

**Table 27–4 Supported Operations**

| Resource                                | Operations (GET, POST, PUT, and DELETE) |
|-----------------------------------------|-----------------------------------------|
| Cloud                                   | GET                                     |
| Service Family Type                     | GET                                     |
| DB Zone                                 | GET, POST                               |
| DB Platform Template [Service Template] | GET, POST                               |
| DB Platform Instance [Service Instance] | GET, DELETE                             |

## 27.3 Database as a Service API Examples

The following sections provide examples of interaction with Enterprise Manager Cloud Control 12.1.0.2 with Cloud Application plug-in 12.1.0.4 and higher. The following examples are covered:

- [Cloud](#)
- [Filtering Output for Specific Resource Attributes](#)
- [Service Family Type Resource](#)
- [Service Instance Type resource](#)
- [Zone](#)
- [DB Platform Template Resource](#)
- [Creating Databases](#)
- [Polling the Database Creation](#)
- [Deleting a Database Instance](#)

### 27.3.1 Cloud

Use the top level /em/cloud to introspect the Cloud resource.

- URL

<https://example.us.oracle.com/em/cloud>

- Headers

Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=X-specification-Version: 10001

- Method

GET

The following returns the Cloud resource:

```
{
 "uri": "/em/cloud",
 "name": "Oracle Cloud by Enterprise Manager",
 "description": "This represents the Cloud resource of the Oracle Enterprise
Manager Cloud Management solution",
 "resource_state": {
 "state": "READY"
 },
 "media_type": "application/oracle.com.cloud.common.Cloud+json",
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud",
 "zones": {
 "media_type": "application/oracle.com.cloud.common.Zone+json",
 "total": "4",
 "elements": [
 {
 "uri": "/em/cloud/jaas/zone/9708491651C090AB4B847C2F3CD177DE",
 "name": "SLC00ECXY_ZONE",
 "media_type": "application/oracle.com.cloud.jaas.Zone+json",
 "service_family_type": "jaas",
 "type": "jaas",
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/zone/9708491
651C090AB4B847C2F3CD177DE"
 },
 {
 "uri": "/em/cloud/jaas/zone/B329742A13BE0D63703BF0C0E4F98FB9",
 "name": "SLC01NBV_ZONE",
 "media_type": "application/oracle.com.cloud.jaas.Zone+json",
 "service_family_type": "jaas",
 "type": "jaas",
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/zone/B329742
A13BE0D63703BF0C0E4F98FB9"
 },
 {
 "uri": "/em/cloud/dbaas/zone/9708491651C090AB4B847C2F3CD177DE",
 "name": "SLC00ECXY_ZONE",
 "description": "",
 "media_type": "application/oracle.com.cloud.common.DbZone+json",
 "service_family_type": "dbaas",
 "type": "self_service_zone",
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/zone/970849
1651C090AB4B847C2F3CD177DE"
 },
 {
 "uri": "/em/cloud/dbaas/zone/B329742A13BE0D63703BF0C0E4F98FB9",
 "name": "SLC01NBV_ZONE",
 "description": "SLC01NBV_ZONE",

```

```
 "media_type": "application/oracle.com.cloud.common.DbZone+json",
 "service_family_type": "dbaas",
 "type": "self_service_zone",
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/zone/B32974
2A13BE0D63703BF0C0E4F98FB9"
 }
],
},
"service_templates": {
 "media_type": "application/oracle.com.cloud.common.ServiceTemplate+json",
 "total": "2",
 "elements": [
 {
 "uri":
"/em/cloud/dbaas/dbplatformtemplate/DEC94FE6D49DAA21E040F20A60511D14",
 "name": "SLC01NBV_SNAPCLONE_ST",
 "description": "",
 "media_type":
"application/oracle.com.cloud.common.DbPlatformTemplate+json",
 "service_family_type": "dbaas",
 "type": "dbaas",
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/dbplatformt
emplate/DEC94FE6D49DAA21E040F20A60511D14"
 },
 {
 "uri":
"/em/cloud/dbaas/dbplatformtemplate/DECCA0E18357F28DE040F20A60514D49",
 "name": "SLC00ECXY_DBCA_ST",
 "description": "",
 "media_type":
"application/oracle.com.cloud.common.DbPlatformTemplate+json",
 "service_family_type": "dbaas",
 "type": "dbaas",
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/dbplatformt
emplate/DECCA0E18357F28DE040F20A60514D49"
 }
]
},
"service_family_types": {
 "media_type": "application/oracle.com.cloud.common.ServiceFamilyType+json",
 "total": "4",
 "elements": [
 {
 "uri": "/em/cloud/service_family_type/jaas",
 "name": "jaas",
 "media_type":
"application/oracle.com.cloud.common.ServiceFamilyType+json",
 "type": "jaas",
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/service_family_
type/jaas"
 },
 {
 "uri": "/em/cloud/service_family_type/dbaas",
 "name": "dbaas",
 "media_type":
"application/oracle.com.cloud.common.ServiceFamilyType+json",
```

```

 "type": "dbaas",
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/service_family_
type/dbaas"
 },
 {
 "uri": "/em/cloud/service_family_type/opc",
 "name": "opc",
 "media_type":
"application/oracle.com.cloud.common.ServiceFamilyType+json",
 "type": "opc",
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/service_family_
type/opc"
 },
 {
 "uri": "/em/cloud/service_family_type/iaas",
 "name": "iaas",
 "media_type":
"application/oracle.com.cloud.iaas.IaasServiceFamilyType+json",
 "type": "iaas",
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/service_family_
type/iaas"
 }
]
},
"service_requests": {
 "media_type": "application/oracle.com.cloud.common.Request+json",
 "total": "1",
 "elements": [
 {
 "uri": "/em/cloud/request/DEC87CE5B1DA175DE040F20A605150F8",
 "name": "SLC01NBV_DBCA_ST - Mon Jun 10 00:20:08 PDT 2013_CREATE_20_25",
 "media_type": "application/oracle.com.cloud.common.Request+json",
 "status": "SUCCESS",
 "lifecycle_type": "CREATE",
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/request/DEC87CE5B
1DA175DE040F20A605150F8"
 }
]
}
}

```

---

**Note:** Even though the Accept type was not specified, the Web service still unambiguously returns the Cloud resource as the "/em/cloud" address uniquely identifies the Cloud as the default resource to be returned.

---

### 27.3.2 Filtering Output for Specific Resource Attributes

The following shows interaction with the Cloud resource picking up specific attributes:

- URL

[https://example.us.oracle.com/em/cloud?service\\_family\\_types,name](https://example.us.oracle.com/em/cloud?service_family_types,name)

- Headers

Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=

- Method

GET

Here are the selective attributes:

```
{
 "name": "Oracle Cloud by Enterprise Manager",
 "service_family_types": {
 "media_type": "application/oracle.com.cloud.common.ServiceFamilyType+json",
 "total": "4",
 "elements": [
 {
 "uri": "/em/cloud/service_family_type/jaas",
 "name": "jaas",
 "media_type":
"application/oracle.com.cloud.common.ServiceFamilyType+json",
 "type": "jaas",
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/service_family_
type/jaas"
 },
 {
 "uri": "/em/cloud/service_family_type/dbaas",
 "name": "dbaas",
 "media_type":
"application/oracle.com.cloud.common.ServiceFamilyType+json",
 "type": "dbaas",
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/service_family_
type/dbaas"
 },
 {
 "uri": "/em/cloud/service_family_type/opc",
 "name": "opc",
 "media_type":
"application/oracle.com.cloud.common.ServiceFamilyType+json",
 "type": "opc",
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/service_family_
type/opc"
 },
 {
 "uri": "/em/cloud/service_family_type/iaas",
 "name": "iaas",
 "media_type":
"application/oracle.com.cloud.iaas.IaasServiceFamilyType+json",
 "type": "iaas",
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/service_family_
type/iaas"
 }
]
 }
}
```

This feature is supported on all the common resources where only selective attributes will be gathered and returned to the client.

### 27.3.3 Service Family Type Resource

The following describes the Service Family Type resource.

- URL  
https://example.us.oracle.com/em/cloud/service\_family\_type/dbaas
- Headers  
Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=
- Method  
GET

The following shows the returned headers and content for this interaction:

```
X-specification-version: 10001
Content-Type: application/oracle.com.cloud.common.ServiceFamilyType+json;
charset=ISO-8859-1
{
 "uri": "/em/cloud/service_family_type/dbaas",
 "name": "dbaas",
 "resource_state": {
 "state": "READY"
 },
 "media_type": "application/oracle.com.cloud.common.ServiceFamilyType+json",
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/service_family_
type/dbaas",
 "type": "dbaas",
 "service_templates": {
 "media_type": "application/oracle.com.cloud.common.ServiceTemplate+json",
 "total": "2",
 "elements": [
 {
 "uri":
"/em/cloud/dbaas/dbplatformtemplate/DEC94FE6D49DAA21E040F20A60511D14",
 "name": "SLC01NBV_SNAPCLONE_ST",
 "description": "",
 "media_type":
"application/oracle.com.cloud.common.DbPlatformTemplate+json",
 "service_family_type": "dbaas",
 "type": "dbaas",
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/dbplatformt
emplate/DEC94FE6D49DAA21E040F20A60511D14"
 },
 {
 "uri":
"/em/cloud/dbaas/dbplatformtemplate/DECCA0E18357F28DE040F20A60514D49",
 "name": "SLC00ECXY_DBCA_ST",
 "description": "",
 "media_type":
"application/oracle.com.cloud.common.DbPlatformTemplate+json",
 "service_family_type": "dbaas",
 "type": "dbaas",
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/dbplatformt
emplate/DECCA0E18357F28DE040F20A60514D49"
 }
]
 }
}
```

```
 },
 "zones": {
 "media_type": "application/oracle.com.cloud.common.Zone+json",
 "total": "2",
 "elements": [
 {
 "uri": "/em/cloud/dbaas/zone/9708491651C090AB4B847C2F3CD177DE",
 "name": "SLC00ECXY_ZONE",
 "description": "",
 "media_type": "application/oracle.com.cloud.common.DbZone+json",
 "service_family_type": "dbaas",
 "type": "self_service_zone",
 "canonicalLink":
 "/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/zone/9708491651C090AB4B847C2F3CD177DE"
 },
 {
 "uri": "/em/cloud/dbaas/zone/B329742A13BE0D63703BF0C0E4F98FB9",
 "name": "SLC01NBV_ZONE",
 "description": "SLC01NBV_ZONE",
 "media_type": "application/oracle.com.cloud.common.DbZone+json",
 "service_family_type": "dbaas",
 "type": "self_service_zone",
 "canonicalLink":
 "/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/zone/B329742A13BE0D63703BF0C0E4F98FB9"
 }
]
 },
 "instance_types": {
 "media_type": "application/oracle.com.cloud.common.InstanceType+json",
 "total": "1",
 "elements": [
 {
 "uri": "/em/cloud/instance_type/dbPlatformInstance%40dbaas",
 "name": "dbPlatformInstance",
 "description": "Db Instance",
 "media_type": "application/oracle.com.cloud.common.InstanceType+json",
 "instance_media_type":
 "application/oracle.com.cloud.common.DbPlatformInstance+json",
 "canonicalLink":
 "/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/instance_type/dbPlatformInstance%40dbaas"
 }
]
 },
 "service_requests": {
 "media_type": "application/oracle.com.cloud.common.Request+json",
 "total": "1",
 "elements": [
 {
 "uri": "/em/cloud/request/DEC87CE5B1DA175DE040F20A605150F8",
 "name": "SLC01NBV_DBCA_ST - Mon Jun 10 00:20:08 PDT 2013_CREATE_20_25",
 "media_type": "application/oracle.com.cloud.common.Request+json",
 "status": "SUCCESS",
 "lifecycle_type": "CREATE",
 "canonicalLink":
 "/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/request/DEC87CE5B1DA175DE040F20A605150F8"
 }
]
 }
 }
```



```
]
 }
}
```

The media type `application/oracle.com.cloud.common.ServiceFamilyType+json` is returned automatically as the most detailed resource.

### 27.3.4 Service Instance Type resource

This is the resource that describes the type of services that are provided by a particular Service Family Type.

- URL  
`https://example.us.oracle.com/em/cloud/instance_type/dbPlatformInstance%40dbaas`
- Headers  
Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=
- Method  
GET

The following shows the Service Instance Type resource:

```
{
 "uri" : "/em/cloud/instance_type/dbPlatformInstance%40dbaas" ,
 "name" : "dbPlatformInstance" ,
 "description" : "Db Instance" ,
 "media_type" : "application/oracle.com.cloud.common.InstanceType+json" ,
 "instance_media_type" :
"application/oracle.com.cloud.common.DbPlatformInstance+json" , ,
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/instance_
type/dbPlatformInstance%40dbaas"
}
```

### 27.3.5 Zone

This describes the Zone resource.

- URL  
`https://example.us.oracle.com/em/cloud/dbaas/zone/82CF1C28FA20A183C99D138FF8065F19`
- Headers  
Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=
- Method  
GET

The following is the Zone resource, automatically cast to the `application/oracle.com.cloud.common.DbZone+json` media type:

```
{
 "uri" : "/em/cloud/dbaas/zone/B49F86387B81DAED9A3D672D722CF29D" ,
 "name" : "RAC SSAZONE" ,
 "description" : "SSA Zone with RAC Nodes" ,
 "resource_state" : {
 "state" : "READY"
 }
}
```

```
 } ,
 "context_id" : "B49F86387B81DAED9A3D672D722CF29D" ,
 "media_type" : "application/oracle.com.cloud.common.DbZone+json" ,
 "service_family_type" : "dbaas" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/zone/B49F86
387B81DAED9A3D672D722CF29D" ,
 "type" : "self_service_zone" ,
 "service_instances" : {
 "media_type" : "application/oracle.com.cloud.common.DbPlatformInstance+json"
 } ,
 "total" : "1" ,
 "elements" :
[
 {
 "uri" : "/em/cloud/dbaas/dbplatforminstance/byrequest/1" ,
 "name" : "db000000.mycompany.com" ,
 "media_type" :
"application/oracle.com.cloud.common.DbPlatformInstance+json" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/dbplatformi
nstance/byrequest/1" ,
 "type" : "rac_database"
 }
]
} ,
"templates" : {
 "media_type" : "application/oracle.com.cloud.common.DbPlatformTemplate+json"
} ,
 "total" : "1" ,
 "elements" :
[
 {
 "uri" :
"/em/cloud/dbaas/dbplatformtemplate/DE039DAC33FFEA6FE040E80A687848A6" ,
 "name" : "12101_RACFileSystem_ServiceTemplate" ,
 "description" : "12101 RAC FileSystem Service Template" ,
 "media_type" :
"application/oracle.com.cloud.common.DbPlatformTemplate+json" ,
 "service_family_type" : "dbaas" ,
 "type" : "dbaas" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/dbplatformt
emplate/DE039DAC33FFEA6FE040E80A687848A6"
 }
]
}
}
```

### 27.3.6 DB Platform Template Resource

As shown in previous examples, the elements in the collection shown all have "media\_type" as an attribute to clearly indicate what its resource type is.

- URL

<https://example.oracle.com/em/cloud/dbaas/dbplatformtemplate/CC3BBB665A6BC6FFE040F00AEF252456>

- Headers

Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=

## ■ Method

### GET

This is the resource returned as the complete DbPlatformTemplate resource:

```
{
 "uri" : "/em/cloud/dbaas/dbplatformtemplate/DE039DAC33FFEA6FE040E80A687848A6" ,
 "name" : "12101_RACFileSystem_ServiceTemplate" ,
 "description" : "12101 RAC FileSystem Service Template" ,
 "resource_state" : {
 "state" : "READY"
 } ,
 "media_type" : "application/oracle.com.cloud.common.DbPlatformTemplate+json" ,
 "service_family_type" : "dbaas" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/dbplatformt
emplate/DE039DAC33FFEA6FE040E80A687848A6" ,
 "type" : "dbaas" ,
 "created" : "2013-05-31T12:35:20+0000" ,
 "zones" : {
 "media_type" : "application/oracle.com.cloud.common.DbZone+json" ,
 "total" : "1" ,
 "elements" :
 [
 {
 "uri" : "/em/cloud/dbaas/zone/B49F86387B81DAED9A3D672D722CF29D" ,
 "name" : "RAC SSAZONE" ,
 "description" : "SSA Zone with RAC Nodes" ,
 "context_id" : "B49F86387B81DAED9A3D672D722CF29D" ,
 "media_type" : "application/oracle.com.cloud.common.DbZone+json" ,
 "service_family_type" : "dbaas" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/zone/B49F86
387B81DAED9A3D672D722CF29D" ,
 "type" : "self_service_zone"
 }
]
 } ,
 "service_instances" : {
 "media_type" : "application/oracle.com.cloud.common.DbPlatformInstance+json"
 } ,
 "total" : "1" ,
 "elements" :
 [
 {
 "uri" : "/em/cloud/dbaas/dbplatforminstance/byrequest/1" ,
 "name" : "db000000.mycompany.com" ,
 "media_type" :
"application/oracle.com.cloud.common.DbPlatformInstance+json" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/dbplatformi
nstance/byrequest/1" ,
 "type" : "rac_database"
 }
]
} ,
"deployment_params" :
[
 {
```

```
 "name" : "username" ,
 "description" : "Master account Username" ,
 "type" : "STRING" ,
 "require" : "true" ,
 "sensitive" : "false"
 } ,
 {
 "name" : "password" ,
 "description" : "Password for the Master account" ,
 "type" : "STRING" ,
 "require" : "true" ,
 "sensitive" : "true"
 }
] ,
"subtype" : "db" ,
"db_type" : "rac_database" ,
"node_count" : "2" ,
"instance_configurable_params" : {
 "username" : "" ,
 "password" : ""
}
}
```

### 27.3.7 Creating Databases

This section describes how to create a service instance using the Cloud resources. As DB is a service instance in the service family type "dbaas", you need to identify the zone of the same family type in which to create the DB. There are two approaches to creating a DB:

1. POST to the Db Zone on which the database is to reside
2. POST to the DbPlatformTemplate of which the database is to be based on

Note the attributes in the submitted body for the POST. The pattern is the same for other service instance creations. Because the intention is clear, the "zone" attribute is not required while POSTing to the Db Zone. Similarly, the "based\_on" attribute is not needed when POSTing to the DbPlatformTemplate. The following operations are described:

- [POST to the Db Zone](#)
- [POST to the DbPlatformTemplate](#)

#### POST to the Db Zone

The following shows the configuration for POSTing to the Db Zone:

- URL  
`https://example.oracle.com/em/cloud/dbaas/zone/82CF1C28FA20A183C99D138FF8065F19`
- Headers  
Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=  
Content-Type: application/oracle.com.cloud.common.DbPlatformInstance+json  
Accept: application/oracle.com.cloud.common.DbPlatformInstance+json
- Body  
{

```

 "based_on":
 "/em/cloud/dbaas/dbplatformtemplate/CC3BBB665A6BC6FFE040F00AEF252456",
 "name": "<Request name>",
 "description": "<Request Description>",
 "params":
 {
 "username": "Master Account username for the DB",
 "password": "Password for the Master Account on the DB"
 }
 }
 }
 }

```

- Method

POST

The following is the response from the above interaction:

```

{
 "uri" : "/em/cloud/dbaas/dbplatforminstance/byrequest/22" ,
 "name" : "Test_create_intance" ,
 "resource_state" : {
 "state" : "INITIATED"
 } ,
 "media_type" : "application/oracle.com.cloud.common.DbPlatformInstance+json" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/dbplatformi
nstance/byrequest/22" ,
 "status" : "SCHEDULED"
}

```

### POST to the DbPlatformTemplate

The following describes the configuration for POSTing to the DbPlatformTemplate:

- URL

<https://example.oracle.com/em/cloud/dbaas/dbplatformtemplate/CC3BBB665A6BC6FFE040F00AEF252456>

- Headers

Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=

Content-Type: application/oracle.com.cloud.common.DbPlatformInstance+json

Accept: application/oracle.com.cloud.common.DbPlatformInstance+json

- Body

```

{
 "zone": "/em/cloud/dbaas/zone/82CF1C28FA20A183C99D138FF8065F19",
 "name": "<Request name>",
 "description": "<Request Description>",
 "params":
 {
 "username": "Master Account username for the DB",
 "password": "Password for the Master Account on the DB"
 }
}

```

- Method

POST

The following is the response from the above interaction:

```
{
 "uri" : "/em/cloud/dbaas/dbplatforminstance/byrequest/22" ,
 "name" : "Test_create_intance" ,
 "resource_state" : {
 "state" : "INITIATED"
 } ,
 "media_type" : "application/oracle.com.cloud.common.DbPlatformInstance+json" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/dbplatformi
nstance/byrequest/23" ,
 "status" : "SCHEDULED"
}
```

This response is the same as the previous version. If the zone is not specified, then it is not possible for the Web service to determine where the DB should be deployed.

### 27.3.8 Polling the Database Creation

Once the database creation has been POSTed, you can GET the resource identified by the return URI to keep track of the status.

- URL

`https://example.oracle.com/em/cloud/dbaas/dbplatforminstance/byrequest/7`

- Headers

Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=

Accept: application/oracle.com.cloud.common.DbPlatformInstance+json

- Method

GET

The following is the response:

```
{
 "uri": "/em/cloud/dbaas/dbplatforminstance/byrequest/22",
 "name": "Test_create_intance_CREATE_55_42",
 "description": "Create Database for test.",
 "resource_state":
 {
 "state": "CREATING"
 },
 "context_id": "22",
 "media_type": "application/oracle.com.cloud.common.DbPlatformInstance+json",
 "status": "IN_PROGRESS",
 "created": "2013-06-07 23:09:04" ,
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/dbplatformi
nstance/byrequest/22"
}
```

Note that the "resource\_state" attribute is showing "CREATING" state to indicate that the database is still being created. Once the database is created, the same GET will return the database resource with "READY" state, as shown below:

```
{
 "uri" : "/em/cloud/dbaas/dbplatforminstance/byrequest/22" ,
 "name" : "db000000.mycompany.com" ,
 "resource_state" : {
 "state" : "READY"
 }
}
```

```

 },
 "context_id" : "22" ,
 "media_type" : "application/oracle.com.cloud.common.DbPlatformInstance+json" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/dbplatformi
nstance/byrequest/22" ,
 "status" : "RUNNING" ,
 "created" : "2013-06-07 23:14:04" ,
 "based_on" :
"/em/cloud/dbaas/dbplatformtemplate/DE9845B82981CF0CE040E80AD87838C9" ,
 "connect_string" : "(DESCRIPTION=(ADDRESS_
LIST=(ADDRESS=(PROTOCOL=TCP) (HOST=cluster-r) (PORT=1521))) (CONNECT_DATA=(SERVICE_
NAME=db000000.mycompany.com)))" ,
 "type" : "rac_database" ,
 "destination_zone" : "/em/cloud/dbaas/zone/B49F86387B81DAED9A3D672D722CF29D" ,
 "up_time" : "0" ,
 "load" : "0.1801962347135334" ,
 "total_sessions" : "155" ,
 "total_sga" : "4778.5546875" ,
 "available_space" : "0.68255615234375" ,
 "instances" :
[
 {
 "name" : "db000000.mycompany.com_db0000001" ,
 "total_memory" : "2672.286 MB" ,
 "active_sessions" : ".037" ,
 "cpu_utilization" : "1.106" ,
 "db_time" : "3.66" ,
 "cpu_count" : "2" ,
 "allocated_sessions" : "63"
 } ,
 {
 "name" : "db000000.mycompany.com_db0000002" ,
 "total_memory" : "2676.052 MB" ,
 "active_sessions" : ".144" ,
 "cpu_utilization" : "4.225" ,
 "db_time" : "14.359" ,
 "cpu_count" : "2" ,
 "allocated_sessions" : "65"
 }
] ,
 "last_backup" : "null" ,
 "master_username" : "InstanceUser1" ,
 "db_version" : "11.2.0.3.0"
}

```

The output in case of single instance database (type:oracle\_database) :

```

{
 "uri" : "/em/cloud/dbaas/dbplatforminstance/byrequest/23" ,
 "name" : "sidb0001.host1.mycompany.com" ,
 "resource_state" : {
 "state" : "READY"
 } ,
 "context_id" : "23" ,
 "media_type" : "application/oracle.com.cloud.common.DbPlatformInstance+json" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/dbplatformi
nstance/byrequest/23" ,
 "status" : "RUNNING" ,
 "created" : "2013-06-09 07:56:57" ,
 "based_on" :

```

```
"/em/cloud/dbaas/dbplatformtemplate/DEB2BDE2F773E356E040E50AD2E56BE9" ,
 "connect_string" : " (DESCRIPTION= (ADDRESS_
LIST= (ADDRESS= (PROTOCOL=TCP) (HOST=host1.mycompany.com) (PORT=1521))) (CONNECT_
DATA= (SID=sidb0001))) " ,
 "type" : "oracle_database" ,
 "destination_zone" : "/em/cloud/dbaas/zone/DEB2E2920C20D2E1E040E50AD2E56BEF" ,
 "up_time" : "0" ,
 "load" : "0.156456056384199" ,
 "total_sessions" : "46" ,
 "total_sga" : "398.203125" ,
 "available_space" : "0.19097900390625" ,
 "last_backup" : "null" ,
 "master_username" : "InstanceUser1" ,
 "db_version" : "11.2.0.1.0" ,
 "active_sessions" : ".156" ,
 "cpu_utilization" : "5.35" ,
 "db_time" : "15.646" ,
 "cpu_count" : "1" ,
 "allocated_sessions" : "39" ,
 "total_memory" : "524.384 MB"
}
```

### 27.3.9 Deleting a Database Instance

The DELETE request on the URI of the resource can be issued to delete a DB Service Instance. For example, the following DELETE request is issued:

- URL  
`https://example.oracle.com/em/em/cloud/dbaas/dbplatforminstance/byrequest/7`
- Headers  
Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=  
Accept: application/oracle.com.cloud.common.DbPlatformInstance+json
- Method  
DELETE

The following Web service response with code 200 is returned:

```
{
 "uri" : "/em/cloud/dbaas/dbplatforminstance/byrequest/22" ,
 "name" : "db000000.mycompany.com" ,
 "resource_state" : {
 "state" : "DESTROYING" ,
 "messages" :
 [
 {
 "text" : "Submit 'DELETE' operation on
'application/oracle.com.cloud.common.DbPlatformInstance+json'
'db000000.mycompany.com'. The process has job id '41' " ,
 "date" : "2013-06-07T23:16:20+0000"
 }
]
 } ,
 "context_id" : "22" ,
 "media_type" : "application/oracle.com.cloud.common.DbPlatformInstance+json" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/dbplatformi
```



```
nstance/byrequest/22" ,
 "type" : "rac_database"
}
```

After some time, the GET request on the database URI will result in the following interaction:

- URL  
https://example.oracle.com/em/cloud/dbaas/dbplatforminstance/byrequest/7
- Headers  
Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=
- Method  
GET

The Web service returns the following:

```
{
 "uri" : "/em/cloud/dbaas/dbplatforminstance/byrequest/23" ,
 "name" : "sidb0001.host1.mycompany.com" ,
 "resource_state" : {
 "state" : "DESTROYING"
 } ,
 "context_id" : "23" ,
 "media_type" : "application/oracle.com.cloud.common.DbPlatformInstance+json" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/dbplatformi
nstance/byrequest/23" ,
 "status" : "RUNNING" ,
 "created" : "2013-06-09 07:56:57" ,
 "based_on" :
"/em/cloud/dbaas/dbplatformtemplate/DEB2BDE2F773E356E040E50AD2E56BE9" ,
 "connect_string" : "(DESCRIPTION=(ADDRESS_
LIST=(ADDRESS=(PROTOCOL=TCP) (HOST=host1.mycompany.com) (PORT=1521))) (CONNECT_
DATA=(SID=sidb0001)))" ,
 "type" : "oracle_database" ,
 "destination_zone" : "/em/cloud/dbaas/zone/DEB2E2920C20D2E1E040E50AD2E56BEF" ,
 "up_time" : "0" ,
 "load" : "0.156456056384199" ,
 "total_sessions" : "46" ,
 "total_sga" : "398.203125" ,
 "available_space" : "0.19097900390625" ,
 "last_backup" : "null" ,
 "master_username" : "InstanceUser1" ,
 "db_version" : "11.2.0.1.0" ,
 "active_sessions" : ".156" ,
 "cpu_utilization" : "5.35" ,
 "db_time" : "15.646" ,
 "cpu_count" : "1" ,
 "allocated_sessions" : "39" ,
 "total_memory" : "524.384 MB"
}
```

The "status" of the database will be "DELETED" once the database is DELETED.

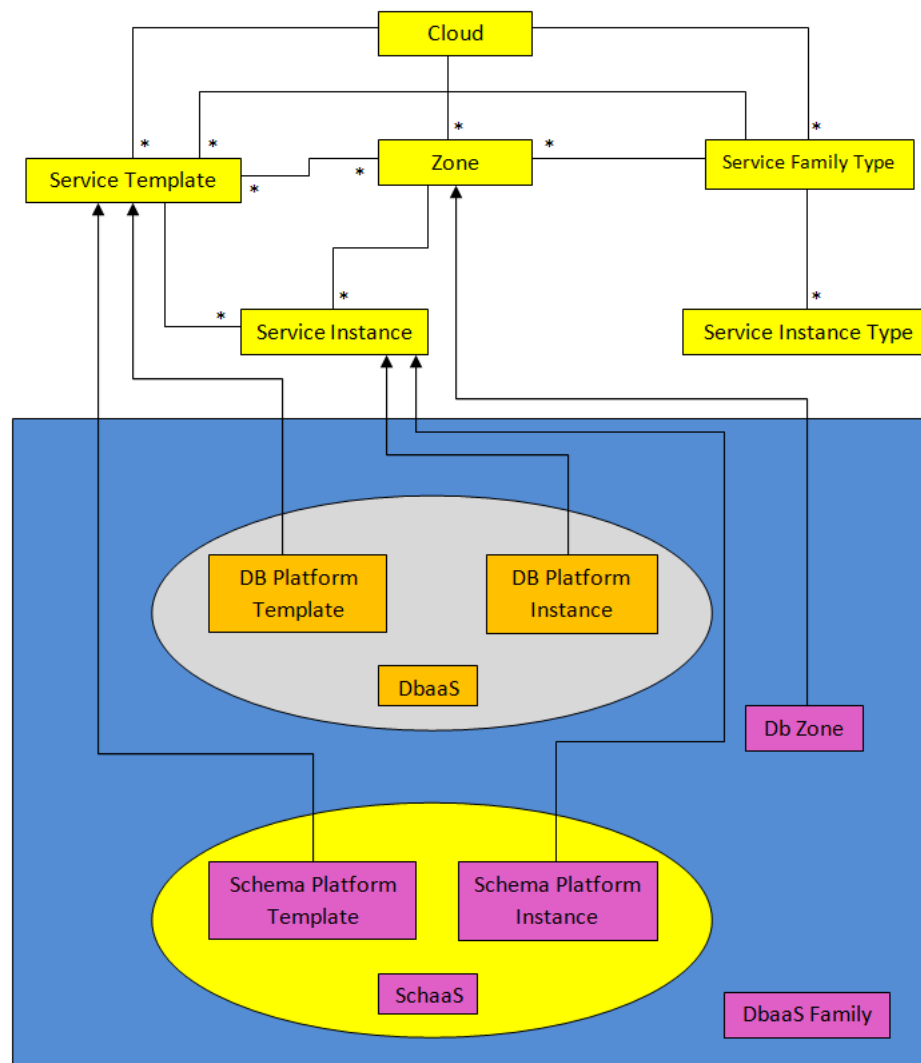
## 27.4 Resource Model for Schema as a Service

This section describes the resource model for Schema as a Service (SchaaS) API and the REST request/response interactions. The following topics are covered:

### 27.4.1 Resource Model for Schema as a Service

Cloud resources are organized by common service entities, and specific services (for example, Schema as a Service) provide resources by extending these common service entities. The figure below shows the SchaaS resource model.

**Figure 27–2 Schema as a Service Resource Model**



The following sections describe the resources that are supported by the Schema as a Service APIs and provide examples of supported operations. The following resource models are described:

- [Schema PlatformTemplate](#)
- [SchemaPlatformInstance](#)

Schema as a Service uses the same Cloud, Service Family, and DB Zone resources as described in [Section 27.1, "Resource Model for Database as a Service"](#).

## 27.4.2 Schema PlatformTemplate

The SchemaPlatformTemplate extends the ServiceTemplate resource. It represents the service template created for Schema Service creation. It can be a template for creating a schema service using profile based schemas (schema with data) or for creating a schemas service containing empty schemas.

It is permissible to accept the ServiceTemplate media type on an instance of SchemaPlatformTemplate. It is a preconfigured deployable service that realizes a SchemaPlatformInstance resource. SchemaPlatformTemplate has the following media type and payload:

- **Media Type:** application/oracle.com.cloud.common.SchemaPlatformTemplate
- **Supported Payload:** json
- **URI Format:** /em/cloud/dbaas/schemaplatformtemplate/<template id>

The following table describes the SchemaPlatformTemplate Data Model.

**Table 27–5 SchemaPlatformTemplate Data Model**

| Field             | Type                                | Occurs | Description                                                                                                                                                         |
|-------------------|-------------------------------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| uri               | URI                                 | 1      | A GET against this URI refreshes the client representation of the resources accessible to this user.                                                                |
| name              | String                              | 1      | A human readable name given to the template.                                                                                                                        |
| description       | String                              | 0..1   | A brief description given to the template.                                                                                                                          |
| created           | Timestamp                           | 1      | Date and time, in ISO 8601 format, when the template was created/last modified.                                                                                     |
| resource_state    | ResourceState                       | 1      | The validity of the fields on a GET should be guaranteed only when the resource state is READY. Otherwise, the client should not assume the validity of the fields. |
| zones             | Collection <DbZone>                 | 0..1   | List of all zones configured for this template.                                                                                                                     |
| service_instances | Collection <SchemaPlatformInstance> | 0..1   | All service instances created based on this SchemaPlatformTemplate.                                                                                                 |
| media_type        | String                              | 1      | The media type of the service template that this type represents.                                                                                                   |
| canonicalLink     | URI                                 | 1      | Can be used to perform cloud interactions like GET and POST.                                                                                                        |
| params            | Object                              | 1      | List of parameters required to provision a Schema Platform Instance from this template.                                                                             |
| deployment_params | List<Object>                        | 1      | Description of the parameters which are required while creating a service instance.                                                                                 |
| type              | String                              | 1      | The type of the service template which is "dbaas" by default.                                                                                                       |

**Table 27–5 (Cont.) SchemaPlatformTemplate Data Model**

| Field                 | Type         | Occurs | Description                                                                                                                             |
|-----------------------|--------------|--------|-----------------------------------------------------------------------------------------------------------------------------------------|
| subtype               | String       | 1      | The sub type of the instance created based on this template.<br>It has a value "schema" in this case.                                   |
| template_type         | String       | 1      | Type of this template based on how the template got created. It can have any of two values : "EMPTY_SCHEMAS" or "SCHEMAS_FROM_PROFILE". |
| max_number_of_schemas | Integer      | 0..1   | Maximum number of schemas that can be created in a single instance request. This field exists only if template_type is "EMPTY_SCHEMAS". |
| workloads             | List<Object> | 0..1   | List all the workloads defined for this template and list the attributes for each one of them.                                          |

### 27.4.3 SchemaPlatformInstance

The SchemaPlatformInstance extends the ServiceInstance resource. It is linked to Oracle Database Service target in Enterprise Manager Cloud Control. SchemaPlatformInstance has the following media type and payload:

- **Media Type:** application/oracle.com.cloud.common.SchemaPlatformInstance
- **Supported Payload:** json
- **URI Format:** /em/cloud/dbaas/schemaplatforminstance/byrequest/<request id>

The following table describes the SchemaPlatformInstance Data Model:

**Table 27–6 SchemaPlatformInstance Data Model**

| Field            | Type                   | Occurs | Description                                                                                                                                                         |
|------------------|------------------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| uri              | URI                    | 1      | A GET against this URI refreshes the client representation of the resources accessible to this user.                                                                |
| name             | String                 | 1      | A human readable name given given to the database service<br>[POST]                                                                                                 |
| destination_zone | Zone                   | 1      | The Zone to which the database service is deployed<br>[POST]                                                                                                        |
| created          | Timestamp              | 1      | Date and time, in ISO 8601 format, when the template was created/last modified.                                                                                     |
| resource_state   | ResourceState          | 1      | The validity of the fields on a GET should be guaranteed only when the resource state is READY. Otherwise, the client should not assume the validity of the fields. |
| based_on         | SchemaPlatformTemplate | 1      | The schema platform template on which this database service is published.                                                                                           |
| media_type       | String                 | 1      | Indicates the additional media type that clients can use to perform a GET.                                                                                          |

**Table 27–6 (Cont.) SchemaPlatformInstance Data Model**

| Field          | Type   | Occurs | Description                                                                 |
|----------------|--------|--------|-----------------------------------------------------------------------------|
| canonicalLink  | URI    | 1      | Can be used to perform cloud interactions like GET and POST.                |
| type           | String | 1      | Type of the database service. It has the default value "oracle_dbsvc".      |
| connect_string | String | 0..1   | Connect String required to establish a connection to this database service. |
| status         | String | 1      | Status of the database service.                                             |
| db_version     | String | 0..1   | Version of the database on which this service is running.                   |
| context_id     | String | 1      | The request id of this database service.                                    |

## 27.5 Supported Operations for Schema as a Service Resources

The following table lists all the operations that are supported on the various Schema as a Service resources.

**Table 27–7 Supported Operations**

| Resource                                    | Operations (GET, POST, PUT, and DELETE) |
|---------------------------------------------|-----------------------------------------|
| Cloud                                       | GET                                     |
| Service Family Type                         | GET                                     |
| DB Zone                                     | GET, POST                               |
| Schema Platform Template [Service Template] | GET, POST                               |
| Schema Platform Instance [Service Instance] | GET, DELETE                             |

## 27.6 Schema as a Service API Examples

The following sections provide examples of interaction with Enterprise Manager Cloud Control 12.1.0.3 with Cloud Application plug-in 12.1.0.6 and higher. The examples for GET on cloud, GET on dbaas family type and GET on Db Zone are covered in the DbaaS APIs section. The following examples are covered:

- [Schema Platform Template Resource](#)
- [Creating Database Services](#)

### 27.6.1 Schema Platform Template Resource

As shown in previous examples (in DbaaS), the elements in the collection showed, all have "media\_type" as an attribute to clearly indicate what its resource type is.

- URL  
https://example.mycompany.com/em/cloud/dbaas/schemaplatformtemplate/C3BBB665A6BC6FFE040F00AEF252456
- Headers  
Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=
- Method  
GET

This is the resource returned as the complete SchemaPlatformTemplate resource:

```
{
 "uri":
"/em/cloud/dbaas/schemaplatformtemplate/DE3F8D882F132F45E040F00AAA2330C5",
 "name": "Template Schema Si",
 "description": "",
 "resource_state": {
 "state": "READY"
 },
 "media_type": "application/oracle.com.cloud.common.SchemaPlatformTemplate+json",
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/schemaplatf
ormtemplate/DE3F8D882F132F45E040F00AAA2330C5",
 "type": "dbaas",
 "created": "2013-06-03T11:51:20+0000",
 "zones": {
 "media_type": "application/oracle.com.cloud.common.DbZone+json",
 "total": "1",
 "elements": [
 {
 "uri": "/em/cloud/dbaas/zone/34C1FC23A7503E1CC9C490C9670635F9",
 "name": "Zone Si",
 "description": "",
 "context_id": "34C1FC23A7503E1CC9C490C9670635F9",
 "media_type": "application/oracle.com.cloud.common.DbZone+json",
 "service_family_type": "dbaas",
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/zone/34C1FC
23A7503E1CC9C490C9670635F9",
 "type": "self_service_zone"
 }
]
 },
 "service_instances": {
 "media_type":
"application/oracle.com.cloud.common.SchemaPlatformInstance+json",
 "total": "1",
 "elements": [
 {
 "uri": "/em/cloud/dbaas/schemaplatforminstance/byrequest/61",
 "name": "Service_1_2AEE494EE0994",
 "media_type":
"application/oracle.com.cloud.common.SchemaPlatformInstance+json",
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/schemaplatf
orminstance/byrequest/61",
 "target_type": "oracle_dbsvc"
 }
]
 },
 "deployment_params": [
 {
 "name": "masterAccount",
 "description": "Master account schema. Not Required for the SCHEMAS_FROM_
PROFILE Schema templates",
 "type": "STRING",
 "require": "false",
 "sensitive": "false"
 }
],
 {
```

```

 "name": "service_name_prefix",
 "description": "Prefix for the service name instances",
 "type": "STRING",
 "require": "true",
 "sensitive": "false"
 },
 {
 "name": "workload_name",
 "description": "WorkLoad name to be associated with the request",
 "type": "STRING",
 "require": "true",
 "sensitive": "false"
 },
 {
 "name": "schema",
 "description": "List of schema to be created",
 "type": "LIST",
 "defaultValue": [
 {
 "name": "username",
 "description": "New schema name",
 "type": "STRING",
 "require": "true",
 "sensitive": "false"
 },
 {
 "name": "original_name",
 "description": "Original schema name. Required only for the SCHEMAS_
FROM_PROFILE Schema templates",
 "type": "STRING",
 "require": "false",
 "sensitive": "false"
 },
 {
 "name": "password",
 "description": "Password for the schema",
 "type": "STRING",
 "require": "true",
 "sensitive": "true"
 }
],
 "require": "true",
 "sensitive": "false"
 }
],
"subtype": "schema",
"template_type": "EMPTY_SCHEMAS",
"max_number_of_schemas": "4",
"workloads": [
 {
 "name": "WL_112",
 "description": "null",
 "attributes": [
 {
 "name": "CPU",
 "value": "0.1",
 "units": "cores"
 },
 {
 "name": "Memory",

```

```
 "value": "0.1",
 "units": "GB"
 },
 {
 "name": "Storage",
 "value": "10",
 "units": "GB"
 }
]
 },
 "params": {
 "masterAccount": "",
 "service_name_prefix": "",
 "workload_name": "",
 "schema": [
 {
 "username": "",
 "password": ""
 },
 {
 "username": "",
 "password": ""
 }
]
 }
}
```

## 27.6.2 Creating Database Services

This section describes how to create a database service using the Cloud resources. As Database service is also a service instance in the service family type "dbaas", you need to identify the zone of the same family type in which to create the database service. There are two approaches to creating a database service:

- POST to the Db Zone on which db service is to be created
- POST to the SchemaPlatformTemplate of which the database service is to be based on

Note the attributes in the submitted body for the POST. The pattern is the same for other service instance creations. Because the intention is clear, the "zone" attribute is not required while POSTing to the Db Zone. Similarly, the "based\_on" attribute is not needed when POSTing to the SchemaPlatformTemplate. The following operations are described:

- [POST to the Db Zone](#)
- [POST on SchemaPlatformTemplate](#)
- [Polling the Schema Service Creation](#)
- [Deleting a Schema Service](#)

### POST to the Db Zone

The following shows the configuration for creating a Schema Service Instance

- URL  
`https://example.mycompany.com/em/cloud/dbaas/zone/82CF1C28FA20A183C99D138FF8065F19`



- Headers

Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=

Content-Type:

application/oracle.com.cloud.common.SchemaPlatformInstance+json

Accept: application/oracle.com.cloud.common.SchemaPlatformInstance+json

- Body for a Profile Based Schema Template

```
{
 "based_on":
"/em/cloud/dbaas/schemaplatfromtemplate/DE3F8D882F132F45E040F00AAA2330C5",
 "name": "TestRequest",
 "params": {
 "workload_name": "WL_112",
 "service_name_prefix": "Service_1",
 "masterAccount": " SchemaUser1",
 "schema": [
 {
 "username": " SchemaUser1",
 "password": "welcome1"
 },
 {
 "username": " SchemaUser2",
 "password": "welcome1"
 }
]
 }
}
```

- Method

POST

The following is the response from the above interaction:

```
{
 "uri" : "/em/cloud/dbaas/schemaplatfrominstance/byrequest/22" ,
 "name" : "TestRequest" ,
 "resource_state" : {
 "state" : "INITIATED"
 } ,
 "media_type" : "application/oracle.com.cloud.common.SchemaPlatformInstance+json"
,
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/schemaplatf
orminstance/byrequest/22" ,
 "status" : "SCHEDULED"
}
```

### POST on SchemaPlatformTemplate

The following shows the configuration for creating a Schema Service Instance

- URL

<https://example.mycompany.com/em/cloud/dbaas/schemaplatfromtemplate/C3BBB665A6BC6FFE040F00AEF252456>

- Headers

Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=

Content-Type:

application/oracle.com.cloud.common.SchemaPlatformInstance+json

Accept: application/oracle.com.cloud.common.SchemaPlatformInstance+json

- Body for a Profile Based Schema Template

```
{
 "zone": "/em/cloud/dbaas/zone/34C1FC23A7503E1CC9C490C9670635F9",
 "name": "TestRequest",
 "params": {
 "workload_name": "WL_112",
 "service_name_prefix": "Service_1",
 "masterAccount": "SchemaUser1",
 "schema": [
 {
 "username": "SchemaUser1",
 "password": "welcome"
 },
 {
 "username": "SchemaUser2",
 "password": "welcome"
 }
]
 }
}
```

- Method

POST

The following is the response from the above interaction:

```
{
 "uri" : "/em/cloud/dbaas/schemaplatforminstance/byrequest/23" ,
 "name" : "TestRequest" ,
 "resource_state" : {
 "state" : "INITIATED"
 } ,
 "media_type" : "application/oracle.com.cloud.common.SchemaPlatformInstance+json"
,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/schemaplatf
orminstance/byrequest/23" ,
 "status" : "SCHEDULED"
}
```

This response is the same as the previous version. If the zone is not specified, then it is not possible for the Web service to determine where the database service should be deployed.

### Polling the Schema Service Creation

Once the database service creation has been POSTed, you can GET the resource identified by the return URI to keep track of the status.

- URL

<https://example.mycompany.com/em/cloud/dbaas/schemaplatforminstance/byrequest/22>

- Headers

Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=

Accept: application/oracle.com.cloud.common.SchemaPlatformInstance+json

- Method

GET

The following is the response:

```
{
 "uri": "/em/cloud/dbaas/schemaplatforminstance/byrequest/22",
 "name": "TestRequest_CREATE_55_42",
 "description": "Create Service for test.",
 "resource_state": {
 "state": "CREATING"
 },
 "context_id": "22",
 "media_type":
"application/oracle.com.cloud.common.SchemaPlatformInstance+json",
 "status": "IN_PROGRESS",
 "created": "2013-06-07 23:09:04" ,
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/schemaplatf
orminstance/byrequest/22"
}
```

Note that the "resource\_state" attribute is showing "CREATING" state to indicate the database Service is still being created. Once the database service is created, the same GET will return the database service resource with "READY" state, as shown below:

```
{
 "uri": "/em/cloud/dbaas/schemaplatforminstance/byrequest/22",
 "name": "Service_1_35E576F22E97E",
 "resource_state": {
 "state": "READY"
 },
 "context_id": "2",
 "media_type": "application/oracle.com.cloud.common.SchemaPlatformInstance+json",
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/schemaplatf
orminstance/byrequest/2",
 "status": "RUNNING",
 "created": "2013-06-10 02:45:32",
 "based_on":
"/em/cloud/dbaas/schemaplatformtemplate/DECA1100F797EF10E040F00AFBB123E7",
 "type": "oracle_dbsvc",
 "connect_string": "(DESCRIPTION=(ADDRESS_
LIST=(ADDRESS=(PROTOCOL=TCP) (HOST=host1.mycompany.com) (PORT=1521))) (CONNECT_
DATA=(SERVICE_NAME=Service_1_35E576F22E97E.mycompany.com)))",
 "destination_zone": "/em/cloud/dbaas/zone/34C1FC23A7503E1CC9C490C9670635F9",
 "db_version": "11.2.0.3.0"
}
```

### Deleting a Schema Service

The DELETE request on the URI of the resource can be issued to delete a database service. For example, the following DELETE request is issued:

- URL

<https://example.mycompany.com/em/cloud/dbaas/schemaplatforminstance/byrequest/22>

- Headers

Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=

Accept: application/oracle.com.cloud.common.SchemaPlatformInstance+json

- Method

DELETE

The following Web service response with code 200 is returned:

```
{
 "uri" : "/em/cloud/dbaas/schemaplatforminstance/byrequest/22" ,
 "name" : "Service_1_35E576F22E97E" ,
 "resource_state" : {
 "state" : "DESTROYING" ,
 "messages" :
 [
 {
 "text" : "Submit 'DELETE' operation on
'application/oracle.com.cloud.common.SchemaPlatformInstance+json' 'Service_1_
35E576F22E97E'. The process has job id '41' " ,
 "date" : "2013-06-07T23:16:20+0000"
 }
]
 } ,
 "context_id" : "22" ,
 "media_type" : "application/oracle.com.cloud.common.SchemaPlatformInstance
+json" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/schemaplatf
orminstance/byrequest/22" ,
 "type" : "oracle_dbsvc"
}
```

After some time, the GET request on the database service URI will result in the following interaction:

- URL

<https://example.mycompany.com/em/cloud/dbaas/schemaplatforminstance/byrebyrequest/22>

- Headers

Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=

- Method

GET

The Web service returns the following:

```
{
 "uri": "/em/cloud/dbaas/schemaplatforminstance/byrequest/22",
 "name": "Service_1_35E576F22E97E",
 "resource_state": {
 "state": "DESTROYING"
 },
 "context_id": "2",
 "media_type": "application/oracle.com.cloud.common.SchemaPlatformInstance+json",
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/schemaplatf
orminstance/byrequest/2",
}
```

```

 "status": "RUNNING",
 "created": "2013-06-10 02:45:32",
 "based_on":
"/em/cloud/dbaas/schemaplatformtemplate/DECA1100F797EF10E040F00AFBB123E7",
 "type": "oracle_dbsvc",
 "connect_string": " (DESCRIPTION= (ADDRESS_
LIST= (ADDRESS= (PROTOCOL=TCP) (HOST=host1.mycompany.com) (PORT=1521))) (CONNECT_
DATA= (SERVICE_NAME=Service_1_35E576F22E97E.mycompany.com))) ",
 "destination_zone": "/em/cloud/dbaas/zone/34C1FC23A7503E1CC9C490C9670635F9",
 "db_version": "11.2.0.3.0"
}

```

Please note the status of the database service will be DELETED once the schema service is deleted.

## 27.7 Using EMCLI to Create Database Profiles

This section explains the different Database Provisioning Profiles used in Database as a Service and Schema as a Service. It also explains how these profiles can be created using the EMCLI verbs exposed in the 12.1.0.3 database plug-in. The following topics are covered:

- [About Database Provisioning Profile Types](#)
- [Creating Provisioning Profiles using EMCLI](#)
- [Sample EMCLI Property Files](#)

### 27.7.1 About Database Provisioning Profile Types

The following are the four types of Database provisioning profiles used in Database as a Service and Schema as a Service:

- [DBCA Template](#)
- [RMAN](#)
- [Snapclone](#)
- [Schema](#)

#### DBCA Template

The DBCA template supports both structure plus data and structure only types. This database provisioning profile is created from the production or source database registered with Enterprise Manager.

---

**Note:** Structure and data requires a shutdown of the database. It should be used for very small databases (< 50GB).

---

#### RMAN

This data content of this database provisioning profile comprises the RMAN Backup pieces. Any of the following methods can be used to create the RMAN profile:

- Hot RMAN Backup
 

This option is enabled only when the archive log mode on the source database is turned ON.
- Cold RMAN Backup

In this case, the source database is brought down to take the RMAN backup

- Existing Backup

This option can be used when the RMAN backup is already taken on the source and you want to use one of the existing RMAN backups to create a profile.

- Existing Backup pieces when there is no access to the source database

Create profile source database using existing backup pieces when there is no access to the source database. This option can be used when the RMAN backup is already taken on the source and kept in a known location. This option is only supported using EMCLI.

### Snapclone

The data content of this database provisioning profile is derived from Snapshot information on the Filer (NetApp or SunZFS filer) of all the volumes where the datafiles for the source database are kept.

### Schema

The data content of this database provisioning profile is compiled from an export dump created using the Data Pump tool. You can choose to export Schema objects (structure only) or the Schema with data (structure + data).

Note that you cannot export empty schemas or include Oracle schemas (these schemas are not available for selection). Schemas that are filtered out are listed below:

**Table 27–8 Filtered Out Schema**

|                             |             |                  |                       |          |
|-----------------------------|-------------|------------------|-----------------------|----------|
| ANONYMOUS                   | DMSYS       | OASPUBLIC        | OWBSYS                | TSMSYS   |
| APEX_030200                 | EXFSYS      | ODM              | OWBSYS_AUDIT          | WEBSYS   |
| APEX_PUBLIC_USER            | FLows_      | ODM_MTR          | SCOTT                 | WK_PROXY |
| APPQOSSYS                   | FLows_03000 | OLAPSYS          | SI_INFORMTN_SCHEMA    | WK_TEST  |
| AURORA\$JIS\$UTILITY\$      | FLows_FILES | ORACLE_OCM       | SPATIAL_CSW_ADMIN_USR | WKPROXY  |
| AURORA\$ORB\$UNAUTENTICATED | LBACSYS     | ORDDATA          | SPATIAL_WFS_ADMIN_USR | WKSYS    |
| BI                          | MDDATA      | ORDPLUGINS       | SYS                   | WMSYS    |
| CTXSYS                      | MDSYS       | ORDSYS           | SYSMAN                | XDB      |
| DBSNMP                      | MGMT_VIEW   | OSE\$HTTP\$ADMIN | SYSTEM                | XS\$NULL |
| DIP                         | MTSSYS      | OUTLN            | TRACESRV              |          |

## 27.7.2 Creating Provisioning Profiles using EMCLI

This method allows the administrators or cloud provisioning operators to create the database provisioning profile using EMCLI verbs. The following is the syntax of the verb introduced in the 12.1.0.3 database plug-in:

```
emcli create_dbprofile -input_file=data:<Prop file name>
```

This command takes in a property file that completely describes the type of profile that will be created and the options used.

### 27.7.3 Sample EMCLI Property Files

The following sections provide sample EMCLI property files for the database provisioning profiles:

- [DBCA Template](#)
- [New RMAN Backup from the Source Database](#)
- [Existing RMAN Backup from the Source Database](#)
- [Existing RMAN Backup on the Stage Location](#)
- [Snapclone Database](#)
- [Schema Database](#)

#### DBCA Template

**Table 27–9 EMCLI Property Files for DBCA Template**

| Name                                            | Required | Description                                                                                                                                                                   | Variable                       |
|-------------------------------------------------|----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------|
| Reference Database                              | true     | The source database from which the profile will be created                                                                                                                    | REFERENCE_DATABASE=            |
| Type of the reference database                  | true     | The target type of the source database (oracle_database or rac_database)                                                                                                      | REFERENCE_DATABASE_TYPE=       |
| Reference Node                                  | false    | The reference node from which the profile will be created. Applicable only for RAC and Cluster targets                                                                        | REF_NODE_SELECTED=             |
| Backup mechanism for data                       | false    | Variable to identify what mode needs to be used for creating backup data. Possible values are EXPORT, DBCA_TEMPLATE, RMAN, THIN_PROVISIONING                                  | DATA_CONTENT_MODE=DBCATEMPLATE |
| The content that will be included in the backup | false    | The data content that gets included in the backup can metadata or data. Possible values are METADATA_ONLY and METADATA_AND_DATA.                                              | DATA_CONTENT=METADATA_AND_DATA |
| Reference Host Credentials                      | true     | The credentials that will be used for OS authentication. Format is <Cred Name:Owner>. If this is not provided, preferred credentials will be checked against NormalHostCreds. | REF_HOST_CREDENTIALS=          |

**Table 27–9 (Cont.) EMCLI Property Files for DBCA Template**

| Name                                 | Required | Description                                                                                                                                                                               | Variable             |
|--------------------------------------|----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------|
| Database Credentials                 | false    | The credentials that will be used for performing database authentication. Format is <Cred Name:Owner>. If this is not provided, preferred credentials will be checked against SysDbuCreds | REF_DB_CREDENTIALS=  |
| Convert to OFA format                | false    | Converts to the OFA format                                                                                                                                                                | CONVERT_TO_OFA=      |
| Profile Name                         | true     | Name of the profile                                                                                                                                                                       | PROFILE_NAME=        |
| Profile Description                  | false    | Name with which the profile should be created.                                                                                                                                            | PROFILE_DESC=        |
| Profile location                     | true     | Location of the profile in software library                                                                                                                                               | PROFILE_LOCATION=    |
| Profile Version                      | false    | Version of the profile. If none is specified, the Database version shall be used.                                                                                                         | PROFILE_VERSION=     |
| Profile Vendor                       | false    | Vendor of the profile. Default is Oracle.                                                                                                                                                 | PROFILE_VENDOR=      |
| Profile notes                        | false    | Notes for the profile.                                                                                                                                                                    | PROFILE_NOTES=       |
| Working Directory                    | false    | Working directory is the location where the files are staged.                                                                                                                             | WORKING_DIRECTORY=   |
| Is the profile initiated by end user | true     | Is the profile initiated by end user as part of backup process or the database.                                                                                                           | USER_INITIATED=false |

### New RMAN Backup from the Source Database

**Table 27–10 EMCLI Property Files for a New RMAN Backup**

| Name                           | Required | Description                                                                                                                                      | Variable                 |
|--------------------------------|----------|--------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------|
| Reference Database             | true     | The source database from which the profile will be created                                                                                       | REFERENCE_DATABASE=      |
| Type of the reference database | true     | The target type of the source database (oracle_database or rac_database)                                                                         | REFERENCE_DATABASE_TYPE= |
| Reference Node                 | false    | The reference node from which the profile will be created. Applicable only for RAC and Cluster targets                                           | REF_NODE_SELECTED=       |
| Backup mechanism for data      | false    | Variable to identify what mode needs to be used for creating backup data. Possible values are EXPORT, DBCA_TEMPLATE, RMAN, and THIN_PROVISIONING | DATA_CONTENT_MODE=RMAN   |



**Table 27–10 (Cont.) EMCLI Property Files for a New RMAN Backup**

| Name                                                      | Required | Description                                                                                                                                                                               | Variable                       |
|-----------------------------------------------------------|----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------|
| The content that will be included in the backup           | false    | The data content that gets included in the backup can be metadata or data. Possible values are METADATA_ONLY and METADATA_AND_DATA.                                                       | DATA_CONTENT=METADATA_AND_DATA |
| Reference Host Credentials                                | true     | The credentials that will be used for OS authentication. Format is <Cred Name:Owner>. If this is not provided, preferred credentials will be checked against NormalHostCreds.             | REF_HOST_CREDENTIALS=          |
| Database Credentials                                      | false    | The credentials that will be used for performing database authentication. Format is <Cred Name:Owner>. If this is not provided, preferred credentials will be checked against SysDbuCreds | REF_DB_CREDENTIALS=            |
| Profile Name                                              | true     | Name of the profile                                                                                                                                                                       | PROFILE_NAME=                  |
| Profile Description                                       | false    | Name with which the profile should be created.                                                                                                                                            | PROFILE_DESC=                  |
| Profile location                                          | true     | Location of the profile in software library                                                                                                                                               | PROFILE_LOCATION=              |
| Profile Version                                           | false    | Version of the profile. If none is specified, the Database version shall be used.                                                                                                         | PROFILE_VERSION=               |
| Profile Vendor                                            | false    | Vendor of the profile. Default is Oracle.                                                                                                                                                 | PROFILE_VENDOR=                |
| Profile notes                                             | false    | Notes for the profile.                                                                                                                                                                    | PROFILE_NOTES=                 |
| Working Directory                                         | false    | Working directory is the location where the files are staged.                                                                                                                             | WORKING_DIRECTORY=             |
| Is the profile initiated by end user                      | true     | Is the profile initiated by end user as part of backup process or the database.                                                                                                           | USER_INITIATED=false           |
| Is Database target available                              | false    | Is Database target available. Must be set to 'true' for this use case.                                                                                                                    | DB_TARGET_AVAILABLE=true       |
| Backup Type                                               | false    | Backup type.                                                                                                                                                                              | RMAN.BACKUP_TYPE=              |
| Number of RMAN Channels to use(Default No of Channels : 2 | false    | Number of Channels to be used. Possible Values: like 1, 2 or 3.                                                                                                                           | RMAN.RMAN_CHANNELS=            |
| Retain backup until                                       | false    | Retain backup until. Format: mm/dd/yyyy.                                                                                                                                                  | RMAN.RETAIN_UNTIL=             |
| Encryption Type                                           | false    | Encryption Type. Possible Values: ON, OFF.                                                                                                                                                | RMAN.ENCRYPTION_TYPE=          |

**Table 27–10 (Cont.) EMCLI Property Files for a New RMAN Backup**

| Name                          | Required | Description                                                                                           | Variable                    |
|-------------------------------|----------|-------------------------------------------------------------------------------------------------------|-----------------------------|
| Encryption Password/Key       | false    | Encryption key which was used to encrypt the backup.                                                  | RMAN.ENCRIPTION_PASSWORD=   |
| Compression Type              | false    | Compression type. Possible Values: HIGH, MEDIUM, LOW.                                                 | RMAN.COMPRESSION_TYPE=      |
| RMAN Backup Location          | true     | Location of the RMAN Backups.                                                                         | RMAN.RMAN_BACKUP_LOCATION=  |
| Backup file format            | false    | Backup file format.                                                                                   | RMAN.BACKUP_FILE_FORMAT=    |
| Backup file tag               | false    | Backup file tag.                                                                                      | RMAN.BACKUP_FILE_TAG=       |
| Source: Control File Location | false    | Location of the Control File.                                                                         | RMAN.CONTROL_FILE_NAME=     |
| Control File Tag              | false    | Control File Tag.                                                                                     | RMAN.CONTROL_FILE_TAG=      |
| Source: Database SID          | false    | Database SID of the source database.                                                                  | RMAN.ORACLE_DB_SID=         |
| Initialize File Name          | false    | Initialize file name.                                                                                 | RMAN.INIT_FILE_NAME=        |
| User List File Name           | false    | Name of the users list file. this file should be kept in location specified for RMAN_BACKUP_LOCATION. | RMAN.USERLIST_FILE_NAME=    |
| GDB Name of source Db         | false    | GDB Name of source Db.                                                                                | RMAN.DATABASE_NAME=         |
| use existing RMAN Backup      | true     | use existing RMAN Backup. This should be set to 'NO' for this use case.                               | RMAN.USE_EXISTING_BACKUP=NO |

### Existing RMAN Backup from the Source Database

**Table 27–11 EMCLI Property Files for a Existing RMAN Backup**

| Name                           | Required | Description                                                                                            | Variable                 |
|--------------------------------|----------|--------------------------------------------------------------------------------------------------------|--------------------------|
| Reference Database             | true     | The source database from which the profile will be created                                             | REFERENCE_DATABASE=      |
| Type of the reference database | true     | The target type of the source database (oracle_database or rac_database)                               | REFERENCE_DATABASE_TYPE= |
| Reference Node                 | false    | The reference node from which the profile will be created. Applicable only for RAC and Cluster targets | REF_NODE_SELECTED=       |

**Table 27–11 (Cont.) EMCLI Property Files for a Existing RMAN Backup**

| Name                                            | Required | Description                                                                                                                                                                               | Variable                       |
|-------------------------------------------------|----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------|
| Backup mechanism for data                       | false    | Variable to identify what mode needs to be used for creating backup data. Possible values are EXPORT, DBCA_TEMPLATE, RMAN, and THIN_PROVISIONING                                          | DATA_CONTENT_MODE=RMAN         |
| The content that will be included in the backup | false    | The data content that gets included in the backup can metadata or data. Possible values are METADATA_ONLY and METADATA_AND_DATA.                                                          | DATA_CONTENT=METADATA_AND_DATA |
| Reference Host Credentials                      | true     | The credentials that will be used for OS authentication. Format is <Cred Name:Owner>. If this is not provided, preferred credentials will be checked against NormalHostCreds.             | REF_HOST_CREDENTIALS=          |
| Database Credentials                            | false    | The credentials that will be used for performing database authentication. Format is <Cred Name:Owner>. If this is not provided, preferred credentials will be checked against SysDbuCreds | REF_DB_CREDENTIALS=            |
| Profile Name                                    | true     | Name of the profile                                                                                                                                                                       | PROFILE_NAME=                  |
| Profile Description                             | false    | Name with which the profile should be created.                                                                                                                                            | PROFILE_DESC=                  |
| Profile location                                | true     | Location of the profile in software library.                                                                                                                                              | PROFILE_LOCATION=              |
| Profile Version                                 | false    | Version of the profile. If none is specified, the Database version shall be used.                                                                                                         | PROFILE_VERSION=               |
| Profile Vendor                                  | false    | Vendor of the profile. Default is Oracle.                                                                                                                                                 | PROFILE_VENDOR=                |
| Profile notes                                   | false    | Notes for the profile.                                                                                                                                                                    | PROFILE_NOTES=                 |
| Working Directory                               | false    | Working directory is the location where the files are staged.                                                                                                                             | WORKING_DIRECTORY=             |
| Is the profile initiated by end user            | true     | Is the profile initiated by end user as part of backup process or the database.                                                                                                           | USER_INITIATED=false           |
| Is Database target available                    | false    | Is Database target available. Must be set to 'true' for this use case.                                                                                                                    | DB_TARGET_AVAILABLE=true       |

**Table 27–11 (Cont.) EMCLI Property Files for a Existing RMAN Backup**

| <b>Name</b>                                               | <b>Required</b> | <b>Description</b>                                                                                    | <b>Variable</b>              |
|-----------------------------------------------------------|-----------------|-------------------------------------------------------------------------------------------------------|------------------------------|
| Platform Name of Source Database                          | false           | Platform Name of Source Database. For example, Linux x86-64 or Oracle Solaris on x86-64 (64-bit).     | SOURCE_DB_PLATFORM_NAME=     |
| Backup Type                                               | false           | Backup type.                                                                                          | RMAN.BACKUP_TYPE=            |
| Number of RMAN Channels to use(Default No of Channels : 2 | false           | Number of Channels to be used. Possible Values: like 1, 2 or 3.                                       | RMAN.RMAN_CHANNELS=          |
| Retain backup until                                       | false           | Retain backup until. Format: mm/dd/yyyy.                                                              | RMAN.RETAIN_UNTIL=           |
| Encryption Type                                           | false           | Encryption Type. Possible Values: ON, OFF.                                                            | RMAN.ENCRYPTION_TYPE=        |
| Source:Encryption Password/Key                            | false           | Encryption key which was used to encrypt the backup.                                                  | RMAN.ENCRYPTION_PASSWORD=    |
| Compression Type                                          | false           | Compression type. Possible Values: HIGH, MEDIUM, LOW.                                                 | RMAN.COMPRESSION_TYPE=       |
| RMAN Backup Location                                      | true            | Location of the RMAN Backups.                                                                         | RMAN.RMAN_BACKUP_LOCATION=   |
| Backup file format                                        | false           | Backup file format.                                                                                   | RMAN.BACKUP_FILE_FORMAT=     |
| Backup file tag                                           | false           | Backup file tag.                                                                                      | RMAN.BACKUP_FILE_TAG=        |
| Source: Control File Location                             | false           | Location of the Control File.                                                                         | RMAN.CONTROL_FILE_NAME=      |
| Control File Tag                                          | false           | Control File Tag.                                                                                     | RMAN.CONTROL_FILE_TAG=       |
| Source: Database SID                                      | false           | Database SID of the source database.                                                                  | RMAN.ORACLE_DB_SID=          |
| Initialize File Name                                      | false           | Initialize file name.                                                                                 | RMAN.INIT_FILE_NAME=         |
| User List File Name                                       | false           | Name of the users list file. this file should be kept in location specified for RMAN_BACKUP_LOCATION. | RMAN.USERLIST_FILE_NAME=     |
| GDB Name of source Db                                     | false           | GDB Name of source Db.                                                                                | RMAN.DATABASE_NAME=          |
| use existing RMAN Backup                                  | true            | use existing RMAN Backup. This should be set to 'YES' for this use case.                              | RMAN.USE_EXISTING_BACKUP=YES |

**Existing RMAN Backup on the Stage Location****Table 27–12 EMCLI Property Files for a Existing RMAN Backup on the Stage Location**

| <b>Name</b>                                     | <b>Required</b> | <b>Description</b>                                                                                                                                                                        | <b>Variable</b>                |
|-------------------------------------------------|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------|
| Reference Database                              | true            | The source database from which the profile will be created                                                                                                                                | REFERENCE_DATABASE=            |
| Type of the reference database                  | true            | The target type of the source database (oracle_database or rac_database)                                                                                                                  | REFERENCE_DATABASE_TYPE=       |
| Reference Node                                  | false           | The reference node from which the profile will be created. Applicable only for RAC and Cluster targets                                                                                    | REF_NODE_SELECTED=             |
| Backup mechanism for data                       | false           | Variable to identify what mode needs to be used for creating backup data. Possible values are EXPORT, DBCA_TEMPLATE, RMAN, and THIN_PROVISIONING                                          | DATA_CONTENT_MODE=RMAN         |
| The content that will be included in the backup | false           | The data content that gets included in the backup can metadata or data. Possible values are METADATA_ONLY and METADATA_AND_DATA.                                                          | DATA_CONTENT=METADATA_AND_DATA |
| Reference Host Credentials                      | true            | The credentials that will be used for OS authentication. Format is <Cred Name:Owner>. If this is not provided, preferred credentials will be checked against NormalHostCreds.             | REF_HOST_CREDENTIALS=          |
| Database Credentials                            | false           | The credentials that will be used for performing database authentication. Format is <Cred Name:Owner>. If this is not provided, preferred credentials will be checked against SysDbuCreds | REF_DB_CREDENTIALS=            |
| Profile Name                                    | true            | Name of the profile                                                                                                                                                                       | PROFILE_NAME=                  |
| Profile Description                             | false           | Name with which the profile should be created.                                                                                                                                            | PROFILE_DESC=                  |
| Profile location                                | true            | Location of the profile in software library.                                                                                                                                              | PROFILE_LOCATION=              |
| Profile Version                                 | false           | Version of the profile. If none is specified, the Database version shall be used.                                                                                                         | PROFILE_VERSION=               |
| Profile Vendor                                  | false           | Vendor of the profile. Default is Oracle.                                                                                                                                                 | PROFILE_VENDOR=                |
| Profile notes                                   | false           | Notes for the profile.                                                                                                                                                                    | PROFILE_NOTES=                 |

**Table 27–12 (Cont.) EMCLI Property Files for a Existing RMAN Backup on the Stage**

| Name                                                      | Required | Description                                                                                           | Variable                   |
|-----------------------------------------------------------|----------|-------------------------------------------------------------------------------------------------------|----------------------------|
| Working Directory                                         | false    | Working directory is the location where the files are staged.                                         | WORKING_DIRECTORY=         |
| Is the profile initiated by end user                      | true     | Is the profile initiated by end user as part of backup process or the database.                       | USER_INITIATED=false       |
| Is Database target available                              | false    | Is Database target available. Must be set to 'false' for this use case.                               | DB_TARGET_AVAILABLE=false  |
| Platform Name of Source Database                          | false    | Platform Name of Source Database. For example, Linux x86-64 or Oracle Solaris on x86-64 (64-bit).     | SOURCE_DB_PLATFORM_NAME=   |
| Backup Type                                               | false    | Backup type.                                                                                          | RMAN.BACKUP_TYPE=          |
| Number of RMAN Channels to use(Default No of Channels : 2 | false    | Number of Channels to be used. Possible Values: like 1, 2 or 3.                                       | RMAN.RMAN_CHANNELS=        |
| Retain backup until                                       | false    | Retain backup until. Format: mm/dd/yyyy.                                                              | RMAN.RETAIN_UNTIL=         |
| Encryption Type                                           | false    | Encryption Type. Possible Values: ON, OFF.                                                            | RMAN.ENCRYPTION_TYPE=      |
| Source:Encryption Password/Key                            | false    | Encryption key which was used to encrypt the backup.                                                  | RMAN.ENCRYPTION_PASSWORD=  |
| Compression Type                                          | false    | Compression type. Possible Values: HIGH, MEDIUM, LOW.                                                 | RMAN.COMPRESSION_TYPE=     |
| RMAN Backup Location                                      | true     | Location of the RMAN Backups.                                                                         | RMAN.RMAN_BACKUP_LOCATION= |
| Backup file format                                        | false    | Backup file format. Possible Values: %U                                                               | RMAN.BACKUP_FILE_FORMAT=   |
| Source: Control File Location                             | false    | Location of the Control File.                                                                         | RMAN.CONTROL_FILE_NAME=    |
| Control File Tag                                          | false    | Control File Tag. Default: control01.ctl                                                              | RMAN.CONTROL_FILE_TAG=     |
| Source: Database SID                                      | false    | Database SID of the source database.                                                                  | RMAN.ORACLE_DB_SID=        |
| Initialize File Name                                      | false    | Initialize file name.                                                                                 | RMAN.INIT_FILE_NAME=       |
| User List File Name                                       | false    | Name of the users list file. this file should be kept in location specified for RMAN_BACKUP_LOCATION. | RMAN.USERLIST_FILE_NAME=   |
| GDB Name of source Db                                     | false    | GDB Name of source Db.                                                                                | RMAN.DATABASE_NAME=        |
| RMDB Version                                              | false    | RMDBS version                                                                                         | RMAN.DATABASE_VERSION=     |

**Snapclone Database****Table 27–13 EMCLI Property Files for Snapclone Database**

| <b>Name</b>                                     | <b>Required</b> | <b>Description</b>                                                                                                                                                                        | <b>Variable</b>                     |
|-------------------------------------------------|-----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------|
| Reference Database                              | true            | The source database from which the profile will be created                                                                                                                                | REFERENCE_DATABASE=                 |
| Type of the reference database                  | true            | The target type of the source database (oracle_database or rac_database)                                                                                                                  | REFERENCE_DATABASE_TYPE=            |
| Reference Node                                  | false           | The reference node from which the profile will be created. Applicable only for RAC and Cluster targets                                                                                    | REF_NODE_SELECTED=                  |
| Backup mechanism for data                       | false           | Variable to identify what mode needs to be used for creating backup data. Possible values are EXPORT, DBCA_TEMPLATE, RMAN, and THIN_PROVISIONING                                          | DATA_CONTENT_MODE=THIN_PROVISIONING |
| The content that will be included in the backup | false           | The data content that gets included in the backup can be metadata or data. Possible values are METADATA_ONLY and METADATA_AND_DATA.                                                       | DATA_CONTENT=METADATA_AND_DATA      |
| Reference Host Credentials                      | true            | The credentials that will be used for OS authentication. Format is <Cred Name:Owner>. If this is not provided, preferred credentials will be checked against NormalHostCreds.             | REF_HOST_CREDENTIALS=               |
| Database Credentials                            | false           | The credentials that will be used for performing database authentication. Format is <Cred Name:Owner>. If this is not provided, preferred credentials will be checked against SysDbuCreds | REF_DB_CREDENTIALS=                 |
| Profile Name                                    | true            | Name of the profile                                                                                                                                                                       | PROFILE_NAME=                       |
| Profile Description                             | false           | Name with which the profile should be created.                                                                                                                                            | PROFILE_DESC=                       |
| Profile location                                | true            | Location of the profile in software library.                                                                                                                                              | PROFILE_LOCATION=                   |
| Profile Version                                 | false           | Version of the profile. If none is specified, the Database version shall be used.                                                                                                         | PROFILE_VERSION=                    |
| Profile Vendor                                  | false           | Vendor of the profile. Default is Oracle.                                                                                                                                                 | PROFILE_VENDOR=                     |
| Profile notes                                   | false           | Notes for the profile.                                                                                                                                                                    | PROFILE_NOTES=                      |

**Table 27–13 (Cont.) EMCLI Property Files for Snapclone Database**

| Name                                 | Required | Description                                                                     | Variable             |
|--------------------------------------|----------|---------------------------------------------------------------------------------|----------------------|
| Working Directory                    | false    | Working directory is the location where the files are staged.                   | WORKING_DIRECTORY=   |
| Is the profile initiated by end user | true     | Is the profile initiated by end user as part of backup process or the database. | USER_INITIATED=false |

### Schema Database

**Table 27–14 EMCLI Property Files for Schema Database**

| Name                                            | Required | Description                                                                                                                                                                               | Variable                       |
|-------------------------------------------------|----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------|
| Reference Database                              | true     | The source database from which the profile will be created                                                                                                                                | REFERENCE_DATABASE=            |
| Type of the reference database                  | true     | The target type of the source database (oracle_database or rac_database)                                                                                                                  | REFERENCE_DATABASE_TYPE=       |
| Reference Node                                  | false    | The reference node from which the profile will be created. Applicable only for RAC and Cluster targets                                                                                    | REF_NODE_SELECTED=             |
| Backup mechanism for data                       | false    | Variable to identify what mode needs to be used for creating backup data. Possible values are EXPORT, DBCA_TEMPLATE, RMAN, and THIN_PROVISIONING                                          | DATA_CONTENT_MODE=EXPORT       |
| The content that will be included in the backup | false    | The data content that gets included in the backup can be metadata or data. Possible values are METADATA_ONLY and METADATA_AND_DATA.                                                       | DATA_CONTENT=METADATA_AND_DATA |
| Reference Host Credentials                      | true     | The credentials that will be used for OS authentication. Format is <Cred Name:Owner>. If this is not provided, preferred credentials will be checked against NormalHostCreds.             | REF_HOST_CREDENTIALS=          |
| Database Credentials                            | false    | The credentials that will be used for performing database authentication. Format is <Cred Name:Owner>. If this is not provided, preferred credentials will be checked against SysDbuCreds | REF_DB_CREDENTIALS=            |
| Profile Name                                    | true     | Name of the profile                                                                                                                                                                       | PROFILE_NAME=                  |



**Table 27–14 (Cont.) EMCLI Property Files for Schema Database**

| <b>Name</b>            | <b>Required</b> | <b>Description</b>                                                                                                               | <b>Variable</b>               |
|------------------------|-----------------|----------------------------------------------------------------------------------------------------------------------------------|-------------------------------|
| Profile Description    | false           | Name with which the profile should be created.                                                                                   | PROFILE_DESC=                 |
| Profile location       | true            | Location of the profile in software library.                                                                                     | PROFILE_LOCATION=             |
| Profile Version        | false           | Version of the profile. If none is specified, the Database version shall be used.                                                | PROFILE_VERSION=              |
| Profile Vendor         | false           | Vendor of the profile. Default is Oracle.                                                                                        | PROFILE_VENDOR=               |
| Profile notes          | false           | Notes for the profile.                                                                                                           | PROFILE_NOTES=                |
| Data gold image name   | true            | Entity name of the data gold image                                                                                               | DATA_GOLD_IMAGE_ENTITY_NAME=  |
| Working Directory      | false           | Working directory is the location where the files are staged.                                                                    | WORKING_DIRECTORY=            |
| Schemas to be exported | true            | List of schemas that needs to be included as part of the export.                                                                 | EXPORT.SCHEMA_INCLUDE_LIST.0= |
| Dump Directories       | false           | List of directory objects that needs to be used for storing export. Format: directory=dir1,file_name=file1%U.dmp[,max_size=1MB]. | EXPORT.DUMP_DIRECTORY_LIST.0= |
| Log Directory          | false           | Log directory location where the log file generated during export.                                                               | PORT.LOG_FILE_DIRECTORY=      |
| Degree of Parallelism  | false           | Degree of Parallelism indicates the number of threads. For example, 1 or 2 or 3...                                               | EXPORT.DEGREE_OF_PARALLELISM= |



---

## Java as a Service APIs

This chapter describes the resource model for Java as a Service (JaaS) API and the REST request/response interactions. Application component filtering, which is supported for JaaS zones, templates, is also described.

This chapter contains the following topics:

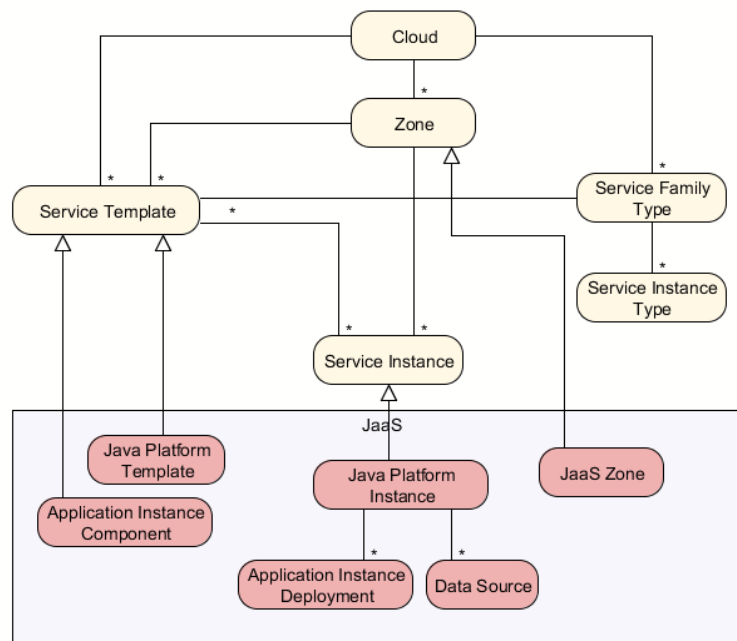
- [Java as a Service Resource Data Models](#)
- [Application Component Filtering](#)

### 28.1 Java as a Service Resource Data Models

The following sections describe the resource models that are supported by the Java as a Service family type and provide examples of supported operations. The following resource models are described:

- [Cloud](#)
- [Zone](#)
- [JavaPlatformTemplate](#)
- [JavaPlatformInstance](#)
- [ApplicationInstanceDeployment](#)
- [DataSource](#)
- [ApplicationInstanceComponent](#)

The figure below shows the Java as a Service resource model.



### 28.1.1 Cloud

The following table describes the Cloud [application/oracle.com.cloud.common.Cloud+json] data model.

**Table 28–1 Cloud Data Model**

| Field               | Type                          | Occurs | Description                                                                                          |
|---------------------|-------------------------------|--------|------------------------------------------------------------------------------------------------------|
| uri                 | URI                           | 1      | A GET against this URI refreshes the client representation of the resources accessible to this user. |
| name                | String                        | 1      | A human readable name. It is a UNICODE string to support different languages.                        |
| description         | String                        | 0..1   | A brief description. This is a UNICODE string to support different languages.                        |
| service_template    | Collection<ServiceTemplates>  | 0..1   | List of service templates that are accessible to the user.                                           |
| zones               | Collection<Zone>              | 0..1   | List of zones that are supported by the cloud and accessible to the user.                            |
| service_family_type | Collection<ServiceFamilyType> | 0..1   | The list of service family types that are supported by the cloud.                                    |

**Table 28–1 (Cont.) Cloud Data Model**

| Field          | Type          | Occurs | Description                                                                                                                                                                                                                                                            |
|----------------|---------------|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| resource_state | ResourceState | 0..1   | A cloud that is online and running has READY as its state. If this field is not returned, the client can assume the cloud is READY. If the state of the returned field is not READY, the client cannot assume the viability of subsequent interactions into the cloud. |
| media_type     | String        | 1      | The value of this media type with the payload format. For example, application/oracle.com.cloud.common.Cloud+json.                                                                                                                                                     |
| canonicalLink  | String        | 1      | GET against this normalized/canonical URI refreshes the client representation of the resources accessible to this user.                                                                                                                                                |

### 28.1.1.1 GET Operation on a Cloud URI to return its representation

The following shows an example of a GET operation on the Cloud URI to return its representation.

#### Example Request:

```
GET /em/cloud
Host: cloudcompany.com
Authorization: Basic xxxxxxxx
Accept: application/oracle.com.cloud.common.Cloud+json
```

#### Example Response:

```
{
 "uri" : "/em/cloud" ,
 "name" : "Oracle Cloud by Enterprise Manager" ,
 "description" : "This represents the Cloud resource of the Oracle Enterprise
Manager Cloud Management solution" ,
 "resource_state" : {
 "state" : "READY"
 } ,
 "media_type" : "application/oracle.com.cloud.common.Cloud+json" ,
 "canonicalLink": "/em/websvcs/restful/extws/cloudservices/service/v0
/ssa/em/cloud" ,
 "zones" : {
 "media_type" : "application/oracle.com.cloud.common.Zone+json" ,
 "total" : "3" ,
 "elements" :
 [
 {
 "uri" : "/em/cloud/jaas/zone/A1B44A4EBCC4563125D9D0A3AAE4FD51" ,
 "name" : "Zone1" ,
 "media_type" : "application/oracle.com.cloud.jaas.Zone+json" ,
 "service_family_type" : "jaas"
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/zone/A1B44A4
EBCC4563125D9D0A3AAE4FD51"
 } ,
 {
 "uri" : "/em/cloud/dbaas/zone/A1B44A4EBCC4563125D9D0A3AAE4FD51" ,
 "name" : "Zone1" ,
 "description" : "Zone for Physical Pool" ,
```

```

 "media_type" : "application/oracle.com.cloud.common.DbZone+json" ,
 "type" : "self_service_zone"
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/dbaas/zone/A1B44A
4EBCC4563125D9D0A3AAE4FD51"
 } ,
 {
 "uri" : "/em/cloud/opc/opczone" ,
 "name" : "OPC Zone" ,
 "media_type" : "application/oracle.com.cloud.opc.OpcZone+json" ,
 "service_family_type" : "opc" ,
 "type" : "opc"
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/service_family_
type/opc"
 }
]
} ,
"service_templates" : {
 "media_type" : "application/oracle.com.cloud.common.ServiceTemplate+json" ,
 "total" : "5" ,
 "elements" :
 [
 {
 "uri" :
"/em/cloud/jaas/javaplatformtemplate/D016AADD465462ACE040F20AD18167BA" ,
 "name" : "PS4_LowHeapTemplate" ,
 "media_type" :
"application/oracle.com.cloud.jaas.JavaPlatformTemplate+json" ,
 "service_family_type" : "jaas"
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/javaplatform
template/ D016AADD465462ACE040F20AD18167BA"
 } ,
 {
 "uri" :
"/em/cloud/jaas/applicationinstancecomponent/oracle%3AdefaultService%3Aem%3Aprovis
ioning%3A1%3Acmp%3ACOMP_Component%3ASUB_
JavaEEApplication%3AD03030FAE80F6B58E040F20AD1812FE2%3A0.1" ,
 "name" : " rtriddappjul18_Deploy_20120718_051319824" ,
 "media_type" :
"application/oracle.com.cloud.jaas.ApplicationInstanceComponent+json" ,
 "service_family_type" : "jaas"
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/applicati
oninstancecomponent/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp%3ACOMP_
Component%3ASUB_JavaEEApplication%3AD03030FAE80F6B58E040F20AD1812FE2%3A0.1"
 } ,
 {
 "uri" :
"/em/cloud/jaas/applicationinstancecomponent/oracle%3AdefaultService%3Aem%3Aprovis
ioning%3A1%3Acmp%3ACOMP_Component%3ASUB_
JavaEEApplication%3AD03517A824EBFCF9E040F20AD1811CD9%3A0.1" ,
 "name" : "FileUploadFromCode" ,
 "media_type" :
"application/oracle.com.cloud.jaas.ApplicationInstanceComponent+json" ,
 "service_family_type" : "jaas"
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/applicati
oninstancecomponent/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp%3ACOMP_

```

```

Component%3ASUB_JavaEEApplication%3AD03517A824EBFCF9E040F20AD1811CD9%3A0.1"
 } ,
 {
 "uri" :
"/em/cloud/jaas/applicationinstancecomponent/oracle%3AdefaultService%3Aem%3Aprovis
ioning%3A1%3Acmp%3ACOMP_Component%3ASUB_
JavaEEApplication%3AD03EFA365D58A674E040F20AD18134AE%3A0.1" ,
 "name" : "Test_1354864500315" ,
 "media_type" :
"application/oracle.com.cloud.jaas.ApplicationInstanceComponent+json" ,
 "service_family_type" : "jaas"
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/applicati
oninstancecomponent/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp%3ACOMP_
Component%3ASUB_JavaEEApplication%3AD03EFA365D58A674E040F20AD18134AE%3A0.1"
 } ,
 {
 "uri" :
"/em/cloud/jaas/applicationinstancecomponent/oracle%3AdefaultService%3Aem%3Aprovis
ioning%3A1%3Acmp%3ACOMP_Component%3ASUB_
JavaEEApplication%3AD1103FE5F6877372E040F20AD1814215%3A0.1" ,
 "name" : "RiddleApp" ,
 "media_type" :
"application/oracle.com.cloud.jaas.ApplicationInstanceComponent+json" ,
 "service_family_type" : "jaas"
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/applicati
oninstancecomponent/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp%3ACOMP_
Component%3ASUB_JavaEEApplication%3AD1103FE5F6877372E040F20AD1814215%3A0.1"
 }
]
 } ,
 "service_family_types" : {
 "media_type" : "application/oracle.com.cloud.common.ServiceFamilyType+json"
 } ,
 "total" : "4" ,
 "elements" :
 [
 {
 "uri" : "/em/cloud/service_family_type/jaas" ,
 "name" : "jaas" ,
 "media_type" :
"application/oracle.com.cloud.common.ServiceFamilyType+json" ,
 "type" : "jaas"
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/service_family_
type/jaas"
 } ,
 {
 "uri" : "/em/cloud/service_family_type/dbaas" ,
 "name" : "dbaas" ,
 "media_type" :
"application/oracle.com.cloud.common.ServiceFamilyType+json" ,
 "type" : "dbaas"
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/service_family_
type/dbaas"
 } ,
 {
 "uri" : "/em/cloud/service_family_type/opc" ,

```

```

 "name" : "opc" ,
 "media_type" :
"application/oracle.com.cloud.common.ServiceFamilyType+json" ,
 "type" : "opc"
 canonicalLink:
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/service_family_
type/opc"
 } ,
 {
 "uri" : "/em/cloud/service_family_type/iaas" ,
 "name" : "iaas" ,
 "media_type" :
"application/oracle.com.cloud.iaas.IaasServiceFamilyType+json" ,
 "type" : "iaas"
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/service_family_
type/iaas"
 }
]
 } ,
 "service_requests" : {
 "media_type" : "application/oracle.com.cloud.common.Request+json" ,
 "total" : "0" ,
 "elements" :
 [
]
 }
}

```

### 28.1.2 JaaS Service Family Type

Jaas service family type is a category of services that are offered by Java as a Service (JaaS). The following table describes the Jaas ServiceFamilyType [application/oracle.com.cloud.common.ServiceFamilyType+json] data model:

**Table 28–2 JaaS Service Family Type**

| Field             | Type                                | Occurs | Description                                                                                                       |
|-------------------|-------------------------------------|--------|-------------------------------------------------------------------------------------------------------------------|
| uri               | URI                                 | 1      | A GET against this URI refreshes the client representation of the Jaas ServiceFamilyType definition to this user. |
| name              | String                              | 1      | Name of the Jaas ServiceFamilyType.                                                                               |
| description       | String                              | 1      | Human readable description of the Service Family Type. Shall be a UNICODE string to support different languages.  |
| resource_state    | ResourceState                       | 0..1   | The resource state of the resource.                                                                               |
| media_type        | String                              | 1      | The media type of the resource.                                                                                   |
| service_templates | Collection<br><ServiceTempl<br>ate> | 0..1   | The collection of service templates that are of this service family type.                                         |
| zones             | Collection<br><Zone>                | 0..1   | The collection of zones that support service instances of this service family type.                               |
| instance_types    | Collection<br><Instance<br>Type>    | 0..1   | The collection of instance types that are supported by this service family type.                                  |



**Table 28–2 (Cont.) JaaS Service Family Type**

| Field         | Type   | Occurs | Description                                                                                                             |
|---------------|--------|--------|-------------------------------------------------------------------------------------------------------------------------|
| canonicalLink | String | 1      | GET against this normalized/canonical URI refreshes the client representation of the resources accessible to this user. |

**28.1.2.1 GET on JaaS Service Family Type URI to Return its Representation**

The following shows an example of a GET operation on the JaaS Service Family Type URI to return its representation.

**Example Request:**

```
GET /em/cloud/service_family_type/jaas
Host: cloudcompany.com
Authorization: Basic xxxxxxxx
Accept: application/oracle.com.cloud.common.ServiceFamilyType+json
```

**Example Response:**

```
{
 "uri" : "/em/cloud/service_family_type/jaas" ,
 "name" : "jaas" ,
 "resource_state" : {
 "state" : "READY"
 } ,
 "media_type" : "application/oracle.com.cloud.common.ServiceFamilyType+json" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/service_family_type/jaas" ,
 "type" : "jaas" ,
 "service_templates" : {
 "media_type" : "application/oracle.com.cloud.common.ServiceTemplate+json" ,
 "total" : "1" ,
 "elements" :
 [
 {
 "uri" :
"/em/cloud/jaas/javaplatformtemplate/D016AADD465462ACE040F20AD18167BA" ,
 "name" : "PS4_LowHeapTemplate" ,
 "media_type" :
"application/oracle.com.cloud.jaas.JavaPlatformTemplate+json" ,
 "service_family_type" : "jaas" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/javaplatformtemplate/D016AADD465462ACE040F20AD18167BA"
 }
]
 }
 "zones" : {
 "media_type" : "application/oracle.com.cloud.common.Zone+json" ,
 "total" : "1" ,
 "elements" :
 [
 {
 "uri" : "/em/cloud/jaas/zone/A1B44A4EBCC4563125D9D0A3AAE4FD51" ,
 "name" : "Zone1" ,
 "media_type" : "application/oracle.com.cloud.jaas.Zone+json" ,
 "service_family_type" : "jaas" ,
```

```
 "type" : "jaas" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/zone/A1B44A4
EBCC4563125D9D0A3AAE4FD51"
 }
}
} ,
"instance_types" : {
 "media_type" : "application/oracle.com.cloud.common.InstanceType+json" ,
 "total" : "0" ,
 "elements" :
 [
]
} ,
"service_requests" : {
 "total" : "0" ,
 "elements" :
 [
]
}
}
```

### 28.1.2.2 POST on the Jaas Service Family Type URI to Return ApplicationInstanceComponent

ApplicationInstanceComponent can be created using a post request on JaaS. Right now this is being done using a Java Program. The following shows a sample program:

---

---

**Note:** <<>> : Provides a hint to an appropriate value.

---

---

```
import java.io.File;
import java.io.FileInputStream;
import java.io.PrintWriter;
import java.io.StringWriter;

import java.util.ArrayList;
import java.util.List;
import java.util.Properties;

import oracle.sysman.emInternalSDK.ssa.cloudapi.client.EMRestService;

public class FileUpload {
 static public void main(String[] args)
 throws Exception {
 String emUser = <<JAAS_USER_NAME>>;
 String emPassword = <<password>>;
 String emHost = <<host_machine_name/ip>>;
 String emPort = <<port>>;
 String emLocation = "https://" + emHost + ":" + emPort;
 String applicationLocation = <<application_location>>;
 String deployPlanLocation = <<deploy_plan_location>>;
 File f = new File(applicationLocation + "deploy_app.ear");
 StringWriter sw = new StringWriter();
 PrintWriter pw = new PrintWriter(sw);
 pw.println("{");
 String fileName = "MyApplication" + emUser + "_" + System.currentTimeMillis();
 pw.println("\nname\":" + "\"" + fileName + "\"");
 pw.println("\ndescription\":" + "\"Riddles App\"");
 pw.println("\noriginate\":" + "\"STREAMING\"");
 pw.println("\nfiles\":[");
```

```

 pw.println("{}");
 //pw.println("\"path\": \"swlib_internal_path\",");
 pw.println("\"file_archive_path\": \"swlib_internal_path\",");
 pw.println("\"content_type\": \"application/octet-stream\",");
 pw.println("\"source\": \"\"+f.getName()+"\"");
 pw.println("{}");
 pw.println("]");
 pw.print("{}");
 pw.flush();
 pw.close();
 String content = sw.toString();
 System.out.println("CONTENT:\n"+content+"\n");
 //String content = "{\n\"name\": \"STOPPED\"\n}";
 List<EMRestService.MultipartItem> items = new
ArrayList<EMRestService.MultipartItem>();
 EMRestService.MultipartItem item1, item2, item3, item4, item5;
 item1 = new EMRestService.MultipartItem("item1", false);
 item1.setContentType("multipart/form-data");
 item1.setContent(content);
 items.add(item1);
 item2 = new EMRestService.MultipartItem("item2", true);
 item2.setFileAlias(f.getName());
 item2.setFile(f);
 items.add(item2);
 File deploymentPlanFile = new File(applicationLocation+"deployment_
plan.xml");
 item3 = new EMRestService.MultipartItem("item3", true);
 item3.setFileAlias(deploymentPlanFile.getName());
 item3.setFile(deploymentPlanFile);
 items.add(item3);
 File preDeployFile = new File(deployPlanLocation+" predeploy_script.py");
 item4 = new EMRestService.MultipartItem("item4", true);
 item4.setFileAlias(preDeployFile.getName());
 item4.setFile(preDeployFile);
 items.add(item4);
 File postDeployFile = new File(deployPlanLocation+" postdeploy_
script.py");
 item5 = new EMRestService.MultipartItem("item5", true);
 item5.setFileAlias(postDeployFile.getName());
 item5.setFile(postDeployFile);
 items.add(item5);

 String uri = emLocation+"/em";
 EMRestService ers = new EMRestService(uri, emUser, emPassword.getBytes());
 ers.setReadTimeout(600*1000);
 Object o = ers.webserviceMultipart(EMRestService.Method.POST,
uri+"/cloud/service_family_type/jaas",
"application/oracle.com.cloud.jaas.ApplicationInstanceComponent", items);
 System.out.println(o);
 }
}

```

### 28.1.3 Zone

The following table describes the Zone  
[application/oracle.com.cloud.common.Zone+json] data model.

**Table 28–3 Zone Data Model**

| Field               | Type                            | Occurs | Description                                                                                                                           |
|---------------------|---------------------------------|--------|---------------------------------------------------------------------------------------------------------------------------------------|
| uri                 | URI                             | 1      | A GET against this URI refreshes the client representation of the zone definition to this user.                                       |
| name                | String                          | 1      | Name of the Zone.                                                                                                                     |
| description         | String                          | 0..1   | Human readable description of the Zone. It is a UNICODE string to support different languages.                                        |
| context_id          | String                          | 1      | This is an ID indication the internal reference of the zone                                                                           |
| service_instances   | Collection<br><ServiceInstance> | 0..1   | Collection of the service instances that are in this zone.                                                                            |
| service_family_type | String                          | 0..1   | The name of the service family type that this zone is associated with.<br><br>Each Zone is contextualized in the service family type. |
| resource_state      | ResourceState                   | 0..1   | The resource state of the resource.                                                                                                   |
| media_type          | String                          | 1      | The media type of the resource.                                                                                                       |
| canonicalLink       | String                          | 1      | GET against this normalized/canonical URI refreshes the client representation of the resources accessible to this user.               |

### 28.1.3.1 GET Operation on a Zone URI to return its representation

The following shows an example of a GET operation on the Zone URI to return its representation.

#### Example Request:

```
GET /em/cloud/jaas/zone/ A1B44A4EBCC4563125D9D0A3AAE4FD51
Host: cloudcompany.com
Authorization: Basic xxxxxxxx
Accept: application/oracle.com.cloud.common.Zone+json
```

#### Example Response:

```
{
 "uri" : "/em/cloud/jaas/zone/A1B44A4EBCC4563125D9D0A3AAE4FD51" ,
 "name" : "Zone1" ,
 "description" : "Zone for Physical Pool" ,
 "resource_state" : {
 "state" : "READY"
 } ,
 "context_id" : "A1B44A4EBCC4563125D9D0A3AAE4FD51" ,
 "media_type" : "application/oracle.com.cloud.jaas.Zone+json" ,
 "service_family_type" : "jaas" ,
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/zone/
A1B44A4EBCC4563125D9D0A3AAE4FD51" ,
 "service_instances" : {
 "media_type" : "application/oracle.com.cloud.common.ServiceInstance+json" ,
 "total" : "1" ,
 "elements" :
[
```

```

 {
 "uri" :
"/em/cloud/jaas/javaplatforminstance/8AF6E65447FBA84E968CCD909BD05D42" ,
 "name" : "MyServer-1" ,
 "media_type" :
"application/oracle.com.cloud.jaas.JavaPlatformInstance+json" ,
 "status" : "RUNNING"
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/javaplatform
instance/8AF6E65447FBA84E968CCD909BD05D42"
 }
]
}
}

```

### 28.1.4 JavaPlatformTemplate

The JavaPlatformTemplate extends the CloudServiceTemplate SDK class. It is permissible to accept the ServiceTemplate media type on an instance of JavaPlatformTemplate. It is a preconfigured deployable entity that realizes a JavaPlatformInstance resource.

The following table describes the JavaPlatformTemplate [application/oracle.com.cloud.jaas.JavaPlatformTemplate+json] data model.

**Table 28–4 JavaPlatformTemplate Data Model**

| Field                       | Type                                 | Occurs | Description                                                                                                                                                         |
|-----------------------------|--------------------------------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| uri                         | URI                                  | 1      | A GET against this URI refreshes the client representation of the resources accessible to this user.                                                                |
| name                        | String                               | 1      | A human readable name given to the template.                                                                                                                        |
| description                 | String                               | 0...1  | A brief description given to the template.                                                                                                                          |
| created                     | Timestamp                            | 1      | Date and time, in ISO 8601 format, when the template was created/last modified.                                                                                     |
| zones                       | Collection<br><Zone>                 | 1      | List of all zones configured for this template.                                                                                                                     |
| service_family_type         | String                               | 1      | Denotes the type of Service Template (for example, "jaas").                                                                                                         |
| resource_state              | ResourceState                        | 1      | The validity of the fields on a GET should be guaranteed only when the resource state is READY. Otherwise, the client should not assume the validity of the fields. |
| service_instances           | Collection<br><JavaPlatformInstance> | 0...1  | All service instances created based on this JavaPlatformTemplate.                                                                                                   |
| media_type                  | String                               | 1      | The media type of the service template that this type represents.                                                                                                   |
| default_instance_media_type | String                               | 1      | Default media type for the request, value "application/oracle.com.cloud.jaas.JavaPlatformInstance+json".                                                            |
| context_id                  | String                               | 1      | This is an ID indication the internal reference of the service template.                                                                                            |

**Table 28–4 (Cont.) JavaPlatformTemplate Data Model**

| Field         | Type   | Occurs | Description                                                                                                             |
|---------------|--------|--------|-------------------------------------------------------------------------------------------------------------------------|
| canonicalLink | String | 1      | GET against this normalized/canonical URI refreshes the client representation of the resources accessible to this user. |

#### 28.1.4.1 GET Operation Supported on a JavaPlatformTemplate

The following shows an example of a GET operation on the JavaPlatformTemplate URI to return its representation.

##### Example Request:

```
GET /em/cloud/jaas/javaplatfromtemplate/D016AADD465462ACE040F20AD18167BA
Host: cloudcompany.com
Authorization: Basic xxxxxxxx
Accept: application/oracle.com.cloud.jaas.JavaPlatformTemplate+json
```

##### Example Response:

```
{
 "uri" : "/em/cloud/jaas/javaplatfromtemplate/D016AADD465462ACE040F20AD18167BA" ,
 "name" : "PS4_LowHeapTemplate" ,
 "description" : "This is a service template using the CDP CDP_PS4_lowHeap_
MWProfile" ,
 "resource_state" : {
 "state" : "READY"
 } ,
 "context_id" : "D016AADD465462ACE040F20AD18167BA" ,
 "media_type" : "application/oracle.com.cloud.jaas.JavaPlatformTemplate+json" ,
 "service_family_type" : "jaas" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/javaplatfrom
template/ D016AADD465462ACE040F20AD18167BA" ,
 "created" : "2012-12-04 23:16:10.0" ,
 "default_instance_media_type" :
"application/oracle.com.cloud.jaas.JavaPlatformInstance+json" ,
 "zones" : {
 "media_type" : "application/oracle.com.cloud.common.Zone+json" ,
 "total" : "1" ,
 "elements" :
[
 {
 "uri" : "/em/cloud/jaas/zone/A1B44A4EBCC4563125D9D0A3AAE4FD51" ,
 "name" : "Zone1" ,
 "media_type" : "application/oracle.com.cloud.jaas.Zone+json" ,
 "service_family_type" : "jaas"
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/service_family_
type/jaas"
 }
]
 } ,
 "service_instances" : {
 "media_type" : "application/oracle.com.cloud.common.ServiceInstance+json" ,
 "total" : "1" ,
 "elements" :
[
 {
 "uri" :
```

```

"/em/cloud/jaas/javaplatforminstance/8AF6E65447FBA84E968CCD909BD05D42" ,
 "name" : "MyServer-1" ,
 "media_type" :
"application/oracle.com.cloud.jaas.JavaPlatformInstance+json" ,
 "status" : "RUNNING"
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/javaplatform
instance/8AF6E65447FBA84E968CCD909BD05D42"
 }
]
 }
}

```

#### 28.1.4.2 POST Operation on a JavaPlatformTemplate

The following shows a sample POST on the JavaPlatformTemplate URI to create an "empty" JavaPlatformInstance by providing:

- Name of JavaPlatformInstance in "name" attribute
- Destination Zone URI in "zone" attribute

##### Example Request:

```

POST /em/cloud/jaas/javaplatformtemplate/ D016AADD465462ACE040F20AD18167BA
Host: cloudcompany.com
Authorization: Basic xxxxxxxx
Accept: application/oracle.com.cloud.jaas.JavaPlatformInstance+json
Content-Type: application/oracle.com.cloud.jaas.JavaPlatformInstance+json
{
 "name" : "rtjul17",
 "zone" : "/em/cloud/jaas/zone/C355810A6"
}

```

##### Example Response (When the operation is in progress):

```

{
 "uri" : "/em/cloud/jaas/javaplatforminstancerequest/1" ,
 "name" : "rtjul17_Create_20120717_045804697" ,
 "resource_state" : {
 "state" : "INITIATED" ,
 "messages" :
 [
 {
 "text" : "null" ,
 "date" : "2012-07-17T09:16:13+0000"
 }
]
 } ,
 "context_id": "1",
 "media_type": "application/oracle.com.cloud.jaas.JavaPlatformInstance+json" ,
 "service_family_type": "jaas" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/javaplatform
instancerequest/1" ,
 "status": "SCHEDULED"
 "created" : "2013-05-09T07:06:57+0000"
}

```

##### Example Response (After the instance has been created):

```

{

```

```

 "uri" : "/em/cloud/jaas/javaplatformtemplate/D016AADD465462ACE040F20AD18167BA" ,
 "name" : "PS4_LowHeapTemplate" ,
 "description" : "This is a service template using the CDP CDP_PS4_lowHeap_
MWProfile" ,
 "resource_state" : {
 "state" : "READY"
 } ,
 "context_id" : "D016AADD465462ACE040F20AD18167BA" ,
 "media_type" : "application/oracle.com.cloud.jaas.JavaPlatformTemplate+json" ,
 "service_family_type" : "jaas" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/javaplatform
template/ D016AADD465462ACE040F20AD18167BA" ,
 "created" : "2012-12-04 23:16:10.0" ,
 "default_instance_media_type" :
"application/oracle.com.cloud.jaas.JavaPlatformInstance+json" ,
 "zones" : {
 "media_type" : "application/oracle.com.cloud.common.Zone+json" ,
 "total" : "1" ,
 "elements" :
 [
 {
 "uri" : "/em/cloud/jaas/zone/A1B44A4EBCC4563125D9D0A3AAE4FD51" ,
 "name" : "Zone1" ,
 "media_type" : "application/oracle.com.cloud.jaas.Zone+json" ,
 "service_family_type" : "jaas"
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/service_family_
type/jaas"
 }
]
 } ,
 "service_instances" : {
 "media_type" : "application/oracle.com.cloud.common.ServiceInstance+json" ,
 "total" : "1" ,
 "elements" :
 [
 {
 "uri" :
"/em/cloud/jaas/javaplatforminstance/8AF6E65447FBA84E968CCD909BD05D42" ,
 "name" : "MyServer-1" ,
 "media_type" :
"application/oracle.com.cloud.jaas.JavaPlatformInstance+json" ,
 "status" : "RUNNING"
 "canonicalLink":
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/javaplatform
instance/8AF6E65447FBA84E968CCD909BD05D42"
 }
]
 }
}

```

### 28.1.4.3 GET on the Java Platform Request URL

The following shows a sample GET operation on the JavaPlatform Request URL:

#### Example Request:

```

GET /em/cloud/jaas/javaplatforminstancerequest/1
Host: cloudcompany.com
Authorization: Basic xxxxxxxx

```



Accept: application/oracle.com.cloud.jaas.JavaPlatformInstance+json

### Example Response (When the operation is in progress):

```
{
 "uri" : "/em/cloud/jaas/javaplatfrominstancerequest/1" ,
 "name" : "rtjul17_Create_20120717_045804697" ,
 "resource_state" : {
 "state" : "INITIATED" ,
 "messages" :
 [
 {
 "text" : "null" ,
 "date" : "2012-07-17T09:16:13+0000"
 }
]
 } ,
 "context_id": "1",
 "media_type": "application/oracle.com.cloud.jaas.JavaPlatformInstance+json" ,
 "service_family_type": "jaas" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/javaplatfrominstancerequest/1" ,
 "status": "SCHEDULED
 "created" : "2013-05-09T07:06:57+0000"
}
```

### Example Response (After the instance has been created):

```
{
 "uri" : "/em/cloud/jaas/javaplatfromtemplate/D016AADD465462ACE040F20AD18167BA" ,
 "name" : "PS4_LowHeapTemplate" ,
 "description" : "This is a service template using the CDP CDP_PS4_lowHeap_MWProfile" ,
 "resource_state" : {
 "state" : "READY"
 } ,
 "context_id" : "D016AADD465462ACE040F20AD18167BA" ,
 "media_type" : "application/oracle.com.cloud.jaas.JavaPlatformTemplate+json" ,
 "service_family_type" : "jaas" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/javaplatfromtemplate/ D016AADD465462ACE040F20AD18167BA" ,
 "created" : "2012-12-04 23:16:10.0" ,
 "default_instance_media_type" :
"application/oracle.com.cloud.jaas.JavaPlatformInstance+json" ,
 "zones" : {
 "media_type" : "application/oracle.com.cloud.common.Zone+json" ,
 "total" : "1" ,
 "elements" :
 [
 {
 "uri" : "/em/cloud/jaas/zone/A1B44A4EBCC4563125D9D0A3AAE4FD51" ,
 "name" : "Zone1" ,
 "media_type" : "application/oracle.com.cloud.jaas.Zone+json" ,
 "service_family_type" : "jaas"
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/service_family_type/jaas"
 }
]
 }
 }
}
```

```

 } ,
 "service_instances" : {
 "media_type" : "application/oracle.com.cloud.common.ServiceInstance+json" ,
 "total" : "1" ,
 "elements" :
 [
 {
 "uri" :
 "/em/cloud/jaas/javaplatforminstance/8AF6E65447FBA84E968CCD909BD05D42" ,
 "name" : "MyServer-1" ,
 "media_type" :
 "application/oracle.com.cloud.jaas.JavaPlatformInstance+json" ,
 "status" : "RUNNING"
 "canonicalLink":
 "/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/javaplatform
 instance/8AF6E65447FBA84E968CCD909BD05D42"
 }
]
 }
}

```

### 28.1.5 JavaPlatformInstance

JavaPlatformInstance will extend the ServiceInstance from the Cloud framework resource model. It is linked to an Oracle Composite Application target in Enterprise Manager.

The following table describes the JavaPlatformInstance [application/oracle.com.cloud.jaas.JavaPlatformInstance+json] data model.

**Table 28–5 JavaPlatformInstance Data Model**

| Field                            | Type                                       | Occurs | Description                                                                                                                                                         |
|----------------------------------|--------------------------------------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| uri                              | URI                                        | 1      | A GET against this URI refreshes the client representation of the resources accessible to this user.                                                                |
| name                             | String                                     | 1      | A human readable name given to the instance.<br>[POST]                                                                                                              |
| zone                             | Zone (uri, name, media_type)               | 1      | The Zone to which the instance is deployed.<br>[POST]                                                                                                               |
| media_type                       | String                                     | 1      | Indicates the additional media type that clients can use to perform a GET.                                                                                          |
| application_instance_deployments | Collection <ApplicationInstanceDeployment> | 0...1  | This is a list of all the ApplicationInstanceDeployments deployed on the JavaPlatformInstance.                                                                      |
| resource_state                   | ResourceState                              | 1      | The validity of the fields on a GET should be guaranteed only when the resource state is READY. Otherwise, the client should not assume the validity of the fields. |
| based_on                         | JavaPlatformTemplate                       | 1      | JavaPlatformTemplate based on which this instance was created.                                                                                                      |

**Table 28–5 (Cont.) JavaPlatformInstance Data Model**

| Field               | Type                   | Occurs | Description                                                                                                             |
|---------------------|------------------------|--------|-------------------------------------------------------------------------------------------------------------------------|
| jdbc_drivers        | String[]               | 0...1  | List of JDBC drivers available to create a data source on this instance.                                                |
| data_sources        | Collection<DataSource> | 0...1  | This is a list of all the data sources configured for this instance.                                                    |
| status              | String                 | 0...1  | Indicates the status of the underlying Weblogic cluster<br>Values are<br>STOPPED, RUNNING, UNREACHABLE, UNKNOWN         |
| server_count        | String                 | 1      | Number of Java servers in the Weblogic cluster of the JavaPlatformInstance<br>[PUT]                                     |
| service_family_type | String                 | 1      | Denotes the type of Service Instance (for example, "jaas").                                                             |
| context_id          | String                 | 1      | This is an ID indication the internal reference of the service template                                                 |
| canonicalLink       | String                 | 1      | GET against this normalized/canonical URI refreshes the client representation of the resources accessible to this user. |

**28.1.5.1 GET Operation Supported on a JavaPlatformInstance**

The following shows a sample GET operation on the JavaPlatformInstance URI to return its representation.

**Example Request:**

```
GET /em/cloud/jaas/javaplatforminstance/960A4A8244145EC7949303B4ECD650D0
Host: cloudcompany.com
Authorization: Basic xxxxxxxx
Accept: application/oracle.com.cloud.jaas.JavaPlatformInstance+json
```

**Example Response:**

```
{
 "uri" : "/em/cloud/jaas/javaplatforminstance/960A4A8244145EC7949303B4ECD650D0" ,
 "name" : "jb_sd_platform_12_24_2012_13_38_28" ,
 "resource_state" : {
 "state" : "READY"
 } ,
 "context_id" : "960A4A8244145EC7949303B4ECD650D0" ,
 "media_type" : "application/oracle.com.cloud.jaas.JavaPlatformInstance+json" ,
 "service_family_type" : "jaas" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/javaplatforminstance/960A4A8244145EC7949303B4ECD650D0" ,
 "status" : "RUNNING" ,
 "zone" : {
 "uri" : "/em/cloud/jaas/zone/A1B44A4EBCC4563125D9D0A3AAE4FD51" ,
 "name" : "Zone1" ,
 "context_id" : "A1B44A4EBCC4563125D9D0A3AAE4FD51" ,
 "media_type" : "application/oracle.com.cloud.jaas.Zone+json" ,
 "service_family_type" : "jaas"
 }
 "canonicalLink" :
```

```

"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/zone/
A1B44A4EBCC4563125D9D0A3AAE4FD51" ,
 "type" : "jaas"
 } ,
 "based_on" : {
 "uri" :
"/em/cloud/jaas/javaplatformtemplate/D016AADD465462ACE040F20AD18167BA" ,
 "name" : "PS4_LowHeapTemplate" ,
 "context_id" : "D016AADD465462ACE040F20AD18167BA" ,
 "media_type" : "application/oracle.com.cloud.jaas.JavaPlatformTemplate+json"
 } ,
 "service_family_type" : "jaas"
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/javaplatform
template/ D016AADD465462ACE040F20AD18167BA "
 } ,
 "data_sources" : {
 "media_type" : "application/oracle.com.cloud.jaas.DataSource+json" ,
 "total" : "0" ,
 "elements" :
 [
]
 } ,
 "jdbc_drivers" :
[
 "com.ddtek.jdbc.oracle.OracleDriver" ,
 "oracle.jdbc.xa.client.OracleXADataSource" ,
 "oracle.jdbc.OracleDriver" ,
 "com.ddtek.jdbcx.oracle.OracleDataSource"
] ,
 "server_count" : "1" ,
 "application_instance_deployments" : {
 "media_type" :
"application/oracle.com.cloud.jaas.ApplicationInstanceDeployment+json" ,
 "total" : "0" ,
 "elements" :
 [
]
 }
}

```

### 28.1.5.2 DELETE Operation Supported on a JavaPlatformInstance

The following shows a sample DELETE operation on the JavaPlatformInstance URI to retire the service instance.

#### Example Request:

```

DELETE /em/cloud/jaas/javaplatforminstance/960A4A8244145EC7949303B4ECD650D0
Host: cloudcompany.com
Authorization: Basic xxxxxxxx
Accept: application/oracle.com.cloud.jaas.JavaPlatformInstance+json

```

#### Example Response:

```

{
 "uri" : "/em/cloud/jaas/javaplatforminstancerequest/121" ,
 "name" : "rtjul17_Delete_20120719_035504227" ,
 "resource_state" : {
 "state" : "INITIATED" ,
 "messages" :
 [
 {

```

```

 "text" : "null" ,
 "date" : "2012-07-19T09:16:13+0000"
 }
}
],
 "context_id": "121",
 "media_type": "application/oracle.com.cloud.jaas.JavaPlatformInstance+json" ,
 "service_family_type": "jaas" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/javaplatfor
instancerequest/121" ,
 "status": "SCHEDULED"
 "created" : "2013-05-09T09:08:38+0000"
}

```

### 28.1.5.3 GET of Java Platform Request URI returned by DELETE operation

#### Example Request:

```

GET /em/cloud/jaas/javaplatforinstancerequest/121
Host: cloudcompany.com
Authorization: Basic xxxxxxxx
Accept: application/oracle.com.cloud.jaas.JavaPlatformInstance+json

```

#### Example Response (When the operation is in progress):

```

{
 "uri" : "/em/cloud/jaas/javaplatforinstancerequest/121" ,
 "name" : "rtjul17_Delete_20120719_035504227" ,
 "resource_state" : {
 "state" : "INITIATED" ,
 "messages" :
 [
 {
 "text" : "null" ,
 "date" : "2012-07-19T09:16:13+0000"
 }
]
 }
} ,
 "context_id": "121",
 "media_type": "application/oracle.com.cloud.jaas.JavaPlatformInstance+json" ,
 "service_family_type": "jaas" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/javaplatfor
instancerequest/121" ,
 "status": "SCHEDULED"
 "created" : "2013-05-09T09:08:38+0000"
}

```

#### Example Response (when the operation is complete):

The Status field will show that the resource has been deleted.

### 28.1.5.4 POST Operations Supported on a JavaPlatformInstance

The following shows a sample POST operation on the JavaPlatformInstance URI to deploy an ApplicationInstanceDeployment.

- Name of ApplicationInstanceDeployment in the "name" attribute
- ApplicationInstanceComponent URI in the "application\_instance\_component" attribute

**Example Request:**

```
POST /em/cloud/jaas/javaplatforminstance/8AF6E65447FBA84E968CCD909BD05D42
Host: cloudcompany.com
Authorization: Basic xxxxxxxx
Accept: application/oracle.com.cloud.jaas.ApplicationInstanceDeployment+json
Content-Type: application/oracle.com.cloud.jaas.
ApplicationInstanceDeployment+json
{
 "name" : "rtriddappjul18",
 "application_instance_component" :
"/em/cloud/jaas/applicationinstancecomponent/oracle%3AdefaultService%3Aem%3Aprovis
ioning%3A1%3Acmp%3ACOMP_Component%3ASUB_
JavaEEApplication%3AC414FFB7A0912357E040F10A716015F9%3A0.1"
}
```

**Example Response:**

```
{
 "uri" : "/em/cloud/jaas/applicationinstancedeploymentrequest/ rtriddappjul18@131"
,
 "name" : " rtriddappjul18_Deploy_20120718_051319824" ,
 "resource_state" : {
 "state" : "INITIATED" ,
 "messages" :
 [
 {
 "date" : "2012-07-18T12:13:28+0000" ,
 "text" : "null"
 }
]
 } ,
 "context_id" : "rtriddappjul18@131" ,
 "media_type" :
"application/oracle.com.cloud.jaas.ApplicationInstanceDeployment+json" ,
 "service_family_type" : "jaas" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/applicati
ninstancedeploymentrequest/ rtriddappjul18@131" ,
 "status" : "SCHEDULED"
}
```

**28.1.5.5 GET of Application Instance Deployment Request**

The following shows a sample GET of an Application Instance Deployment Request returned by POST to create Application Instance.

**Example Request:**

```
GET /em/cloud/jaas/applicationinstancedeploymentrequest/rtriddappjul18@131
Host: cloudcompany.com
Authorization: Basic xxxxxxxx
Accept: application/oracle.com.cloud.jaas.JavaPlatformInstance+json
```

**Example Response (When the operation is in progress)**

```
{
 "uri" : "/em/cloud/jaas/applicationinstancedeploymentrequest/ rtriddappjul18@131"
,
 "name" : " rtriddappjul18_Deploy_20120718_051319824" ,
 "resource_state" : {
 "state" : "INITIATED" ,
```

```

 "messages" :
 [
 {
 "date" : "2012-07-18T12:13:28+0000" ,
 "text" : "null"
 }
]
 } ,
 "context_id" : "rtriddappjul18@131" ,
 "media_type" :
"application/oracle.com.cloud.jaas.ApplicationInstanceDeployment+json" ,
 "service_family_type" : "jaas" ,
 "status" : "SCHEDULED"
}

```

### Example Response (When the operation is complete)

```

{
 "uri" : "/em/cloud/jaas/applicationinstancedeployment/9D8868C9632E3" ,
 "name" : "rtriddappjul19" ,
 "resource_state" : {
 "state" : "READY"
 } ,
 "context_id" : "9D8868C9632E3" ,
 "media_type" :
"application/oracle.com.cloud.jaas.ApplicationInstanceDeployment+json" ,
 "service_family_type" : "jaas" ,
 "status" : "RUNNING" ,
 "http_application_invocation_url" : {
 "ms_3" : "http://slc02jyq.us.oracle.com:32463/RiddlesWeb" ,
 "ms_4" : "http://slc02jyq.us.oracle.com:40298/RiddlesWeb"
 } ,
 "https_application_invocation_url" : {
 "ms_3" : "https://slc02jyq.us.oracle.com:3550/RiddlesWeb" ,
 "ms_4" : "https://slc02jyq.us.oracle.com:41467/RiddlesWeb"
 } ,
 "contained_in" : {
 "uri" : "/em/cloud/jaas/javaplatforminstance/CF4339D922C" ,
 "name" : "rtjul17" ,
 "media_type" :
"application/oracle.com.cloud.jaas.JavaPlatformInstance+json" ,
 "status" : "RUNNING"
 }
}

```

#### 28.1.5.6 POST on the JavaPlatformInstance URI

The following shows a sample POST operation on the JavaPlatformInstance URI to create a DataSource on the JavaPlatformInstance

- Name of DataSource in the "name" attribute
- JNDI Name in "jndi\_name" attribute
- JDBC Driver in "jdbc\_driver" attribute
- Database connect string in "database\_connect\_string" attribute
- Database type in "database\_type" attribute
- username and password

**Example Request:**

```
POST /em/cloud/jaas/javaplatforminstance/8AF6E65447FBA84E968CCD909BD05D42
Host: cloudcompany.com
Authorization: Basic xxxxxxxx
Accept: application/oracle.com.cloud.jaas.DataSource+json
Content-Type: application/oracle.com.cloud.jaas.DataSource+json
{
 "name" : "tarunDS_wthParamchanges" ,
 "jndi_name" : ["jndi_1Paramchanges"] ,
 "jdbc_driver" : "oracle.jdbc.OracleDriver" ,
 "database_connect_string" :
"jdbc:oracle:thin:sysman/sysman@slc04jqf.us.oracle.com:15044:sasdec2" ,
 "database_type" : "Oracle" ,
 "username": "sysman",
 "password": "sysman"
}
```

**Example Response:**

```
{
 "uri" : "/em/cloud/jaas/datasourcerequest/rtrepodsjul23@181" ,
 "name" : "rtrepodsjul23_CreatedS_20120722_112705396" ,
 "resource_state" : {
 "state" : "INITIATED" ,
 "messages" :
 [
 {
 "date" : "2012-07-23T06:27:06+0000" ,
 "text" : "null"
 }
]
 },
 "context_id" : "rtrepodsjul23@181" ,
 "media_type" : "application/oracle.com.cloud.jaas.DataSource+json" ,
 "service_family_type" : "jaas" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/
datasourcerequest/rtrepodsjul23@181" ,
 "status" : "SCHEDULED"
}
```

**28.1.5.7 Get of Data Source Request**

The following shows a sample ET of Data Source request returned by POST to create Data Source

**Example Request:**

```
GET /em/cloud/jaas/datasourcerequest/rtrepodsjul23@181
Host: cloudcompany.com
Authorization: Basic xxxxxxxx
Accept: application/oracle.com.cloud.jaas.JavaPlatformInstance+json
```

**Example Response (When the operation is in progress):**

```
{
 "uri" : "/em/cloud/jaas/datasourcerequest/rtrepodsjul23@181" ,
 "name" : "rtrepodsjul23_CreatedS_20120722_112705396" ,
 "resource_state" : {
 "state" : "INITIATED" ,
 "messages" :
```



```

 [
 {
 "date" : "2012-07-23T06:27:06+0000" ,
 "text" : "null"
 }
]
 },
 "context_id" : "rtrepodsjul23@181" ,
 "media_type" : "application/oracle.com.cloud.jaas.DataSource+json" ,
 "service_family_type" : "jaas" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/
datasourcerequest/rtrepodsjul23@181" ,
 "status" : "SCHEDULED"
 }
 }

```

#### Example Response (When the operation is complete):

```

{
 "uri" : "/em/cloud/jaas/datasource/ rtrepodsjul18@CF4339D922C" ,
 "name" : "rtrepodsjul18" ,
 "resource_state" : {
 "state" : "READY"
 } ,
 "context_id" : " rtrepodsjul18@CF4339D922C" ,
 "media_type" : "application/oracle.com.cloud.jaas.DataSource+json" ,
 "service_family_type" : "jaas" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/datasource/r
trepodsjul18@CF4339D922C " ,
 "jndi_name" :
[
 "jndi_1"
] ,
 "jdbc_driver" : "oracle.jdbc.OracleDriver" ,
 "username" : "sysam",
 "database_connect_string" :
"jdbc:oracle:thin:sysman/sysman@slc02jyq.us.oracle.com:15044:sjuly14" ,
 "contained_in" : {
 "media_type" :
"application/oracle.com.cloud.jaas.JavaPlatformInstance+json" ,
 "context_id" : "CF4339D922C51" ,
 "name" : "rtjul17" ,
 "uri" : "/em/cloud/jaas/javaplatforminstance/CF4339D922C51"
 }
}

```

#### 28.1.5.8 PUT Operations Supported on a JavaPlatformInstance

The following shows a sample PUT operation on JavaPlatformInstance URI to scale up or scale down the service instance. The final server count is contained in "server\_count" attributes. A Scale up or scale down request will be submitted according to Current server count being greater than or lesser than the Final server count.

#### Example Request:

```

PUT /em/cloud/jaas/javaplatforminstance/CF4339D922C5
Host: cloudcompany.com
Authorization: Basic xxxxxxxx
Accept: application/oracle.com.cloud.jaas.JavaPlatformInstance+json

```

```
Content-type: application/oracle.com.cloud.jaas.JavaPlatformInstance+json
{
 "server_count" : "3"
}
```

### Example Response:

```
{
 "uri" : "/em/cloud/jaas/javaplatforminstancerequest/211" ,
 "name" : "rtjul17_Scaleup_20120719_045804697" ,
 "resource_state" : {
 "state" : "INITIATED" ,
 "messages" :
 [
 {
 "text" : "null" ,
 "date" : "2012-07-19T09:16:13+0000"
 }
]
 } ,
 "context_id": "211",
 "media_type": "application/oracle.com.cloud.jaas.JavaPlatformInstance+json" ,
 "service_family_type": "jaas" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/
javaplatforminstancerequest/211" ,
 "status": "SCHEDULED"
 "created" : "2013-05-09T11:46:27+0000"
}
```

## 28.1.6 ApplicationInstanceDeployment

This represents the deployment of an Application Instance in a JavaPlatformInstance.

The following table describes the ApplicationInstanceDeployment [application/oracle.com.cloud.jaas.ApplicationInstanceDeployment+json] data model.

**Table 28–6 ApplicationInstanceDeployment Data Model**

| Field                            | Type                                               | Occurs | Description                                                                                                           |
|----------------------------------|----------------------------------------------------|--------|-----------------------------------------------------------------------------------------------------------------------|
| uri                              | URI                                                | 1      | A GET against this URI refreshes the client representation of the resources accessible to this user.                  |
| name                             | String                                             | 1      | A human readable name given to the instance.<br>[POST]                                                                |
| application_instance_component   | ApplicationInstanceComponent<br>(uri,name,version) | 1      | The ApplicationInstanceComponent in the Software library which was used during deployment of this instance.<br>[POST] |
| http_application_invocation_url  | Map<String,String>                                 | 0...1  | Map of server name and HTTP URL of the deployed application. Will be populated after the instance has been created.   |
| https_application_invocation_url | Map<String,String>                                 | 0...1  | Map of server name and HTTPS URL of the deployed application. Will be populated after the instance has been created.  |

**Table 28–6 (Cont.) ApplicationInstanceDeployment Data Model**

| Field               | Type                               | Occurs | Description                                                                                                                                                         |
|---------------------|------------------------------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| contained_in        | JavaPlatformInstance<br>(uri,name) | 1      | JavaPlatformInstance on which the application is deployed.                                                                                                          |
| deployment_plan     | String                             | 0...1  | Vendor specific configuration parameters for this deployment in the prescribed XML format.<br>[POST]                                                                |
| status              | String                             | 0...1  | Indicates the status of the ApplicationInstanceDeployment.<br>Values are<br>STOPPED, RUNNING, UNREACHABLE, UNKNOWN<br>[PUT]                                         |
| metrics             | Collection<br><Metric>             | 0...1  | The list of metrics (uri, name, type current_value) that are pertinent to this service instance.                                                                    |
| resource_state      | ResourceState                      | 1      | The validity of the fields on a GET should be guaranteed only when the resource state is READY. Otherwise, the client should not assume the validity of the fields. |
| service_family_type | String                             | 1      | Denotes the type of ApplicationInstanceDeployment (for example, "jaas").                                                                                            |
| media_type          | String                             | 1      | Indicates the additional media type that clients can use to perform a GET.                                                                                          |
| context_id          | String                             | 1      | This is an ID indication of the internal reference of the application deployment.                                                                                   |
| canonicalLink       | String                             | 1      | GET against this normalized/canonical URI refreshes the client representation of the resources accessible to this user.                                             |

**28.1.6.1 GET Operation supported on an ApplicationInstanceDeployment**

The following shows a sample GET operation on the ApplicationInstanceDeployment URI to return its representation.

**Example Request:**

```
GET /em/cloud/jaas/applicationinstancedeployment/9D8868C9632E3
Host: cloudcompany.com
Authorization: Basic xxxxxxxx
Accept: application/oracle.com.cloud.jaas.ApplicationInstanceDeployment+json
```

**Example Response:**

```
HTTP/1.1 200 OK
Content-Type: application/oracle.com.cloud.jaas.
ApplicationInstanceDeployment+json
Content-Location: /em/cloud/jaas/applicationinstancedeployment/9D8868C9632E3
Cache-Control: public
Content-Length: nnn
{
 "uri" : "/em/cloud/jaas/applicationinstancedeployment/9D8868C9632E3" ,
```

```

 "name" : "rtriddappjul19" ,
 "resource_state" : {
 "state" : "READY"
 } ,
 "media_type" :
"application/oracle.com.cloud.jaas.ApplicationInstanceDeployment+json" ,
 "service_family_type" : "jaas" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/applicationi
nstancedeployment/9D8868C9632E3" ,
 "status" : "RUNNING" ,
 "http_application_invocation_url" : {
 "ms_3" : "http://slc02jyq.us.oracle.com:32463/RiddlesWeb" ,
 "ms_4" : "http://slc02jyq.us.oracle.com:40298/RiddlesWeb"
 } ,
 "https_application_invocation_url" : {
 "ms_3" : "https://slc02jyq.us.oracle.com:3550/RiddlesWeb" ,
 "ms_4" : "https://slc02jyq.us.oracle.com:41467/RiddlesWeb"
 } ,
 "contained_in" : {
 "uri" : "/em/cloud/jaas/javaplatforminstance/CF4339D922C" ,
 "name" : "rtj17" ,
 "media_type" :
"application/oracle.com.cloud.jaas.JavaPlatformInstance+json" ,
 "status" : "RUNNING"
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/javaplatfor
minstance/ CF4339D922C "
 }
}

```

### 28.1.6.2 DELETE Operation supported on an ApplicationInstanceDeployment

The following shows a sample DELETE operation on the ApplicationInstanceDeployment URI which undeploys the application from JavaPlatformInstance.

#### Example Request:

```

DELETE /em/cloud/jaas/applicationinstancedeployment/9D8868C9632E3
Host: cloudcompany.com
Authorization: Basic xxxxxxxx
Accept: application/oracle.com.cloud.jaas.ApplicationInstanceDeployment+json

```

#### Example Response:

```

HTTP/1.1 200 OK
Content-Type: application/oracle.com.cloud.jaas.
ApplicationInstanceDeployment+json
Content-Location:
/em/cloud/jaas/applicationinstancerequest/rtriddappjul19@184
Cache-Control: public
Content-Length: nnn
{
 "uri" :
"/em/cloud/jaas/applicationinstancerequest/rtriddappjul19@184" ,
 "name" : "rtriddappjul19_Undeploy_20120723_011757881" ,
 "resource_state" : {
 "state" : "INITIATED" ,
 "messages" :
[

```

```

 {
 "date" : "2012-07-23T08:18:01+0000" ,
 "text" : "null"
 }
]
} ,
"media_type" :
"application/oracle.com.cloud.jaas.ApplicationInstanceDeployment+json" ,
"service_family_type" : "jaas" ,
canonicalLink:
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/applicationi
nstancedeploymentrequest/rtriddappjul19@184" ,
"status" : "SCHEDULED"
}

```

### 28.1.6.3 Sample GET on Application Instance Deployment Request

The following shows a sample GET operation on an Application Instance Deployment Request.

#### Example Request:

```

GET /em/cloud/jaas/applicationinstancedeploymentrequest/rtriddappjul19@184
Host: cloudcompany.com
Authorization: Basic xxxxxxxx
Accept: application/oracle.com.cloud.jaas.ApplicationInstanceDeployment+json

```

#### Example Response: (operation in progress):

```

{
 "uri" :
"/em/cloud/jaas/applicationinstancedeploymentrequest/rtriddappjul19@184" ,
 "name" : "rtriddappjul19_Undeploy_20120723_011757881" ,
 "resource_state" : {
 "state" : "INITIATED" ,
 "messages" :
 [
 {
 "date" : "2012-07-23T08:18:01+0000" ,
 "text" : "null"
 }
]
 }
} ,
"context_id" : "rtriddappjul19@184" ,
"media_type" :
"application/oracle.com.cloud.jaas.ApplicationInstanceDeployment+json" ,
"service_family_type" : "jaas" ,
canonicalLink:
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/applicationi
nstancedeploymentrequest/rtriddappjul19@184" ,
"status" : "SCHEDULED"
}

```

#### Example Response: (operation is complete):

The status field will show that the resource has been deleted.

### 28.1.6.4 Sample PUT Operation supported on an ApplicationInstanceDeployment

The following shows a sample PUT operation on the ApplicationInstanceDeployment URI to start or stop the instance.

- STOPPED to RUNNING  
Put on status attribute, value = "START"
- RUNNING to STOPPED  
Put on status attribute, value = "STOP"
- REDEPLOY  
Put on status attribute, value = "REDEPLOY" and ApplicationInstanceComponent URI in the "application\_instance\_component" attribute

**Example Request: START/STOP**

```
PUT /em/cloud/jaas/applicationinstancedeployment/9D8868C9632E3
Host: cloudcompany.com
Authorization: Basic xxxxxxxx
Accept: application/oracle.com.cloud.jaas.ApplicationInstanceDeployment+json
Content-type:
application/oracle.com.cloud.jaas.ApplicationInstanceDeployment+json
{
 "status" : "STOP"
}
```

**Example Response:**

```
HTTP/1.1 200 OK
Content-Type: application/oracle.com.cloud.jaas.
ApplicationInstanceDeployment+json
Content-Location:
/em/cloud/jaas/applicationinstancedeploymentrequest/rtriddappjul19@182
Cache-Control: public
Content-Length: nnn
{
 "uri" :
"/em/cloud/jaas/applicationinstancedeploymentrequest/rtriddappjul19@182" ,
 "name" : "rtriddappjul19_Stop_20120723_123327380" ,
 "resource_state" : {
 "state" : "INITIATED" ,
 "messages" :
 [
 {
 "date" : "2012-07-23T07:33:27+0000" ,
 "text" : "null"
 }
]
 } ,
 "media_type" :
"application/oracle.com.cloud.jaas.ApplicationInstanceDeployment+json" ,
 "service_family_type" : "jaas" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/applicationi
nstedeploymentrequest/rtriddappjul19@182" ,
 "status" : "SCHEDULED"
}
```

**Example Request for REDEPLOY:**

```
PUT /em/cloud/jaas/applicationinstancedeployment/9D8868C9632E3
Host: cloudcompany.com
Authorization: Basic xxxxxxxx
Accept: application/oracle.com.cloud.jaas.ApplicationInstanceDeployment+json
```

```

Content-type:
application/oracle.com.cloud.jaas.ApplicationInstanceDeployment+json
{
 "status" : "STOP" ,
 "application_instance_component" :
"/em/cloud/jaas/applicationinstancecomponent/oracle%3AdefaultService%3Aem%3Aprovis
ioning%3A1%3Acmp%3ACOMP_Component%3ASUB_
JavaEEApplication%3AC414FFB7A0912357E040F10A716015F9%3A0.1"
}

```

### Example Response:

```

{
 "uri" :
"/em/cloud/jaas/applicationinstancedeploymentrequest/rtriddappjul19@183" ,
 "name" : "rtriddappjul19_Undeploy_20120723_123327380" ,
 "resource_state" : {
 "state" : "INITIATED" ,
 "messages" :
 [
 {
 "date" : "2012-07-23T07:33:27+0000" ,
 "text" : "null"
 }
]
 } ,
 "context_id" : "rtriddappjul19@183" ,
 "media_type" :
"application/oracle.com.cloud.jaas.ApplicationInstanceDeployment+json" ,
 "service_family_type" : "jaas" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/applicationi
nstedeploymentrequest/rtriddappjul19@183" ,
 "status" : "SCHEDULED"
}

```

## 28.1.7 DataSource

The DataSource resource model has the following characteristics:

The following table describes the DataSource

[application/oracle.com.cloud.jaas.DataSource+json] data model.

**Table 28–7 DataSource Data Model**

| Field     | Type     | Occurs | Description                                                                                                                                                       |
|-----------|----------|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| uri       | URI      | 1      | A GET against this URI refreshes the client representation of the resources accessible to this user.                                                              |
| name      | String   | 1      | A human readable name given to the data source.<br>[POST][PUT]                                                                                                    |
| jndi_name | String[] | 1      | JNDI path to where this data source will be bound. Applications look up the data source on the JNDI tree by this name when reserving a connection.<br>[POST][PUT] |

**Table 28–7 (Cont.) DataSource Data Model**

| Field                   | Type                                | Occurs | Description                                                                                                                                                         |
|-------------------------|-------------------------------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| database_type           | String                              | 0...1  | The DBMS of the database that the data source represents.<br>[POST]                                                                                                 |
| jdbc_driver             | String                              | 1      | The database driver used to connect to the database.<br>[POST]                                                                                                      |
| contained_in            | JavaPlatformInstance<br>(uri, name) | 1      | JavaPlatformInstance on which the data source exists.                                                                                                               |
| resource_state          | ResourceState                       | 1      | The validity of the fields on a GET should be guaranteed only when the resource state is READY. Otherwise, the client should not assume the validity of the fields. |
| database_connect_string | String                              | 0...1  | Connect String of the database can be provided to create the data source.<br>[POST][PUT]                                                                            |
| service_family_type     | String                              | 1      | Denotes the type of Data Source (for example, "jaas").                                                                                                              |
| media_type              | String                              | 1      | Indicates the additional media type that clients can use to perform a GET.                                                                                          |
| username                | String                              | 1      | Username of the DB connection.                                                                                                                                      |
| password                | String                              | 1      | Password of the DB connection.                                                                                                                                      |
| context_id              | String                              | 1      | This is an ID indication the internal reference of the data source.                                                                                                 |
| canonicalLink           | String                              | 1      | GET against this normalized/canonical URI refreshes the client representation of the resources accessible to this user.                                             |

**28.1.7.1 GET Operation supported on a DataSource**

The following shows a sample GET operation on the DataSource URI to return its representation.

**Example Request:**

```
GET /em/cloud/jaas/datasource/rtrepodsjul18@CF4339D922C
Host: cloudcompany.com
Authorization: Basic xxxxxxxx
Accept: application/oracle.com.cloud.jaas.DataSource+json
```

**Example Response:**

```
{
 "uri" : "/em/cloud/jaas/datasource/rtrepodsjul18@CF4339D922C" ,
 "name" : "rtrepodsjul18" ,
 "resource_state" : {
 "state" : "READY"
 } ,
 "context_id" : "rtrepodsjul18@CF4339D922C" ,
 "media_type" : "application/oracle.com.cloud.jaas.DataSource+json" ,
 "service_family_type" : "jaas" ,
 "canonicalLink" :
```



```

"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/datasource/r
trepodsjul18@CF4339D922C " ,
 "jndi_name" :
 [
 "jndi_1"
] ,
 "jdbc_driver" : "oracle.jdbc.OracleDriver" ,
 "username" : "sysam",
 "database_connect_string" :
"jdbc:oracle:thin:sysman/sysman@slc02jyq.us.oracle.com:15044:sjuly14" ,
 "contained_in" : {
 "media_type" :
"application/oracle.com.cloud.jaas.JavaPlatformInstance+json" ,
 "context_id" : "CF4339D922C51" ,
 "name" : "rtjul17" ,
 "uri" : "/em/cloud/jaas/javaplatforminstance/CF4339D922C51"
 }
}

```

### 28.1.7.2 DELETE Operation supported on a DataSource

The following shows a sample DELETE operation on the DataSource URI removes it from the JavaPlatformInstance.

#### Example Request:

```

DELETE /em/cloud/jaas/datasource/rtrepodsjul18@CF4339D922C
Host: cloudcompany.com
Authorization: Basic xxxxxxxx
Accept: application/oracle.com.cloud.jaas.DataSource+json

```

#### Example Response:

```

HTTP/1.1 200 OK
Content-Type: application/oracle.com.cloud.jaas.DataSource+json
Content-Location: /em/cloud/jaas/datasourcerequest/rtrepodsjul18@201
Cache-Control: public
Content-Length: nnn
{
 "uri" : "/em/cloud/jaas/datasourcerequest/rtrepodsjul18@201" ,
 "name" : "rtrepodsjul18_DeleteDS_20120723_024836692" ,
 "resource_state" : {
 "state" : "INITIATED" ,
 "messages" :
 [
 {
 "date" : "2012-07-23T09:48:37+0000" ,
 "text" : "null"
 }
]
 } ,
 "media_type" : "application/oracle.com.cloud.jaas.DataSource+json" ,
 "service_family_type" : "jaas" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/
datasourcerequest/rtrepodsjul18@201" ,
 "status" : "SCHEDULED"
}

```

### 28.1.7.3 GET on Data source Request

The following shows a sample GET operation on a DataSource request.

#### Example Request:

```
DELETE / em/cloud/jaas/datasourcerequest/rtrepodsjul18@201
Host: cloudcompany.com
Authorization: Basic xxxxxxxx
Accept: application/oracle.com.cloud.jaas.DataSource+json
```

#### Example Response: (operation in progress):

```
{
 "uri" : "/em/cloud/jaas/datasourcerequest/rtrepodsjul18@201" ,
 "name" : "rtrepodsjul18_DeleteDS_20120723_024836692" ,
 "resource_state" : {
 "state" : "INITIATED" ,
 "messages" :
 [
 {
 "date" : "2012-07-23T09:48:37+0000" ,
 "text" : "null"
 }
]
 } ,
 "context_id" : "rtrepodsjul18@201" ,
 "media_type" : "application/oracle.com.cloud.jaas.DataSource+json" ,
 "service_family_type" : "jaas" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/
datasourcerequest/rtrepodsjul18@201" ,
 "status" : "SCHEDULED"
}
```

#### Example Response: (operation is complete)

The status field will show that the resource has been deleted.

### 28.1.7.4 PUT Operation Supported on a DataSource

The following shows a sample PUT operation on the DataSource URI to update the DataSource to add/remove JNDI names and edit the database connect string. An example request for TEST DataSource is also provided.

#### Example Request:

```
PUT /em/cloud/jaas/datasource/rtrepodsjul18@CF4339D922C
Host: cloudcompany.com
Authorization: Basic xxxxxxxx
Accept: application/oracle.com.cloud.jaas.DataSource+json
Content-type: application/oracle.com.cloud.jaas.DataSource+json
{
 "jndi_name" : ["jndi_2"],
 "database_connect_string" :
"jdbc:oracle:thin:sysman/sysman@slc02knx.us.oracle.com:15044:sep25cc" ,
 "username" : "sysman",
 "password" : "sysman"
}
```

#### Example Response:

```
{
```

```

"uri" : "/em/cloud/jaas/datasourcerequest/rtrepodsjul18@CF4339D922C" ,
"name" : "rtrepodsjul18_UpdatedDS_20121216_062909975" ,
"resource_state" : {
 "state" : "READY" ,
 "messages" :
 [
 {
 "date" : "2012-12-16T14:29:16+0000" ,
 "text" : "null"
 }
]
} ,
"context_id" : "rtrepodsjul18@681" ,
"media_type" : "application/oracle.com.cloud.jaas.DataSource+json" ,
"service_family_type" : "jaas" ,
"canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/datasourcerequest/rtrepodsjul18@191" ,
"status" : "SCHEDULED"
}

```

### Example Request for TEST Data Source

```

PUT /em/cloud/jaas/datasource/rtrepodsjul18@CF4339D922C
Host: cloudcompany.com
Authorization: Basic xxxxxxxx
Accept: application/oracle.com.cloud.jaas.DataSource+json
Content-type: application/oracle.com.cloud.jaas.DataSource+json
{
 "status" : "TEST"
}

```

### Example Response

```

{
 "uri" : "/em/cloud/jaas/datasource/rtrepodsjul18@CF4339D922C" ,
 "name" : "rtrepodsjul18" ,
 "resource_state" : {
 "state" : "READY"
 } ,
 "context_id" : " rtrepodsjul18@CF4339D922C" ,
 "media_type" : "application/oracle.com.cloud.jaas.DataSource+json" ,
 "service_family_type" : "jaas" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/datasource/rtrepodsjul18@CF4339D922C " ,
 "jndi_name" :
 [
 "jndi_1"
] ,
 "jdbc_driver" : "oracle.jdbc.OracleDriver" ,
 "username" : "sysam",
 "database_connect_string" :
"jdbc:oracle:thin:sysman/sysman@slc02jyq.us.oracle.com:15044:sjuly14" ,
 "contained_in" : {
 "media_type" :
"application/oracle.com.cloud.jaas.JavaPlatformInstance+json" ,
 "context_id" : "CF4339D922C51" ,
 "name" : "rtjul17" ,
 "uri" : "/em/cloud/jaas/javaplatforminstance/CF4339D922C51"
 }
}

```

```
}
```

### 28.1.7.5 GET on Data source Request

The following shows a sample GET on DataSource request.

#### Example Request:

```
DELETE / em/cloud/jaas/datasourcerequest/rtrepodsjul18@201
Host: cloudcompany.com
Authorization: Basic xxxxxxxx
Accept: application/oracle.com.cloud.jaas.DataSource+json
```

#### Example Response (operation in progress):

```
{
 "uri" : "/em/cloud/jaas/datasourcerequest/rtrepodsjul18@201" ,
 "name" : "rtrepodsjul18_DeleteDS_20120723_024836692" ,
 "resource_state" : {
 "state" : "INITIATED" ,
 "messages" :
 [
 {
 "date" : "2012-07-23T09:48:37+0000" ,
 "text" : "null"
 }
]
 } ,
 "context_id" : "rtrepodsjul18@201" ,
 "media_type" : "application/oracle.com.cloud.jaas.DataSource+json" ,
 "service_family_type" : "jaas" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/
datasourcerequest/rtrepodsjul18@201" ,
 "status" : "SCHEDULED"
}
```

#### Example Response (operation is complete):

```
{
 "uri" : "/em/cloud/jaas/datasource/ rtrepodsjul18@CF4339D922C" ,
 "name" : "rtrepodsjul18" ,
 "resource_state" : {
 "state" : "READY"
 } ,
 "context_id" : " rtrepodsjul18@CF4339D922C" ,
 "media_type" : "application/oracle.com.cloud.jaas.DataSource+json" ,
 "service_family_type" : "jaas" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/datasource/r
trepodsjul18@CF4339D922C " ,
 "jndi_name" :
[
 "jndi_1"
] ,
 "jdbc_driver" : "oracle.jdbc.OracleDriver" ,
 "username" : "sysam",
 "database_connect_string" :
"jdbc:oracle:thin:sysman/sysman@slc02jyq.us.oracle.com:15044:sjuly14" ,
 "contained_in" : {
 "media_type" :
"application/oracle.com.cloud.jaas.JavaPlatformInstance+json" ,

```

```

 "context_id" : "CF4339D922C51" ,
 "name" : "rtjul17" ,
 "uri" : "/em/cloud/jaas/javaplatforminstance/CF4339D922C51"
 }
}

```

### 28.1.8 ApplicationInstanceComponent

This represents the application component in the software library. It can be created by POST-ing to the Jaas service type.

The following table describes the ApplicationInstanceComponent [application/oracle.com.cloud.jaas.ApplicationInstanceComponent+json] data model.

**Table 28–8 ApplicationInstanceComponent Data Model**

| Field               | Type          | Occurs | Description                                                                                                                                                         |
|---------------------|---------------|--------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| uri                 | URI           | 1      | A GET against this URI refreshes the client representation of the resources accessible to this user.                                                                |
| name                | String        | 1      | A human readable name given to the instance.<br>[POST]                                                                                                              |
| creator             | String        | 1      | Creator of the ApplicationInstanceComponent in the Software Library.                                                                                                |
| created             | Timestamp     | 1      | Date and time, in ISO 8601 format, when the instance was created.                                                                                                   |
| version             | String        | 1      | Version of the ApplicationInstanceComponent in the Software library.<br>[POST]                                                                                      |
| resource_state      | ResourceState | 1      | The validity of the fields on a GET should be guaranteed only when the resource state is READY. Otherwise, the client should not assume the validity of the fields. |
| context_id          | String        | 1      | This is an ID indication the internal reference of the ApplicationInstanceComponent.                                                                                |
| service_family_type | String        | 1      | Denotes the type of ApplicationInstanceComponent (for example, "jaas").                                                                                             |
| media_type          | String        | 1      | Indicates the additional media type that clients can use to perform a GET.                                                                                          |
| status              | String        | 1      | Status of the ApplicationInstanceComponent in the Software Library.                                                                                                 |
| maturity            | String        | 1      | Maturity of the ApplicationInstanceComponent in the Software Library.                                                                                               |
| componenttype       | String        | 1      | Component type of ApplicationInstanceComponent (for example, Java EE Application).                                                                                  |

**Table 28–8 (Cont.) ApplicationInstanceComponent Data Model**

| Field         | Type   | Occurs | Description                                                                                                             |
|---------------|--------|--------|-------------------------------------------------------------------------------------------------------------------------|
| canonicalLink | String | 1      | GET against this normalized/canonical URI refreshes the client representation of the resources accessible to this user. |

### 28.1.8.1 GET Operation supported on an ApplicationInstanceComponent

The following shows a sample GET operation on the ApplicationInstanceComponent URI to return its representation.

#### Example Request:

```
GET
/em/cloud/jaas/applicationinstancecomponent/oracle%3AdefaultService%3Aem%3Aprovisi
oning%3A1%3Acmp%3ACOMP_Component%3ASUB_
JavaEEApplication%3AC514D5CF9FC00CD7E040F10A716020D0%3A0.1
Host: cloudcompany.com
Authorization: Basic xxxxxxxx
Accept: application/oracle.com.cloud.jaas.ApplicationInstanceComponent+json
```

#### Example Response:

```
HTTP/1.1 200 OK
Content-Type: application/oracle.com.cloud.jaas.
ApplicationInstanceComponent+json
Content-Location:
/em/cloud/jaas/applicationinstancecomponent/oracle%3AdefaultService%3Aem%3Aprovisi
oning%3A1%3Acmp%3ACOMP_Component%3ASUB_
JavaEEApplication%3AC514D5CF9FC00CD7E040F10A716020D0%3A0.1
Cache-Control: public
Content-Length: nnn
{
 "uri" :
"/em/cloud/jaas/applicationinstancecomponent/oracle%3AdefaultService%3Aem%3Aprovis
ioning%3A1%3Acmp%3ACOMP_Component%3ASUB_
JavaEEApplication%3AC514D5CF9FC00CD7E040F10A716020D0%3A0.1" ,
 "name" : "Riddles" ,
 "resource_state" : {
 "state" : "READY"
 } ,
 "context_id" : "oracle:defaultService:em:provisioning:1:cmp:COMP_
Component:SUB_JavaEEApplication:C514D5CF9FC00CD7E040F10A716020D0:0.1" ,
 "media_type" :
"application/oracle.com.cloud.jaas.ApplicationInstanceComponent+json" ,
 "service_family_type" : "jaas" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/
applicationinstancecomponent/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp
%3ACOMP_Component%3ASUB_
JavaEEApplication%3AC514D5CF9FC00CD7E040F10A716020D0%3A0.1" ,
 "created" : "2012-07-17 22:31:33.0" ,
 "version" : "0.1" ,
 "creator" : "SSA_USER_1" ,
 "componenttype" : "Java EE Application" ,
 "status" : "Ready" ,
 "maturity" : "Untested"
}
```

### 28.1.8.2 DELETE Operation supported on an ApplicationInstanceComponent

The following shows a sample DELETE operation on the ApplicationInstanceComponent URI.

#### Example Request:

```
DELETE
/em/cloud/jaas/applicationinstancecomponent/oracle%3AdefaultService%3Aem%3Aprovisi
oning%3A1%3Acmp%3ACOMP_Component%3ASUB_
JavaEEApplication%3AC514D5CF9FC00CD7E040F10A716020D0%3A0.1
Host: cloudcompany.com
Authorization: Basic xxxxxxxx
Accept: application/oracle.com.cloud.jaas.ApplicationInstanceComponent+json
```

#### Example Response:

```
{
 "uri" :
"/em/cloud/jaas/applicationinstancecomponent/oracle%3AdefaultService%3Aem%3Aprovisi
oning%3A1%3Acmp%3ACOMP_Component%3ASUB_
JavaEEApplication%3AC514D5CF9FC00CD7E040F10A716020D0%3A0.1" ,
 "resource_state" : {
 "state" : "DESTROYED"
 } ,
 "context_id" : "oracle:defaultService:em:provisioning:1:cmp:COMP_
Component:SUB_JavaEEApplication:C514D5CF9FC00CD7E040F10A716020D0:0.1" ,
 "media_type" :
"application/oracle.com.cloud.jaas.ApplicationInstanceComponent+json" ,
 "canonicalLink" :
"/em/websvcs/restful/extws/cloudservices/service/v0/ssa/em/cloud/jaas/applicati
oninstancecomponent/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp%3ACOMP_
Component%3ASUB_JavaEEApplication%3AC514D5CF9FC00CD7E040F10A716020D0%3A0.1"
}
```

## 28.2 Application Component Filtering

Application component filtering is supported for JaaS zones, templates, and application components.

### 28.2.1 Service Template Finds

The filtering attributes supported are name, service\_template\_type, owner, and media\_type. The following are sample requests and responses to Service Template Finds:

#### Example Request

```
GET /em/cloud?filters={"filters":{"name":"RiddleApp"}}
Host: cloudcompany.com
Authorization: Basic xxxxxxxx
Accept: application/oracle.com.cloud.common.ServiceTemplateFinds+json
```

#### Example Response

```
{
 "uri" : "/em/cloud/finds/service_
template/bWVkaWFfdHlwZSUyNTNEYXBwbGljYXRpb241MjUyRm9yYWNSZS5jb20uY2xvdWQuY29tbW9uL
lNl%0AcnZpY2VUZWlwbGF0ZSUyRm5hbWU1MjUzRFJpZGRsZUFwcA%3D%3D" ,
 "name" : "Service Template Finds" ,
 "media_type" : "application/oracle.com.cloud.common.ServiceTemplateFinds+json" ,
```

```

 "description" : "This is a Resource Finds for 'Service Template Finds' with
media type 'application/oracle.com.cloud.common.ServiceTemplateFinds+json' " ,
 "resource_state" : {
 "state" : "READY"
 } ,
 "filters" : {
 "media_type" : "application/oracle.com.cloud.common.ServiceTemplate" ,
 "name" : "RiddleApp"
 } ,
 "finds" : {
 "media_type" : "application/oracle.com.cloud.common.ServiceTemplate" ,
 "total" : "1" ,
 "elements" :
 [
 {
 "uri" :
"/em/cloud/jaas/applicationinstancecomponent/oracle%3AdefaultService%3Aem%3Aprovis
ioning%3A1%3Acmp%3ACOMP_Component%3ASUB_
JavaEEApplication%3AD1103FE5F6877372E040F20AD1814215%3A0.1" ,
 "name" : "RiddleApp" ,
 "media_type" :
"application/oracle.com.cloud.jaas.ApplicationInstanceComponent+json" ,
 "service_family_type" : "jaas"
 }
]
 }
}

```

## 28.2.2 Zone Finds

The filtering attributes supported are name, owner, service\_family\_type, and media\_type.

### Example Zone Finds GET Request

```

GET /em/cloud?filters={"filters":{"name":"Zone1"}}
Host: cloudcompany.com
Authorization: Basic xxxxxxxx
Accept: application/oracle.com.cloud.common.ZoneFinds+json

```

### Example Zone Finds GET Response

```

{
 "uri" :
"/em/cloud/finds/zone/bWVkaWFFdHlwZSUyNTNEYXBwbGljYXRpb24lMjUyRm9yYWNSZS5jb20uY2xv
dWQuY29tbW9uLlpr%0AbmU1MkZuYW11JTl1M0Rab25lMQ%3D%3D" ,
 "name" : "Zone Finds" ,
 "media_type" : "application/oracle.com.cloud.common.ZoneFinds+json" ,
 "description" : "This is a Resource Finds for 'Zone Finds' with media type
'application/oracle.com.cloud.common.ZoneFinds+json' " ,
 "resource_state" : {
 "state" : "READY"
 } ,
 "filters" : {
 "media_type" : "application/oracle.com.cloud.common.Zone" ,
 "name" : "Zone1"
 } ,
 "finds" : {
 "media_type" : "*/*" ,
 "total" : "1" ,
 "elements" :

```



```
[
 {
 "uri" : "/em/cloud/jaas/zone/A1B44A4EBCC4563125D9D0A3AAE4FD51" ,
 "name" : "Zone1" ,
 "media_type" : "application/oracle.com.cloud.jaas.Zone+json" ,
 "service_family_type" : "jaas"
 }
]
```



---

## Chargeback and Metering EMCLI Verbs

---

This chapter describes the EMCLI verbs used for Chargeback and Metering and provides details on verb names, descriptions, parameters and syntax. Examples of verb output and commands are also provided. This chapter contains the following sections:

- [Using Custom Charge Items](#)
- [Retrieving Metering Data from Enterprise Manager](#)

### 29.1 Using Custom Charge Items

The following sections provide some sample commands along with descriptions and examples. These verbs are used in the following circumstances:

- You would like to meter/charge based on the metric extension/custom configuration that you have defined. Custom Charge Items can also be used for out of box metrics and target properties. Note that Chargeback only supports metric extensions without key columns.
- These Custom Charge Items, when created in Chargeback, allow you to meter or charge for the custom item like any existing charge items that are shipped along with a plug-in, such as CPU Utilization. Note that the EM\_CBA\_ADMIN role is required to access these verbs.

#### 29.1.1 list\_charge\_item\_candidates

This section describes the list\_charge\_item\_candidate data verb.

- **Description:** This verb lists the items that can be registered to Chargeback.
- **Verb Syntax**

```
emcli list_charge_item_candidates
 -target_type=<target type>
 -source_data_type=<metric|config|property>
 -target_name=<target name>
 -config_name=<config name>
 -config_data_source=<target name>
 -all
```

The following properties can be defined:

- target\_type: Required parameter. Name of target type. In Enterprise Manager release 12.1.0.2, you can specify the target\_type as **oracle\_databases**, **oracle\_vm\_guest**, **host**, **oracle\_pdb**, or **weblogic\_j2eeserver**.

- `source_data_type`: Valid values are metric, config, and property.
- `target_name`: If you specify a `target_name`, the metering /charge data will be retrieved only for this target. If you do not specify a valid `target_name`, or if the specified target has not been enabled in the given date-range, then no data is generated. If this parameter is not specified, **All** targets for the specified target-type will be included.
- `config_name`: Required parameter if `source_data_type=config`.
- `config_data_source`: Data Source of the configuration. Required parameter if `source_data_type=config`.
- `all`: Displays all items, including out of box metrics of target type. Without this option, only user defined Metric Extensions, and Custom Configurations will be displayed.

### Examples

```
emcli list_charge_item_candidates -target_type="oracle_database" -source_data_type="metric"
```

```
emcli list_charge_item_candidates -target_type="oracle_database" -source_data_type="config" -target_name="myDatabase" -config_name="myCustomCCS" -config_data_source="ccsfile"
```

## 29.1.2 create\_charge\_item

This section describes the `create_charge_item` data verb.

- **Description:** This creates a charge item for Chargeback.
- **Verb Syntax**

```
emcli create_charge_item -input_file="property_file:<full path of property file>"
```

The option `[-input_file]` is the full path of file and contains the item properties. The following properties can be defined in the file:

- `target_type`: Required parameter. Name of target type. Supported target types for this release are `oracle_database`, `oracle_vm_guest`, `host`, `oracle_pdb`, and `weblogic_j2eeserver`.
- `source_data_type`: Required parameter. Type of source data. Valid values are metric, config, and property.
- `item_name`: Required parameter. The name of the item.
- `metric_group`: Metric group name as listed in `list_item_candidates`. This is a required property if `source_data_type=metric`.
- `metric_column`: Metric column name as listed in `list_item_candidates`. Required property if `source_data_type=metric`.
- `config_view`: Config view name as listed in `list_item_candidates`. Required property if `source_data_type=config`.
- `config_key`: Config key name as listed in `list_item_candidates`. Required property if `source_data_type=config`.
- `config_column`: Config column name as listed in `list_item_candidates`. Required property if `source_data_type=config`.

- `config_data_source`: Data source of configuration metric. Required property if `source_data_type=config`.
- `property`: Property name as listed in `list_item_candidates`. Required property if `source_data_type=property`.
- `item_displayname`: Required property. Display name of item
- `unit`: Display name of unit
- `aggregation_type`: Type of aggregation to be used for this item, only applicable for numeric data type. Valid values are `sum` and `avg`. Default value is `avg`.
- `is_config_condition`: Is item will be used for conditions in Chargeplan. Valid values are 0, 1. Default value is 0.
- `item_category`: Category of item. Default value is `instance`. Valid values are `cpu`, `storage`, `memory`, `network`, and `instance`.
- `data_type`: Type of data. Default value is `number`. Valid values are `string` and `number`.

### Examples

```
emcli create_charge_item -input_file="property_file:/home/user/property_file"
```

Contents of `/home/user/property_file`:

```
target_type=host
source_data_type=metric
item_name=total_proc
metric_group=Load
metric_column=noOfProcs
item_displayname=Total Processes
unit=process
aggregation_type=avg
item_category=instance
data_type=number
```

```
emcli create_charge_item -input_file="property_file:/home/user/property_file"
```

Contents of `/home/user/property_file`:

```
target_type=oracle_database
source_data_type=config
item_name=custom_config
config_view=myCustomCCS
config_key=region
config_column=country
config_data_source=regionList.txt
item_display_name=Region of Instance
item_category=instance
data_type=string
```

### 29.1.3 delete\_charge\_item

This section describes the `delete_charge_item` data verb.

- **Description:** This deletes the custom item from Chargeback
- **Verb Syntax**

```
emcli delete_charge_item -target_type=<target type> -item_name=<Item name>
```

The following properties can be defined in the file:

- target\_type is the name of target type and is a required parameter.. Supported target types for this release are oracle\_database, oracle\_vm\_guest, host, oracle\_pdb, and weblogic\_j2eeserver.
- item\_name is the name of the item and is a required parameter.

Examples

```
emcli delete_charge_item -target_type="oracle_database" -item_name="SampleMetricExtension"

emcli delete_charge_item -target_type="host" -item_name="SampleCustomConfig"
```

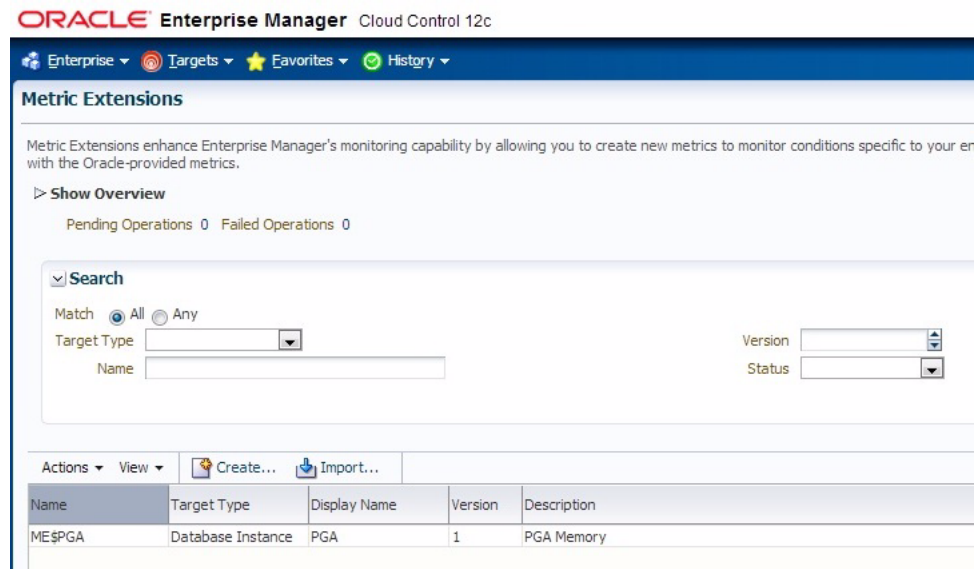
29.2 Examples of Custom Charge Item Verb Use

The following sections provide end-to-end examples in the correct use of Custom Charge Item verbs for custom configurations.

29.2.1 Metric Extension (PGA)


In this example, a metric extension "PGA" with a metric column named "PGA Memory" is created for a database instance target to collect PGA size.

Figure 29–1 PGA Memory



1. Log in to EMCLI and use the Chargeback EMCLI verb list\_charge\_item\_candidates to list the items that can be registered to Chargeback.

Figure 29–2 list\_charge\_item\_candidates





**Figure 29–5 Create Plan: Custom Plan**

Chargeback

Create Plan: Custom Plan

Charge Plan: Custom Plan

Effective Date: May 1, 2013 - Onward

Target Types

+ Add... - Remove

Database Instance

Database Instance

Database Instance

+ Add Item... Edit Item... - Remove Item Setup Configurations...

| Item                       | Default Configuration |
|----------------------------|-----------------------|
|                            | Charge                |
| Universal Rate Adjustments |                       |
| CPU Rate Factor            | 1x                    |
| Memory Rate Factor         | 1x                    |
| Storage Rate Factor        | 1x                    |

Add Item

Item Name: PGA Memory

Charge Type: Usage

Unit: GB

As the aggregation\_type is set to "avg", the charge rate is related to a time period (for example, per GB per timeperiod).

**Figure 29–6 Confirmation**

Chargeback

Confirmation

Custom Plan has been saved successfully.

Home Charge Plans Cost Centers Targets Reports

Create Set Rates... Delete...

Charge Plans

View

- Universal Charge Plan
- Custom Plan
  - May 1, 2013 - Onward
  - Database Instance
  - Sample Charge Plan

Custom Plan: May 1, 2013 - Onward: Database Instance

| Item                       | Default Configuration |
|----------------------------|-----------------------|
|                            | Charge                |
| PGA Memory                 | \$10.00 / GB / Day    |
| Universal Rate Adjustments |                       |
| CPU Rate Factor            | 1x                    |
| Memory Rate Factor         | 1x                    |
| Storage Rate Factor        | 1x                    |

### 29.2.2 Metric Extension (Application A Orders)

A metric extension named "Application A Orders" with a metric column called "Orders Processed" is created for a database instance target. This metric collects the number of orders processed by Application A.



Figure 29–7 Metric Extensions

Metric Extensions

Metric Extensions enhance Enterprise Manager's monitoring capability by allowing you to create new metrics to monitor conditions specific to your environment with the Oracle-provided metrics.

Show Overview

Pending Operations 0Failed Operations 0

Search

Match

☒ All☐ Any

Target Type

Name

Version

Status

Actions

View

Create...

Import...

| Name      | Target Type       | Display Name         | Version | Description                        |
|-----------|-------------------|----------------------|---------|------------------------------------|
| ME\$Order | Database Instance | Application A Orders | 1       | Orders Processed for Application A |
| ME\$PGA   | Database Instance | PGA                  | 1       | PGA Memory                         |

1. Log in to EMCLI and use the Chargeback EMCLI verb `list_charge_item_candidates` to list the items that can be registered to Chargeback.

Figure 29–8 list\_charge\_item\_candidates

xterm

```
$emcli list_charge_item_candidates --source_data_type=metric --target_type=oracle_database
```

| Metric Group | Metric Group | Display Name         | Metric Column   | Metric Column | Display Name |
|--------------|--------------|----------------------|-----------------|---------------|--------------|
| ME\$Order    |              | Application A Orders | OrdersProcessed | Orders        | Processed    |

```
$
```

2. Create a property file `prop_file2` to generate a custom charge item for the "Orders Processed" metric column.

For this metric, the charge rate will be based on total number of orders processed, with `aggregation_type` set to "sum".

Figure 29–9 Set Aggregation Type

xterm

```
target_type=oracle_database
item_displayname=Orders Processed
source_data_type=metric
metric_group=ME$Order
metric_column=OrdersProcessed
data_type=number
unit=Order
aggregation_type=sum
item_name=charge_orders

```

```
"prop_file2" 10L, 212C
```

3. Create a custom item in Chargeback using the `create_charge_item` EMCLI verb.

**Figure 29–10** *create-charge\_item*

```
xterm
$emcli create_charge_item -input_file="property_file:/tmp/prop_file2"
Custom item "Orders Processed" is successfully created in Chargeback.
$
```

4. Once the custom charge item is created, the Charge Plans page shows the new custom item. Rates can be set against a custom item similar to those set for out-of-box charge items.

**Figure 29–11** *Set Rates: Custom Plan*

Chargeback

Set Rates: Custom Plan

Charge Plan: Custom Plan  
Effective Date: May 1, 2013 - Onward

**Warning**  
Changing rates for the current report cycle will compute charges based on the n

**Target Types**  
+ Add... - Remove  
Database Instance

**Database Instance**  
+ Add Item... Edit Item... - Remove Item Setup Configurations...

| Item                       | Default Configuration | Charge         |
|----------------------------|-----------------------|----------------|
| PGA Memory                 |                       | \$ 10.00 / Day |
| Universal Rate Adjustments |                       |                |
| CPU Rate Factor            | 1x                    |                |
| Memory Rate Factor         | 1x                    |                |
| Storage Rate Factor        | 1x                    |                |

**Add Item**  
Item Name: Orders Proces  
Charge Type: Usage  
Unit: Order

As the `aggregation_type` is set to "sum", the charge rate does not depend upon time periods but rather on the total value of "Orders Processed".

**Figure 29–12** *Confirmation*

Chargeback

**Confirmation**  
Custom Plan has been saved successfully.

Home Charge Plans Cost Centers Targets Reports

Create Set Rates... Delete...

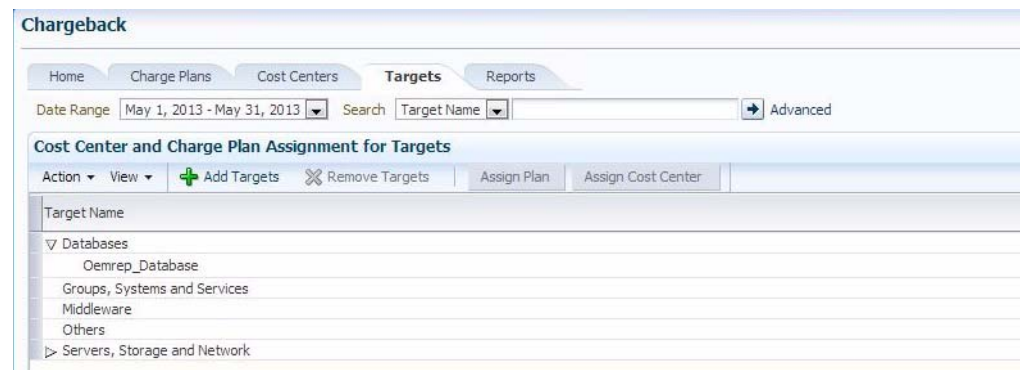
**Charge Plans**  
View  
Universal Charge Plan  
Custom Plan  
  May 1, 2013 - Onward  
    Database Instance  
Sample Charge Plan

**Custom Plan: May 1, 2013 - Onward: Database Instance**

| Item                       | Default Configuration | Charge             |
|----------------------------|-----------------------|--------------------|
| Orders Processed           |                       | \$2.00 / Order     |
| PGA Memory                 |                       | \$10.00 / GB / Day |
| Universal Rate Adjustments |                       |                    |
| CPU Rate Factor            | 1x                    |                    |
| Memory Rate Factor         | 1x                    |                    |
| Storage Rate Factor        | 1x                    |                    |

5. Add the database target to Chargeback and assign the "Custom Plan".

**Figure 29–13 Cost Center and Charge Plan Assignment**



6. Following the daily ETL run, the charge/usage is calculated against new custom items. The charge/usage of custom items can be viewed from the Chargeback report page.

**Figure 29–14 Details**

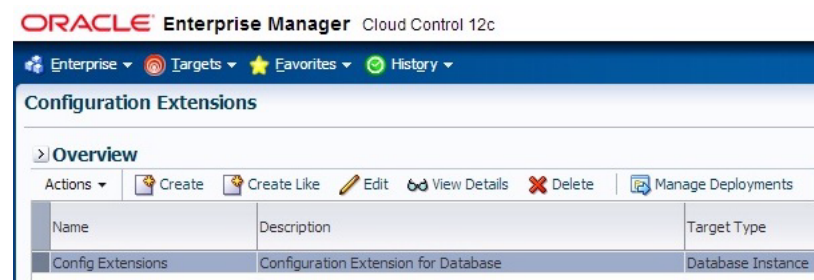
The screenshot shows the 'Details' tab in the Chargeback interface. It includes a 'View' dropdown, a 'Detail Level' dropdown set to 'All', and an 'Export' button. Below these is a table with the following data:

| Date         | Cost Center         | Target Type       | Target Name     | Plan Name   | Plan Configuration |
|--------------|---------------------|-------------------|-----------------|-------------|--------------------|
| May 22, 2013 | Default Cost Center | Database Instance | Oemrep_Database | Custom Plan | Default            |
| May 23, 2013 | Default Cost Center | Database Instance | Oemrep_Database | Custom Plan | Default            |
| May 24, 2013 | Default Cost Center | Database Instance | Oemrep_Database | Custom Plan | Default            |
| May 23, 2013 | Default Cost Center | Database Instance | Oemrep_Database | Custom Plan | Default            |
| May 24, 2013 | Default Cost Center | Database Instance | Oemrep_Database | Custom Plan | Default            |
| May 21, 2013 | Default Cost Center | Database Instance | Oemrep_Database | Custom Plan | Default            |

### 29.2.3 Configuration Extensions

A configuration extension named "Config Extensions" is created for the database instance target. This collects the number of high available tablespaces and the name of each high available tablespace as shown below. In this example, the high available tablespaces value is used as the basis of the new custom charge item.

**Figure 29–15 Configuration Extensions**



1. Log in to EMCLI and use the Chargeback EMCLI verb `list_charge_item_candidates` to list the configuration items that can be registered to Chargeback.

Note that `config_name` is the name of the configuration extension shown on the UI ("Config Extensions" in this case). Similarly, `config_data_source` is the name of the config column shown on the UI (in this case "tmp/dbconfig.xml").

**Figure 29–16** *list\_charge\_item\_candidates*

```

xterm
$emcli list_charge_item_candidates -source_data_type=config -target_type=oracle_database -config_name="Co
g_data_source="tmp/dbconfig.xml" -target_name="Oemrep_Database"
Custom Config Name Custom Config Key Custom Config Column

Config Extensions tablespaces/numberHATablespace count
Config Extensions tablespaces/haTablespace name
Config Extensions tablespaces/haTablespace{2} name
$

```

2. Create a property file `prop_file1` to generate a custom charge item for the "Number of High Available Tablespaces" metric column.

Note that the `config_key` value comes from the output of the "Custom Config Key" column in the output of the `list_charge_item_candidates` verb described previously. Similarly, the value of the `config_view` should be the value of "Custom Config Name" above.

**Figure 29–17** *Number of High Available Tablespaces*

```

xterm
target_type=oracle_database
item_displayname=Number of High Available Tablespaces
source_data_type=config
config_data_source=tmp/dbconfig.xml
config_view=Config Extensions
config_key=tablespaces/numberHATablespace
config_column=count
target_name=Oemrep_Database
data_type=number
unit=tablespace
item_name=config_ext_ha_tblspcs
~
~

```

3. Create a custom item in Chargeback using the `create_charge_item` EMCLI verb.

**Figure 29–18** *create\_charge\_item*

xterm
\$vi prop\_file1
\$emcli create\_charge\_item -input\_file="property\_file:/tmp/prop\_file1"
Custom item "Number of High Available Tablespaces" is successfully created in Chargeback.
\$

4. Once the custom charge item is created, the Charge Plans page shows the new custom item. Rates can be set against a custom item similar to those set for regular out-of-box charge items.

Figure 29–19 Set Rates: New Plan

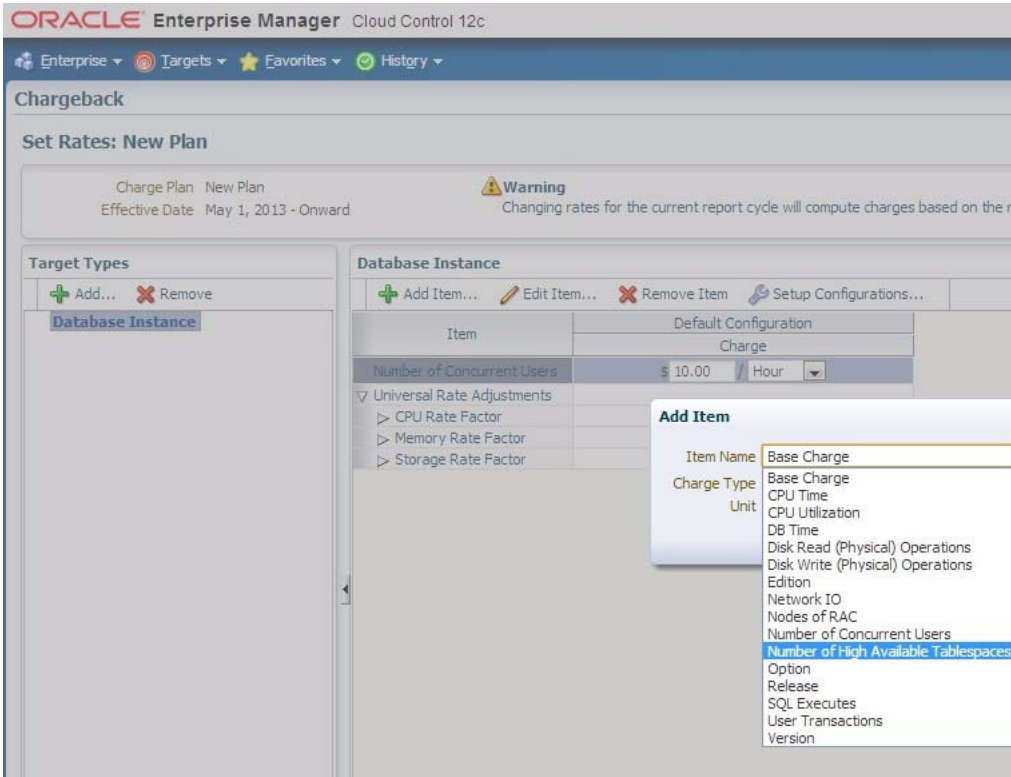
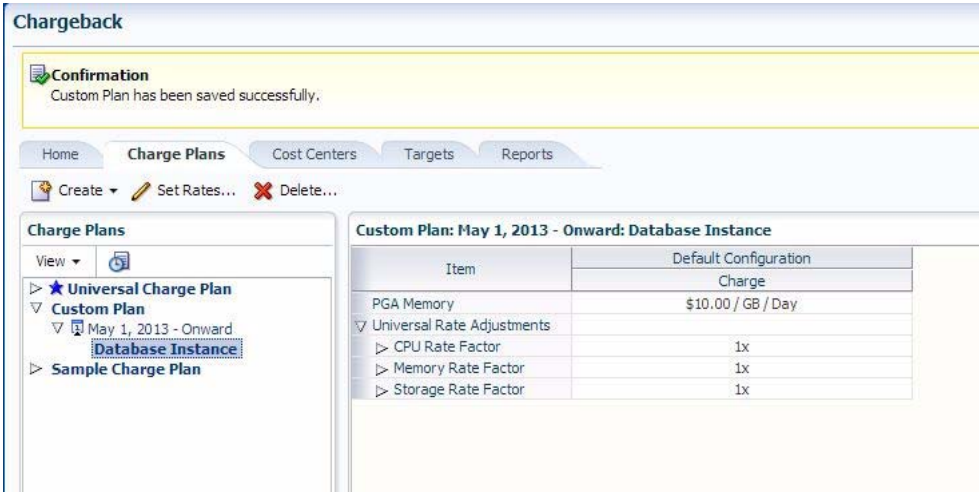
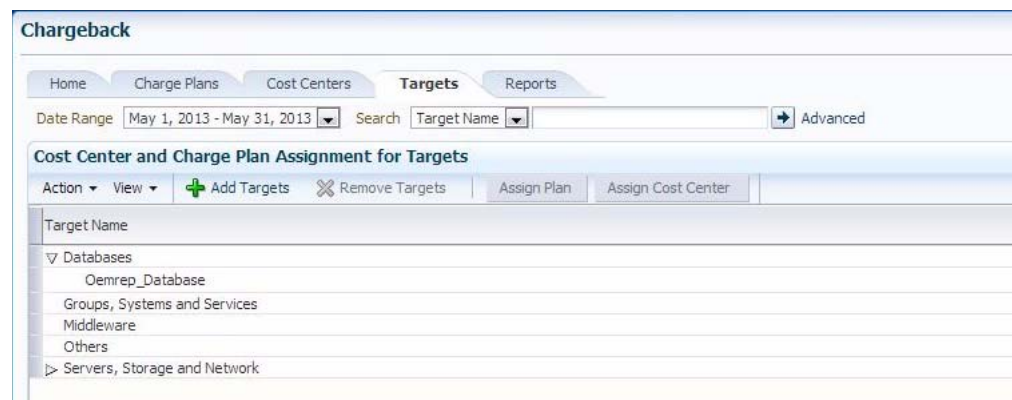


Figure 29–20 Confirmation



5. Add the database target to Chargeback and assign the "New Plan".

**Figure 29–21 New Plan**

6. After the daily ETL run, the charge/usage is calculated against new custom items. The charge/usage of custom items can be viewed from the Chargeback report page.

## 29.3 Retrieving Metering Data from Enterprise Manager

The Chargeback APIs can be used to integrate Enterprise Manager with Billing and Revenue Management systems (including Oracle BRM) and also to provide flexible chargeback mechanisms. Some examples include:

- Integration with a billing system with different rate plans for small, medium, and large configurations.
- To calculate chargeback for a flexible time period. For example, months starting from 15th and ending on the 15th.
- To rationalize chargeback based on other attributes that are not captured by Enterprise Manager.
- To adjust or round up chargeback based on usage. For example charge is calculated for the full-day even if the usage is for the partial day.

This section describes the process of retrieving metering data using the `get_metering_data` verb.

### `get_metering_data`

- **Description:** This verb generates comma-separated output with each row or line containing usage (and optionally charge) information for the specified parameters.
- **Verb Syntax:**

```
emcli get_metering_data
[-start_date=<start date in mmddyyyy> [-end_date=<end_date in mmddyyyy>]]
[-target_type=<target type> [-target_name= <target name>]]
[-cost_center=<cost center name>]
[-charge]
```

All the parameters of this verb are optional. If no parameters are specified, the following default values will be used:

- Start and End Date: The start and end date of the current report cycle.
- Target Type: The default value for this parameter is 'All'.



- Cost Center: The name of the user who has logged in.
- Charge: If this option is not specified, the metering data will be retrieved.

The parameters are described below:

- start\_date: The value for this parameter must be specified in **mmddyyyy** format. If you specify this parameter, the output metering/charge information will be filtered accordingly. If the start\_date is not specified, the start date for current report cycle will be used.

In this scenario the start\_date means midnight on the start date.

- end\_date: This parameter must be used along with the start\_date parameter. The value for this parameter must be given in **mmddyyyy** format. If you specify this parameter, the output metering/charge information will be filtered accordingly. If the end\_date is not specified, the end date for current report cycle will be used.

In this scenario the end\_date means midnight on the end date.

- target\_type: In Enterprise Manager release 12.1.0.2, you can specify the target\_type as **oracle\_databases**, **oracle\_vm\_guest**, **host**, or **weblogic\_j2eeserver**. If the targets of the specified target\_type have been enabled within the specified date range, appropriate metering or charge data (data for all targets of the specified target\_type) will be retrieved. If this parameter is not specified, **All** target types will be included.
- target\_name: If you specify a target\_name, the metering /charge data will be retrieved only for this target. If you do not specify a valid target\_name, or if the specified target has not been enabled in the given date-range, then no data is generated. If this parameter is not specified, **All** targets for the specified target-type will be included.
- cost\_center: If specified, the value of this parameter must be the same as the internal cost-center (as displayed on Cost-Center tab in the Chargeback application). If the specified value is a valid cost-center in the given date range, appropriate metering/charge data will be retrieved. If not, then no data will be generated. If this option is not specified, the default value for the parameter is considered as the logged in user. To retrieve metering/charge information for all cost-centers, specify "All Users" in quotes. Each user is also a consumer (for example, cost-center) in Chargeback. Dependent on their privileges, the logged-in user is shown relevant metering/charge information across different targets.

---

**Note:** The privileges required to view target information in Chargeback are:

- VIEW\_CAT\_TARGET: Allows you to view information for a specific target (active or inactive).
  - VIEW\_ANY\_CAT\_TARGET: Allows you to view information for any Chargeback target (active or inactive).
- 

- -charge: If this parameter is not specified, only the metering data will be retrieved. If this option is specified, both metering and charge information will be retrieved.

### 29.3.1 get\_metering\_data Output

This verb generates the following output:

- **CONSUMER\_NAME**: Name of the cost-center.
- **TARGET\_TYPE**: Type of target.
- **TARGET\_NAME**: Name of the target.
- **ITEM\_TYPE**: Type of the item or metric. Possible values are: **config**, **fixed**, **metric**, **property**, and **usage**.
- **CATEGORY\_NAME**: Possible values are: **cpu**, **memory**, **storage**, **activity**, **instance**, **network**, **service**, and **software**.
- **ENTITY\_NAME**: Name of the shared entity (valid when the target is enabled in Chargeback in shared mode).
- **ITEM\_DISPLAY\_NAME**: Display name of the item or metric (in English language). **Note**: Translation support is currently not available.
- **VALUE\_AVERAGE**: Average value for the metric on the given date (valid for numeric metrics).
- **STRING\_VALUE**: Value for the metric data (valid for string-based metrics).
- **DATA\_TYPE**: Data type of the metric. This can be **string** or **number**.
- **UNIT**: Unit of the metric data (for example, req and GB).
- **COLLECTION\_DATE**: Date on which the data is collected from the Enterprise Manager metric or configuration tables.
- **PLAN\_NAME**: Name of the Charge Plan associated with the particular target.
- **CHARGE**: Charge value for the specific metric of the target on the particular date.
- **DEFINED\_RATE**: Charge rate defined in the charge plan associated with the target.
- **RATE\_TYPE**: Type of the rate. Possible values are: **No value (blank)**, **Flat**, **config**, and **usage**.
- **RATE\_FACTOR**: Adjustment rate for universal metrics, as defined in the associated extended Chargeplan.
- **CHARGE\_RATE\_UNIT**: Unit, corresponding to the metric or item, as defined in the Charge Plan for the specific rate. For example, for an item or metric named **CPU Utilization (%)** or **CPU Utilization (%) per service**, the value will be "CPU". For a metric or item named **Base Charge**, the value will be "instance", a metric named 'User Requests', the value will be "req"). If the rate defined in Chargeplan is \$1 per MB per day, the value for this column will be "MB" .
- **CHARGE\_TIME\_UNIT**: Time unit as defined in the Charge Plan for the specific rate. (such as, hourly, daily, weekly, monthly, yearly). For example, if the rate defined in Charge Plan is \$1 per MB per Day, value for this column will be *daily*.

#### Example: Sample Output

```
"CONSUMER_NAME" , "TARGET_TYPE" , "TARGET_NAME" , "ITEM_TYPE" , "CATEGORY_NAME" , "ENTITY_NAME" , "ITEM_DISPLAY_NAME" , "STRING_VALUE" , "PLAN_NAME" , "COLLECTION_DATE" , "CHARGE" , "VALUE_AVERAGE" , "DEFINED_RATE" , "RATE_TYPE" , "RATE_FACTOR" , "CHARGE_RATE_UNIT" , "CHARGE_TIME_UNIT" , "NORMALIZE_RATE" , "ADJUST_RATE" , "DATA_TYPE" , "UNIT"
```

```
"linbo" , "oracle_vm_guest" , "mySite/myWls/AdminServer:agent_"
```



```

push","fixed","instance"," ","Base Charge","","zone_plan",15-OCT-11
00:00:00,72,1,"3","flat",1,"instance","hourly",1,1,"number",""
"linbo","oracle_vm_guest","mySite/myWls/AdminServer:agent_
push","fixed","instance"," ","Base Charge","","zone_plan",16-OCT-11
00:00:00,72,1,"3","flat",1,"instance","hourly",1,1,"number",""
"linbo","oracle_vm_guest","mySite/myWls/AdminServer:agent_
push","fixed","instance"," ","Base Charge","","zone_plan",17-OCT-11
00:00:00,72,1,"3","flat",1,"instance","hourly",1,1,"number",""
"linbo","oracle_vm_guest","mySite/myWls/AdminServer:agent_
push","fixed","instance"," ","Base Charge","","zone_plan",18-OCT-11
00:00:00,72,1,"3","flat",1,"instance","hourly",1,1,"number",""
"linbo","oracle_vm_guest","mySite/myWls/Cluster-0_
vm0:assembly1","fixed","instance"," ","Base Charge","","zone_plan",15-OCT-11
00:00:00,72,1,"3","flat",1,"instance","hourly",1,1,"number",""

```

### 29.3.1.1 Examples

A few examples are shown in the following section:

- `emcli get_metering_data`: Returns metering information for all targets (active or enabled in Chargeback) for the current report cycle for the logged in user.
- `emcli get_metering_data -charge`: Returns metering and charge information for all targets (active or enabled in Chargeback) for the current report cycle for the logged in user.
- `emcli get_metering_data -start_date=01202011-cost_center=ORG1`: Returns metering information for all targets (active or enabled in Chargeback) starting from 20th January 2011 until the end of the month for the **ORG1** cost-center.
- `emcli get_metering_data -start_date=01152011 -end_date=02152011 -target_type=oracle_database`: Returns metering information for all Oracle DB targets (active or enabled in Chargeback) that are owned by the logged in user, starting from 15th January 2011 until the 15th February 2011.
- `emcli get_metering_data -target_type=host target_name=my_host -cost_center=organization1`: Returns metering and charge information for "my\_host" target (of type *host*) for the current report cycle for "organization1" cost-center.
- `emcli get_metering_data -cost_center="All Users"`: Returns metering and charge information for all targets (active or enabled in Chargeback) in the current report cycle for all cost centers.
- Log in as `cba_admin_user` (who is the Chargeback Administrator) and enter the following commands:
  - `emcli get_metering_data`: Returns metering information for all targets (active or enabled in Chargeback) for the current report cycle for the `cba_admin_user`.
  - `emcli get_metering_data -cost_center=ssa_user1`: Returns metering information for all targets (active or enabled in Chargeback) that are owned by the `ssa_user1` in the cost center in the current report cycle.
- If the value specified for the cost center, target type, and / or target name is incorrect, no data is generated. For example, the following commands will not generate data:
  - `emcli get_metering_data -target_type=unknown`
  - `emcli get_metering_data -target_type=oracle_database -target_name=unknown_target`

**Note:** There is no target with the name **unknown\_target** configured in Enterprise Manager.

---

## SSA Administration APIs

This chapter provides details of the Cloud (RESTful) APIs for SSA Administration functions. In this release, APIs for PaasZone Management and SoftwarePool Management have been implemented. Other APIs used to fetch the API catalog, API output, and metadata and expansion of a resource are also described.

This chapter also provides details on the Enterprise Manager Command Line Interface (EMCLI) verbs for PaaS Infrastructure Zone and Software Pool operations for Enterprise Manager Self Service Applications Admin users. The EMCLI enables you to access Enterprise Manager Cloud Control functionality from text-based consoles (shells and command windows) for a variety of operating systems.

This chapter has the following sections:

- [SSA Admin Resources](#)
- [PaasZones](#)
- [PaasZone](#)
- [SoftwarePools](#)
- [SoftwarePool](#)
- [SoftwarePoolMetadata](#)
- [SoftwarePoolFilteredTargets](#)
- [SoftwarePoolCapacity](#)
- [Composite Resource Attributes](#)
- [EMCLI Verbs for Self Service Applications](#)

### 30.1 SSA Admin Resources

The Cloud APIs supported in this release can be categorized as follows:

- Collection Resources
  - [PaasZones](#)
  - [SoftwarePools](#)
- Singular Resources
  - [PaasZone](#)
  - [SoftwarePool](#)

Every singular resource has its own attributes and methods supported on it, but in general all singular resources have some common attributes, as described in the following table.

**Table 30–1 Common Resource Attributes for Singular Resources**

| Field         | Type      | Occurs | Supported Method | Description                                                      |
|---------------|-----------|--------|------------------|------------------------------------------------------------------|
| name          | String    | 1      | Any              | Name of Resource                                                 |
| type          | String    | 1      | Any              | Resource Type                                                    |
| description   | String    | 0..1   | Any              | Description                                                      |
| id            | String    | 0..1   | GET              | Resource Identifier                                              |
| createdOn     | TimeStamp | 0..1   | GET              | Date and time, in W3C format when the resource was created.      |
| createdBy     | String    | 0..1   | GET              | User who created resource                                        |
| modifiedOn    | TimeStamp | 0..1   | GET              | Date and time, in W3C format when the resource was last modified |
| ModifiedBy    | Date      | 0..1   | GET              | User who last modified resource                                  |
| owner         | String    | 0..1   | GET              | User who owns resource                                           |
| etag          | String    | 0..1   | GET              | HTTP Entity tag value                                            |
| lastModified  | String    | 0..1   | GET              | HTTP Last-Modified value                                         |
| canonicalLink | URI       | 1      | GET              | Preferred Version of Resource                                    |
| selfLink      | URI       | 1      | GET              | Refers to resource equivalent to the containing element          |

Where:

- Supported Method: Indicates in which method type this attribute will be available
- Supported Method type 'Any': Indicates that it might be available with GET, PUT, or POST methods, but not for DELETE.
- Occurs: An occurrence of 0..1 indicates it is an optional parameter. An occurrence of 1 means the parameter is mandatory.

### 30.1.1 GET to get API Entry Point Resources

The following shows a sample GET request and response:

#### Example Request:

- URL  
`https://example.oracle.com/em/websvcs/restful/extws/cloudservices/admin/cfw/api`
- Headers  
Authorization: basic xxxxxxxxx
- Method  
GET

#### Example Response:

```
{
```

```

v1:
{
canonicalLink:
"https://example.oracle.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1"
paasZones:
"https://example.oracle.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/pa
aszones"
softwarePools:
"https://example.oracle.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/so
ftwarepools"
}
}

```

## 30.2 PaasZones

PaasZones is a collection resource representing a Cloud Administrator's view of all the accessible PaasZone resources. The following table describes the PaasZones [application/oracle.com.cloud.common.PaaSZones+json] data model.

**Table 30–2 PaasZones Data Model**

| Field         | Type                     | Occurs | Description                                                                                   |
|---------------|--------------------------|--------|-----------------------------------------------------------------------------------------------|
| name          | String                   | 1      | Display name of this collection resource                                                      |
| type          | String                   | 1      | Type of this collection resource                                                              |
| hasMore       | Boolean                  | 1      | Indicates whether there are more elements in the collection                                   |
| count         | Integer                  | 1      | Number of elements returned                                                                   |
| items         | Collection<br><PaasZone> | 1      | List of Paas Zone resources.<br>In case no instance are there items will be present but empty |
| totalCount    | Integer                  | 1      | Total number of Paas Zone resources                                                           |
| canonicalLink | URI                      | 1      | A GET against this URI refreshes the client representation of this resource.                  |
| selfLink      | URI                      | 1      | Refers to the resource equivalent to the containing elements.                                 |

The following table describes the Supported Methods for PaasZones:

**Table 30–3 Supported Methods for PaasZones**

| Method       | Query Parameters | Request Payload         | URI                                                                                 | Description                                   |
|--------------|------------------|-------------------------|-------------------------------------------------------------------------------------|-----------------------------------------------|
| GET          | NA               | NA                      | <code>em/websvcs/restful/extws/cloudservices/admin/cfw/v1/paaszones/</code>         | Returns PaasZones collection resource details |
| GET/metadata | NA               | NA                      | <code>em/websvcs/restful/extws/cloudservices/admin/cfw/v1/paaszones/metadata</code> | Returns PaasZone Metadata                     |
| PUT          | NA               | PaasZone (Json Payload) | <code>em/websvcs/restful/extws/cloudservices/admin/cfw/v1/paaszones/</code>         | Updates PaasZone resource                     |

**Table 30–3 (Cont.) Supported Methods for PaasZones**

| Method | Query Parameters | Request Payload         | URI                                                                    | Description               |
|--------|------------------|-------------------------|------------------------------------------------------------------------|---------------------------|
| POST   | NA               | PaasZone (Json Payload) | <i>em/websvcs/restful//extws/cloudservices/admin/cfw/v1/paaszones/</i> | Creates PaasZone resource |

## 30.2.1 PaasZones API Examples

The following sections provide example PaasZones API operations. Note that you can also use EMCLI commands to perform a range of different operations (including, create, update, and delete operations) and these commands are described in [Section 30.10.2, "Using EMCLI Verbs for SSA Operations"](#).

### 30.2.1.1 Retrieving PaaS Zones Metadata

The GET request with metadata path is issued to fetch Metadata Information.

#### Example Request:

- URL
 

```
https://hostname/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/paaszones/metadata
```
- Headers
 

```
Authorization: basic xxxxxxxxx
Content-Type: application/json
```
- Method
 

```
GET
```

#### Example Response:

```
Status 200 OK
Content-Type: application/json
{
 "constraints" : [{
 "name" : "MAX_CPU_UTILIZATION",
 "description" : "Maximum CPU Utilization (%)",
 "value" : "80"
 }, {
 "name" : "MAX_MEMORY_ALLOCATION",
 "description" : "Maximum Memory Allocation (%)",
 "value" : "80"
 }],
 "memberTargetTypes" : ["oracle_vm_zone", "host"],
 "message" : "When creating Paas Infrastructure zone: (1) member_type should be set to one of the values shown in member target types. All Zone members have to be of the same type. (2) Please change the value of the constraints as required."
}
```

### 30.2.1.2 Listing PaaS Zones

The GET request can be issued to get list of PaasZone Instances.

**Example Request:**

## ■ URL

```
https://example.oracle.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/paaszones/
```

## ■ Headers

Authorization: basic xxxxxxxxx

Content-Type: application/json

## ■ Method

GET

**Example Response:**

Status 200 OK

Content-Type: application/json

```
{
 "name" : "Paas Zones",
 "type" : "Paas Infrastructure Zone",
 "hasMore" : false,
 "count" : 4,
 "items" : [{
 "name" : "pzone1",
 "id" : "A3CF2D49CFF3F3E664D073303EA51F8E",
 "type" : "self_service_zone",
 "canonicalLink" :
"https://example.oracle.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/paaszones/A3CF2D49CFF3F3E664D073303EA51F8E"
 }, {
 "name" : "Test Zone 2",
 "id" : "EF3830C71CC54B50B963376F9217AB95",
 "type" : "self_service_zone",
 "description" : "with MAX_CPU_UTILIZATION ",
 "canonicalLink" :
"https://example.oracle.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/paaszones/EF3830C71CC54B50B963376F9217AB95"
 }, {
 "name" : "test paas zone",
 "id" : "5D7548C1B879A51CFD894CEA8D5FB19B",
 "type" : "self_service_zone",
 "canonicalLink" :
"https://example.oracle.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/paaszones/5D7548C1B879A51CFD894CEA8D5FB19B"
 }, {
 "name" : "Test Zone",
 "id" : "CACBB07A674F280DF84FC4CBFA76DFA7",
 "type" : "self_service_zone",
 "description" : "with MAX_CPU_UTILIZATION ",
 "canonicalLink" :
"https://example.oracle.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/paaszones/CACBB07A674F280DF84FC4CBFA76DFA7"
 }],
 "totalCount" : 4,
 "canonicalLink" :
"https://example.oracle.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/paaszones",
 "selfLink" :
"https://example.oracle.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/pa"
```

```
aszones"
}
```

### 30.3 PaasZone

This resource represents a PaaS Infrastructure Zone. The following table describes the PaaSZone [application/oracle.com.cloud.common.PaaSZone+json] data model.

**Table 30–4 PaaSZone Data Model**

| Field           | Type                           | Supported Methods | Occurs | Description                                                                                                                                                |
|-----------------|--------------------------------|-------------------|--------|------------------------------------------------------------------------------------------------------------------------------------------------------------|
| name            | String                         | Any               | 1      | A human readable name given to the PaaS Zone.<br><br>It is non editable.                                                                                   |
| id              | String                         | GET               | 1      | Zone ID. GET method will have value for this element.                                                                                                      |
| description     | String                         | Any               | 0..1   | A brief description given to the zone.                                                                                                                     |
| credential      | String                         | Any               | 1      | Global Named Credentials to be used for provisioning in this PaaS Infrastructure Zone                                                                      |
| members         | Collection <String>            | GET<br>POST       | 1      | Members of PaaS Zone. They can be either Host members or Oracle VM Zone members.                                                                           |
| memberType      | String                         | GET<br>POST       | 1      | Target type of the members of PaaS Zone. It can be either "host" for Host members or "oracle_vm_zone" for Oracle VM Zone members.                          |
| constraints     | Collection <EntityValueHolder> | Any               | 0..1   | Placement Policy Constraints allow the cloud administrator to set maximum thresholds for any host.<br><br>Refer to zone metadata API for more information. |
| roles           | Collection <String>            | GET<br>POST       | 0..1   | SSA roles that can access this PaaS Infrastructure Zone.                                                                                                   |
| pools           | Collection <Software Pool>     | GET               | 0..1   | Software Pools associated with this PaaS Zone                                                                                                              |
| membersToAdd    | Collection <String>            | PUT               | 0..1   | Host or Oracle VM Zone targets to be added to an existing PaaS Zone while editing it.                                                                      |
| membersToRemove | Collection <String>            | PUT               | 0..1   | Host or Oracle VM Zone targets to be removed from an existing PaaS Zone while editing it.                                                                  |
| rolesToAdd      | Collection <String>            | PUT               | 0..1   | SSA roles to be added while editing a PaaS Zone                                                                                                            |
| rolesToRemove   | Collection <String>            | PUT               | 0..1   | SSA roles to be removed while editing a PaaS Zone.                                                                                                         |

Where:

- Supported Methods: Value 'Any' indicates GET/POST/PUT



Depending on the Supported Methods column value, the respective attributes need to be specified in the Request Payload or are available in the Response Payload after the operation in complete.

For example, while executing a PUT operation, the name, description, credential, membersToAdd, membersToRemove, constraints, rolesToAdd, and rolesToRemove attributes can be specified in the Request payload.

- The constraints attribute is not defined for a PaasZone when its memberType is oracle\_vm\_zone.
- The MAX\_CPU\_UTILIZATION and MAX\_MEMORY\_ALLOCATION are the two placement constraints supported for a PaaS Zone when the memberType is "host". If they are not provided, then their default value is taken to be 80 percent.

The following table describes the supported methods on the PaaSZone resource:

**Table 30–5 Supported Methods on PaaSZone Resources**

| Method | Query Parameters                                                                                                                                                        | Request Payload | Response Payload              | URI                                                                                 | Description                               |
|--------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|-------------------------------|-------------------------------------------------------------------------------------|-------------------------------------------|
| GET    | Possible values for this optional query parameter: <ul style="list-style-type: none"> <li>■ constraints</li> <li>■ pools</li> <li>■ members</li> <li>■ roles</li> </ul> | NA              | PaasZone (Json Payload)       | <code>em/websvcs/restful/extws/cloudservices/admin/cfw/v1/paaszones/{zoneId}</code> | Retrieve PaaS Infrastructure Zone details |
| DELETE | NA                                                                                                                                                                      | NA              | Status Message (json payload) | <code>em/websvcs/restful/extws/cloudservices/admin/cfw/v1/paaszones/{zoneId}</code> | Deletes a PaaS Infrastructure Zone        |

### 30.3.1 PaaS Zone API Examples

The following sections provide example Paas Zone API operations. Note that you can also use EMCLI commands to perform a range of different operations on a PaaS Zone (including, create, update, and delete) and these commands are described in [Section 30.10.2, "Using EMCLI Verbs for SSA Operations"](#).

#### 30.3.1.1 Creating a PaaS Zone

A PaaS Zone is created with a POST request on the paaszones collection resource.

##### Example Request:

- URL  
`https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/paaszones/`
- Headers  
Authorization: basic xxxxxxxxx  
Content-Type: application/json
- Method  
POST

```
{
 "name": "My Test Zone",
 "description": "This is a TestZone",
 "credential": "NCAIME",
 "memberType": "host",
 "members": [
 "slc05gvt.us.oracle.com"
],
 "constraintItems": [
 {
 "name": "MAX_CPU_UTILIZATION",
 "value": "85"
 },
 {
 "name": "MAX_MEMORY_ALLOCATION",
 "value": "95"
 }
],
 "rolesItems": [
 "SSACOREROLE1",
 "SSACOREROLE2"
],
}
```

**Example Response:**

Status 201 Created

Content-Type: application/json

```
{
 "name" : "My Test Zone",
 "description" : "This is a TestZone",
 "type" : "PaaS Infrastructure Zone",
 "id" : "DDBFEFDAD2AE6490E040F00AA37D4C67",
 "credential" : "NCAIME",
 "members" : {
 "canonicalLink" :
 "https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/paas
 zones/DDBFEFDAD2AE6490E040F00AA37D4C67/members",
 "memberType" : "host",
 "numberOfHostMembers" : "1"
 },
 "constraints" : {
 "canonicalLink" :
 "https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/paas
 zones/DDBFEFDAD2AE6490E040F00AA37D4C67/constraints",
 "numberOfConstraints" : "2"
 },
 "roles" : {
 "canonicalLink" :
 "https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/
 paaszones/DDBFEFDAD2AE6490E040F00AA37D4C67/roles",
 "numberOfUserRoles" : "2"
 },
 "pools" : {
 "canonicalLink" :
 "https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/paas
 zones/DDBFEFDAD2AE6490E040F00AA37D4C67/pools",
 "numberOfSoftwarePools" : "0"
 },
 "createdOn" : "2013-05-27 20:57:18",
}
```

```

 "createdBy" : null,
 "modifiedOn" : null,
 "modifiedBy" : null,
 "owner" : "SYSMAN",
 "etag" : null,
 "lastModified" : null,
 "canonicalLink" :
 "https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/paas
zones/DDBFEFDAD2AE6490E040F00AA37D4C67",
 "selfLink" :
 "https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/paas
zones/DDBFEFDAD2AE6490E040F00AA37D4C67"
}

```

### 30.3.1.2 Updating a PaaS Zone

The PUT request with payload updates an existing PaaS Infrastructure Zone.

#### Example Request:

- URL

```
https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/paaszones/
```

- Headers

Authorization: basic xxxxxxxxx

Content-Type: application/json

- Method

PUT

```

{
 "name": "My Test Zone",
 "credential": "NCAIME2",
 "description": "Description_01",
 "membersToAdd": [
 "somehost.com",
 "someotherhost.com"
],
 "membersToRemove": [
 "removehost.com",
 "removeotherhost2.com"
],
 "rolesToAdd": [
 "ROLE_1",
 "ROLE_2"
],
 "rolesToRemove": [
 "ROLE_1_1",
 "ROLE_2_1"
],
 "constraints": [
 {
 "name": "MAX_MEMORY_ALLOCATION",
 "value": "80"
 },
 {
 "name": "MAX_CPU_UTILIZATION",

```

```
 "value": "80"
 }
]
 }
}
```

**Example Response:**

```
Status 200 OK
Content-Type: application/json
{
 "message" : "PaaS Infrastructure Zone "My Test Zone" updated successfully."
}
```

**30.3.1.3 Deleting a PaaS Zone**

The DELETE request on the URI of the PAASZONE can be issued to delete a PaasZone Instance.

**Example Request:**

- URL  
  
https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/paaszones/DDBFEFDAD2AE6490E040F00AA37D4C67
- Headers  
  
Authorization: basic xxxxxxxxx
- Method  
  
DELETE

**Example Response:**

```
Status 200 OK
Content-Type: application/json
{
 "message" : "PaaS Infrastructure Zone "DDBFEFDAD2AE6490E040F00AA37D4C67"
deleted successfully."
}
```

---

---

**Note:** A PaaS Infrastructure Zone cannot be deleted if there are Software Pools associated with it.

---

---

**30.3.1.4 Retrieving PaasZone details**

The GET request on the PaasZone resource can be issued to retrieve its details.

**Example Request:**

- URL  
  
https://example.oracle.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/paaszones/A3CF2D49CFF3F3E664D073303EA51F8E
- Headers  
  
Authorization: basic xxxxxxxxx
- Method  
  
GET

**Example Response:**

Status 200 OK

Content-Type: application/json

```
{
 "name" : "pzone1",
 "description" : "This zone contains Host members",
 "type" : "PaaS Infrastructure Zone",
 "id" : "A3CF2D49CFF3F3E664D073303EA51F8E",
 "credential" : "SXAAJUNITHOSTCRED",
 "members" : {
 "canonicalLink" :
 "https://example.oracle.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/pa
 aszones/A3CF2D49CFF3F3E664D073303EA51F8E/members",
 "memberType" : "host",
 "numberOfHostMembers" : "2"
 },
 "constraints" : {
 "canonicalLink" :
 "https://example.oracle.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/pa
 aszones/A3CF2D49CFF3F3E664D073303EA51F8E/constraints",
 "numberOfConstraints" : "2"
 },
 "roles" : {
 "canonicalLink" :
 "https://example.oracle.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/pa
 aszones/A3CF2D49CFF3F3E664D073303EA51F8E/roles",
 "numberOfUserRoles" : "1"
 },
 "pools" : {
 "canonicalLink" :
 "https://example.oracle.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/pa
 aszones/A3CF2D49CFF3F3E664D073303EA51F8E/pools",
 "numberOfSoftwarePools" : "0"
 },
 "createdOn" : "2013-05-22 01:50:30",
 "createdBy" : null,
 "modifiedOn" : null,
 "modifiedBy" : null,
 "owner" : "SYSMAN",
 "etag" : null,
 "lastModified" : null,
 "canonicalLink" :
 "https://example.oracle.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/pa
 aszones/A3CF2D49CFF3F3E664D073303EA51F8E",
 "selfLink" :
 "https://example.oracle.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/pa
 aszones/A3CF2D49CFF3F3E664D073303EA51F8E"
}
```

**30.3.1.5 Getting a PaasZone (Expanded)**

The get request supports expands query parameter. Based on this query parameter the relevant attribute for a PaasZone is expanded.

The following query parameters are supported:

- members
- role
- constraints

## ■ pools

### URL is of the format

`https://example.oracle.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/paaszones/{zoneid}/?expands=<paramname>`

### Example:

`https://hostname/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/paaszones/DDBFEFDAD2AE6490E040F00AA37D4C67/?expands=constraints`

### Example Response

```
{
 "name" : "My Test Zone",
 "description" : "Description_02",
 "type" : "PaaS Infrastructure Zone",
 "id" : "DDBFEFDAD2AE6490E040F00AA37D4C67",
 "credential" : "C",
 "members" : {
 "canonicalLink" :
 "https://example.oracle.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/paaszones/DDBFEFDAD2AE6490E040F00AA37D4C67/members",
 "memberType" : "host",
 "numberOfHostMembers" : "1"
 },
 "constraints" : {
 "canonicalLink" :
 "https://example.oracle.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/paaszones/DDBFEFDAD2AE6490E040F00AA37D4C67/constraints",
 "items" : [{
 "name" : "MAX_MEMORY_ALLOCATION",
 "value" : "25"
 }, {
 "name" : "MAX_CPU_UTILIZATION",
 "value" : "25"
 }]
 },
 "roles" : {
 "canonicalLink" :
 "https://example.oracle.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/paaszones/DDBFEFDAD2AE6490E040F00AA37D4C67/roles",
 "numberOfUserRoles" : "0"
 },
 "pools" : {
 "canonicalLink" :
 "https://example.oracle.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/paaszones/DDBFEFDAD2AE6490E040F00AA37D4C67/pools",
 "numberOfSoftwarePools" : "0"
 },
 "createdOn" : "2013-05-27 20:57:18",
 "createdBy" : null,
 "modifiedOn" : null,
 "modifiedBy" : null,
 "owner" : "SYSMAN",
 "etag" : null,
 "lastModified" : null,
 "canonicalLink" :
 "https://example.oracle.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/paaszones/DDBFEFDAD2AE6490E040F00AA37D4C67",
 "selfLink" :
```

```
"https://example.oracle.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/pa
aszones/DDBFEFDAD2AE6490E040F00AA37D4C67"
}
```

## 30.4 SoftwarePools

SoftwarePools is a collection resource representing an SSA Administrator's view of all the accessible SoftwarePool resources. The following table describes the SoftwarePools [application/oracle.com.cloud.common.SoftwarePools+json] data model.

**Table 30–6 SoftwarePools Data Model**

| Field         | Type                          | Occurs | Description                                                                  |
|---------------|-------------------------------|--------|------------------------------------------------------------------------------|
| name          | String                        | 1      | Display name of this collection resource                                     |
| type          | String                        | 1      | Type of this collection resource                                             |
| hasMore       | Boolean                       | 1      | Indicates whether there are more elements in the collection                  |
| count         | Integer                       | 1      | Number of elements returned                                                  |
| items         | Collection<br><SoftwarePools> | 1      | The elements of this collection                                              |
| totalCount    | Integer                       | 1      | Total number of elements in the collection                                   |
| canonicalLink | URI                           | 1      | A GET against this URI refreshes the client representation of this resource. |
| selfLink      | URI                           | 1      | Refers to the resource equivalent to the containing elements.                |

The following table describes the supported methods on the SoftwarePools resource:

**Table 30–7 Supported Methods on SoftwarePools Resources**

| Method | Query Parameters     | Request Payload                   | URI                                                                                       | Description                                       |
|--------|----------------------|-----------------------------------|-------------------------------------------------------------------------------------------|---------------------------------------------------|
| GET    | zonename<br>pooltype | NA                                | <i>em/websvcs/restful/<br/>extws/cloudservices<br/>/admin/cfw/v1/softw<br/>arepools/</i>  | Returns the collection of existing Software Pools |
| PUT    | NA                   | SoftwarePo<br>ol(Json<br>Payload) | <i>em/websvcs/restful/<br/>extws/cloudservices<br/>/admin/cfw/v1/softw<br/>arepools/</i>  | Updates an existing Software Pool resource.       |
| POST   | NA                   | SoftwarePo<br>ol(Json<br>Payload) | <i>em/websvcs/restful/<br/>/extws/cloudservice<br/>s/admin/cfw/v1/soft<br/>warepools/</i> | Creates a Software Pool                           |

### 30.4.1 SoftwarePools API Example

The following provides an example of a SoftwarePools API operation. Note that you can also use EMCLI commands to perform a range of different operations on Software Pools (including, create, update, and delete) and these commands are described in [Section 30.10.2, "Using EMCLI Verbs for SSA Operations"](#)

### 30.4.1.1 Listing Software Pools

A GET request on SoftwarePools resource is issued to list all the accessible Software Pools.

#### Example Request:

- URL  
`https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools`
- Headers  
Authorization: basic xxxxxxxxx
- Method  
GET

#### Example Response:

```
Status 200 OK
Content-Type: application/json
{
 name: "Software Pools"
 type: "Software Pool"
 hasMore: false
 count: 2
 items: [2]
 0: {
 name: "IT_MW_Pool_2"
 type: "mwaas_zone"
 id: "E284FAAA7FBA6A06F7090115A3E07299"
 zoneName: "IT_MW_Zone1"
 canonicalLink:
 "https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/E284FAAA7FBA6A06F7090115A3E07299"
 description: "Middleware Pool of Oracle Homes of Version 10.3.5.0"
 }-
 1: {
 name: "IT_MW_Pool1"
 type: "mwaas_zone"
 id: "DD73A46E9E3C9866E040F00AE5235A7F"
 zoneName: "IT_MW_Zone1"
 canonicalLink:
 "https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/DD73A46E9E3C9866E040F00AE5235A7F"
 description: "Middleware Pool of Oracle Homes"
 }-
 -
 totalCount: 2
 canonicalLink:
 "https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/"
 selfLink:
 "https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/"
}
```



### 30.4.1.2 Filtering Output based on Search Query Parameters

Search on the Software Pools is supported on the following by the following query parameters:

- zonename
- pooltype

Single or multiple search parameters may be specified at once. Wildcard search is supported using '%' character.

#### Example Request:

- URL
 

```
https://
cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools? zonename=IT_MW%&pooltype=mwaas_zone
```
- Headers
 

```
Authorization: basic xxxxxxxxx
```
- Method
 

```
GET
```

#### Example Response:

```
Status 200 OK
Content-Type: application/json
{
 name: "Software Pools"
 type: "Software Pool"
 hasMore: false
 count: 1
 items: [1]
 0: {
 name: "IT_MW_Pool1"
 type: "mwaas_zone"
 id: "DD73A46E9E3C9866E040F00AE5235A7F"
 zoneName: "IT_MW_Zone1"
 canonicalLink:
 "https://cloudcompany/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/DD73A46E9E3C9866E040F00AE5235A7F"
 description: "Middleware Pool of Oracle Homes"
 }

 totalCount:1
 canonicalLink:
 "https://cloudcompany/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/"
 selfLink:
 "https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/"
}
```

## 30.5 SoftwarePool

A SoftwarePool resource represents a collection of homogeneous servers or clusters that are used to service requests within a PaaS Infrastructure Zone. All members

within a SoftwarePool must be of the same type and must belong to the same PaaS Infrastructure Zone.

In addition, to maintain homogeneity, the members must satisfy a set of filter criteria that restrict the addition of members to a SoftwarePool. The following table describes the SoftwarePool [application/oracle.com.cloud.common.SoftwarePool+json] data model.

**Table 30–8 SoftwarePool Data Model**

| Field           | Type                              | Supported Methods | Occurs | Description                                                                                                                                                              |
|-----------------|-----------------------------------|-------------------|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| name            | String                            | Any               | 1      | A human readable name given to the Software Pool.<br>It is non editable.                                                                                                 |
| id              | String                            | GET               | 1      | Pool GUID that uniquely identifies the resource.                                                                                                                         |
| description     | String                            | Any               | 0..1   | A brief description given to the Software Pool.                                                                                                                          |
| poolTargetType  | String                            | Any               | 1      | Target type of the Software Pool to be created.<br><br>Example : "mwaas_zone" for Middleware Pool, "oracle_cloud_zone" for Database Pool, "schaas_pool" for Schema Pool. |
| paasZone        | String                            | GET<br>POST       | 1      | Name of the PaaS Zone associated with this Software Pool.                                                                                                                |
| members         | Collection<br><String>            | GET<br>POST       | 1      | Members of PaaS Zone. They can be either Host members or Oracle VM Zone members.                                                                                         |
| constraints     | Collection<br><EntityValueHolder> | Any               | 0..1   | Placement constraints for a Software Pool that allow the self service administrator to set maximum ceilings for resource utilization.                                    |
| filters         | Collection<br><EntityValueHolder> | GET<br>POST       | 1      | Filters on a Software Pool restrict the addition of member targets to it with a set criteria.                                                                            |
| properties      | Collection<br><EntityValueHolder> | Any               | 1      | Additional properties that need to be specified for a specific pool target type.                                                                                         |
| membersToAdd    | Collection<br><String>            | PUT               | 0..1   | Targets to be added to an existing Software Pool while editing it.                                                                                                       |
| membersToRemove | Collection<br><String>            | PUT               | 0..1   | Targets to be removed from an existing Software Pool while editing it.                                                                                                   |
| canonicalLink   | URI                               | GET               | 1      | A GET against this URI refreshes the client representation of this resource.                                                                                             |

**Note:** Supported methods value 'Any' indicates GET/POST/PUT .

The following table describes the Supported Methods for Software Pool resources:

**Table 30–9 Supported Methods on Software Pool Resources**

| Method | Query Parameters                                                                                                                                                               | Request Payload | Response Payload              | URI                                                                               | Description                         |
|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------|-------------------------------|-----------------------------------------------------------------------------------|-------------------------------------|
| GET    | param-name : expands<br>param-value : members/<br>filters/<br>constraints/all<br><br>Expands a particular attribute to show its data.<br><br>"all" expands all the attributes. | NA              | SoftwarePool (Json Payload)   | <i>em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/{poolId}</i> | Retrieves the Software Pool details |
| DELETE | NA                                                                                                                                                                             | NA              | Status Message (json payload) | <i>em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/{poolId}</i> | Deletes a Software Pool             |

## 30.5.1 SoftwarePool API Examples

The following sections provide example SoftwarePool API operations. Note that you can also use EMCLI commands to perform a range of different operations on a Software Pool (including, create, update, and delete operations) and these commands are described in [Section 30.10.2, "Using EMCLI Verbs for SSA Operations"](#)

### 30.5.1.1 Creating a Software Pool

A POST operation is issued on Software Pools Collection resource to create a Software Pool.

#### Example Request:

- URL
 

```
https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools
```
- Headers
 

```
Authorization: basic xxxxxxxxx
Content-Type : application/json
```
- Method
 

```
POST

{
 "name": "IT_MW_Pool1",
 "poolTargetType": "mwaas_zone",
 "description": "Middleware Pool of Oracle Homes",
 "paasZone": "IT_MW_Zone1",
 "members": ["WebLogicServer10_3_6_0_slc01afx_6775", "WebLogicServer10_3_6_0_adc2201439_4731"],
 "constraints":
```

```
 {"items": [{"name": "MAX_INSTANCES", "value": "10"}]},
 "filters":
 {"items": [{"name": "VERSION", "value": "10.3.6.0"}]}
}
```

**Example Response:**

```
Status 201 Created
Content-Type: application/json
{
 name: "IT_MW_Pool1"
 description: "Middleware Pool of Oracle Homes"
 type: "Middleware Pool"
 id: "687FE169ED3556CED38D11DC9BDD5CCD"
 zoneName: "IT_MW_Zone1"
 members: {
 canonicalLink:
 "https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/687FE169ED3556CED38D11DC9BDD5CCD/members"
 numberOfPoolMembers: "2"
 }-
 constraints: {
 canonicalLink:
 "https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/687FE169ED3556CED38D11DC9BDD5CCD/constraints"
 numberOfConstraints: "1"
 }-
 filters: {
 canonicalLink:
 "https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/687FE169ED3556CED38D11DC9BDD5CCD/filters"
 numberOfFilters: "1"
 }-
 properties: {
 canonicalLink:
 "https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/687FE169ED3556CED38D11DC9BDD5CCD/properties"
 numberOfPropertiess: "0"
 }-
 createdOn: "2013-05-24 00:26:42"
 createdBy: null
 modifiedOn: null
 modifiedBy: null
 owner: "CLOUD_ADMIN1"
 etag: null
 lastModified: null
 canonicalLink:
 "https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/687FE169ED3556CED38D11DC9BDD5CCD"
 selfLink:
 "https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/687FE169ED3556CED38D11DC9BDD5CCD"
}
```

**30.5.1.2 Updating a Software Pool**

A PUT operation is issued on Software Pools Collection resource to update an Software Pool.

**Example Request:**

## ■ URL

```
https://
cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools
```

## ■ Headers

Authorization: basic xxxxxxxxx

Content-Type : application/json

## ■ Method

PUT

```
{
 "name": "IT_MW_Pool1",
 "poolTargetType": "mwaas_zone",
 "description": "Middleware Pool of Oracle Homes Updated",
 "membersToAdd": ["WebLogicServer10_3_6_0_slc0lafx_6776", "WebLogicServer10_3_6_0_adc2201439_4732"],
 "constraints":
 {"items": [{"name": "MAX_INSTANCES", "value": "10"}]},
}
```

**Example Response:**

Status 200 OK

Content-Type: application/json

```
{
 message: "Software Pool " IT_MW_Pool1" updated successfully."
}
```

**30.5.1.3 Deleting a Software Pool**

The following shows an example request and response:

**Example Request:**

## ■ URL

```
https://
cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/687FE169ED3556CED38D11DC9BDD5CCD
```

## ■ Headers

Authorization: basic xxxxxxxxx

## ■ Method

DELETE

**Example Response:**

Status 200 OK

Content-Type: application/json

```
{
 message: "Software Pool "687FE169ED3556CED38D11DC9BDD5CCD" deleted successfully."
}
```

---

---

**Note:** Software Pool cannot be deleted if there are Service Templates associated with the Pool

---

---

#### 30.5.1.4 Retrieving Software Pool Details

A GET request on SoftwarePool resource is issued to get a specific Software Pool details.

##### Example Request:

- URL

```
https://
cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/687FE169ED3556CED38D11DC9BDD5CCD
```

- Headers

Authorization: basic xxxxxxxxx

Content-Type : application/json

- Method

GET

##### Example Response:

```
Status 200 OK
Content-Type: application/json
{
 name: "IT_MW_Pool1"
 description: "Middleware Pool of Oracle Homes"
 type: "Middleware Pool"
 id: "687FE169ED3556CED38D11DC9BDD5CCD"
 zoneName: "IT_MW_Zone1"
 members: {
 canonicalLink:
 "https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/687FE169ED3556CED38D11DC9BDD5CCD/members"
 numberOfPoolMembers: "2"
 }-
 constraints: {
 canonicalLink:
 "https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/687FE169ED3556CED38D11DC9BDD5CCD/constraints"
 numberOfConstraints: "1"
 }-
 filters: {
 canonicalLink:
 "https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/687FE169ED3556CED38D11DC9BDD5CCD/filters"
 numberOfFilters: "1"
 }-
 properties: {
 canonicalLink:
 "https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/687FE169ED3556CED38D11DC9BDD5CCD/properties"
 numberOfPropertyess: "0"
 }-
 createdOn: "2013-05-24 00:26:42"
```

```

createdBy: null
modifiedOn: null
modifiedBy: null
owner: "CLOUD_ADMIN1"
etag: null
lastModified: null
canonicalLink:
"https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/687FE169ED3556CED38D11DC9BDD5CCD"
selfLink:
"https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/687FE169ED3556CED38D11DC9BDD5CCD"
}

```

### 30.5.1.5 Expands Query Parameter

GET operation on SoftwarePool resource supports the expands query parameter to expand the following link resource attributes:

- **members:** Expands the members attribute to show all the members of this Software Pool.
- **filters:** Expands the filters attribute to show the filter criteria of the members of this Software Pool.
- **constraints:** Expands the constraints attribute to show the placement constraints associated with this Software Pool.
- **all :** Expands all the attributes of the Software Pool when returning the Software Pool details.

The URL has the following format:

```
https://hostname/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/{poolid}?expands=<paramname>
```

Example:

```
https://hostname/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/687FE169ED3556CED38D11DC9BDD5CCD?expands=members
```

If no query parameter is specified then software pool details data is returned to the user without any attributes expanded.

### 30.5.1.6 Retrieving Software Pool Capacity

A GET on SoftwarePoolCapacity resource is issued to get the Software Pool capacity details.

#### Example Request:

- **URL**

```
https://
cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/687FE169ED3556CED38D11DC9BDD5CCD/capacity
```
- **Headers**

```
Authorization: basic xxxxxxxxx
```
- **Method**

```
GET
```

**Example Response:**

```
Status 200 OK
Content-Type: application/json
{
 name: "mySchemaPool"
 type: "schaas_pool"
 capacity: {
 items: [1]
 0: {
 entityName: "Oemrep_Database"
 entityType: "oracle_database"
 entityCapacity: {
 items: [4]
 0: {
 resourceType: "Storage Utilization"
 resourceUsed: "64.74"
 resourceThreshold: "100.0"
 resourceUnit: "%"
 }-
 }-
 }-
 1: {
 resourceType: "CPU Utilization"
 resourceUsed: "0.0"
 resourceThreshold: "0.0"
 resourceUnit: "%"
 }-
 2: {
 resourceType: "Memory Utilization"
 resourceUsed: "0.0"
 resourceThreshold: "0.0"
 resourceUnit: "%"
 }-
 3: {
 resourceType: "Instances"
 resourceUsed: "0"
 resourceThreshold: "100"
 }-
 }-
 canonicalLink:
 "https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/2E7C9C632B18C3FEB4FEA5A614E51BE7/capacity"
}
```

## 30.6 SoftwarePoolMetadata

This provides the metadata information for all the software pool target types that is needed while creating a Software Pool. The following table describes the SoftwarePoolMetadata [application/oracle.com.cloud.common.SoftwarePools+json] data model.

**Table 30–10** *SoftwarePoolMetadata Data Model*

| Field         | Type | Occurs | Description                                                                  |
|---------------|------|--------|------------------------------------------------------------------------------|
| canonicalLink | URI  | 1      | A GET against this URI refreshes the client representation of this resource. |



**Table 30–10 (Cont.) SoftwarePoolMetadata Data Model**

| Field                | Type                        | Occurs | Description                                                                                          |
|----------------------|-----------------------------|--------|------------------------------------------------------------------------------------------------------|
| poolType             | String                      | 1      | Target type of the Software Pool                                                                     |
| poolTypeDisplay Name | String                      | 0..1   | Display name of this Software Pool target type.                                                      |
| poolTypeDescription  | String                      | 0..1   | Description of the Software Pool target type.                                                        |
| serviceFamily        | String                      | 1      | Represents the service family to which this pool type belongs.                                       |
| serviceType          | String                      | 1      | Represents the service type to which this pool type belongs.                                         |
| constraints          | Collection<ValueDescriptor> | 0..1   | Represents the metadata information about the possible placement constraints for this pool type.     |
| filters              | Collection<ValueDescriptor> | 0..1   | Represents the metadata information about the possible filters (member constraints) for a pool type. |

The GET method for the SupportedPoolMetadata resource has the following characteristics:

- Response Payload:  
SoftwarePoolMetadata (Json Payload)
- URI:  
`em/websvcs/restful//extws/cloudservices/admin/cfw/v1/softwarepools/metadata`
- Descriptions:  
Retrieves the metadata information for all the software pool types like service type, service family, available constraints, and available filters.

### 30.6.1 Retrieving Software Pool Metadata API Example

A GET request is issued to obtain the SoftwarePoolMetadata URI, which is then used to retrieve the metadata information of all the software pool types. Note that you can also use EMCLI commands to retrieve Software Pool information and these commands are described in [Section 30.10.2, "Using EMCLI Verbs for SSA Operations"](#)

#### Example Request:

- URL  
`https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/metadata`
- Headers  
Authorization: basic xxxxxxxxx
- Method  
DELETE

**Example Response:**

Status 200 OK

Content-Type: application/json

```
{
 softwarePoolsMetadata:
 {
 items: [1]
 0:
 {
 poolType: "mwaas_zone"
 poolTypeDisplayName: "Middleware Pool"
 poolTypeDescription: "Middleware Pool is a set of homogeneous Middleware Oracle
 Homes."
 serviceFamily: "MWAAS"
 serviceType: "PhysicalWLS"
 constraints:
 {
 items: [1]
 0:
 {
 name: "MAX_INSTANCES"
 description: "Maximum Number of Java Servers (per host)"
 defaultValue: "1"
 isRequiredValue: true
 }
 }
 filters:
 {
 items: [1]
 0:
 {
 name: "VERSION"
 description: "Version"
 defaultValue: "10.3.5.0"
 isRequiredValue: true
 possibleValues: [7]
 0:
 {value: "12.1.1.0"
 description: "12.1.1.0"}
 1:
 {
 value: "10.3.6.0"
 description: "10.3.6.0"
 }
 2:
 {
 value: "10.3.5.0"
 description: "10.3.5.0"
 }
 3:
 {
 value: "10.3.4.0"
 description: "10.3.4.0"
 }
 4:
 {
 value: "10.3.3.0"
 description: "10.3.3.0"
 }
 }
 }
 }
 }
}
```

```

5:
{
 value: "10.3.2.0"
 description: "10.3.2.0"
}
6:
{
 value: "10.3.1.0"
 description: "10.3.1.0"
}
}
}
}
}
canonicalLink:
"https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/soft
warepools/metadata"
}

```

## 30.7 SoftwarePoolFilteredTargets

This gives the available targets in a given PaaS Infrastructure Zone based on the filter criteria given, that can be used in POST operation on SoftwarePool resource. The following table describes the SoftwarePoolFilteredTargets [application/oracle.com.cloud.common.SoftwarePoolFilteredTargets+json] data model.

**Table 30–11 SoftwarePoolFilteredTargets Data Model**

| Field         | Type                      | Occurs | Description                                                                                                                                                                                                                                                        |
|---------------|---------------------------|--------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| canonicalLink | URI                       | 1      | A GET against this URI refreshes the client representation of this resource.                                                                                                                                                                                       |
| targetType    | String                    | 1      | Target type of the Software Pool .<br>Example : "mwaas_zone" for Middleware Pool, "oracle_cloud_zone" for Database Pool, "schaas_pool" for Schema Pool.<br><br>Refer SoftwarePoolsMetadata resource to get the available filters for a pool target type.<br>[POST] |
| paasZone      | String                    | 1      | Name of the PaaS Zone representing a filter criteria.<br>[POST]                                                                                                                                                                                                    |
| targets       | Map<String, List<String>> | 1      | Filtered targets after applying the filter criteria.                                                                                                                                                                                                               |
| filters       | Collection<Filters>       | 1      | Filter criteria that needs to be applied to get the filtered targets.<br><br>Refer SoftwarePoolsMetadata resource to get the available filters for a pool target type.<br>[POST]                                                                                   |

Note that filtered targets satisfy the following:

- Are part of the given PaaS Infrastructure zone.

- Satisfy the filter criteria passed.
- Are not part of any other Software Pool.

The POST method for the SupportedPoolFilteredTargets resource has the following characteristics:

- Request Payload:  
SoftwarePoolFilteredTargets (Json Payload)
- Response Payload  
SoftwarePoolFilteredTargets (Json Payload)
- URI:  
*em/websvcs/restful//extws/cloudservices/admin/cfw/v1/software  
pools/filteredtargets*
- Descriptions:  
Retrieves the metadata information for all the software pool types, such as service type, service family, available constraints, and available filters.

### 30.7.1 Retrieving Available Filtered Targets API Example

A POST operation on SoftwarePoolFilteredTargets resource is issued to fetch the filtered targets based on the paas zone and filters given. Note that you can also use EMCLI commands to retrieve Filtered Targets information and these commands are described in [Section 30.10.2, "Using EMCLI Verbs for SSA Operations"](#)

#### Example Request:

- URL  

```
https://
cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools
```
- Headers  
Authorization: basic xxxxxxxxx  
Content-Type : application/json
- Method  
POST  

```
{
 "targetType": "mwaas_zone",
 "paasZone": "IT_MW_Zone1",
 "filters": {"items": [{"name": "VERSION", "value": "10.3.6.0"}]}
}
```

#### Example Response:

```
Status 201 Created
Content-Type: application/json

{
 targets: {
 oracle_home: [2]
 0: "WebLogicServer10_3_6_0_slc01afx_6775"
 1: "WebLogicServer10_3_6_0_adc2201439_4731"
```

```

}
canonicalLink:
"https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/filteredtargets"
}

```

## 30.8 SoftwarePoolCapacity

This resource retrieves the software pool capacity details like CPU utilization, memory allocation, and number of instances for the pool entities. The following table describes the SoftwarePoolCapacity [application/oracle.com.cloud.common.SoftwarePoolCapacity+json] data model.

**Table 30–12 SoftwarePoolCapacity Data Model**

| Field         | Type                | Occurs | Description                                                                                                                                                                      |
|---------------|---------------------|--------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| canonicalLink | URI                 | 1      | A GET against this URI refreshes the client representation of this resource.                                                                                                     |
| name          | String              | 1      | Name of the Software Pool.<br>[POST]                                                                                                                                             |
| type          | String              | 1      | Target type of the Software Pool.<br>Example : "mwaas_zone" for Middleware Pool, "oracle_cloud_zone" for Database Pool, "schaas_pool" for Schema Pool.                           |
| capacity      | Collection<Filters> | 1      | Filter criteria that needs to be applied to get the filtered targets.<br><br>Refer SoftwarePoolsMetadata resource to get the available filters for a pool target type.<br>[POST] |

The GET method for the SupportedPoolCapacity resource has the following characteristics:

- Response Payload  
SoftwarePoolCapacity (Json Payload)
- URI:  
`em/websvcs/restful//extws/cloudservices/admin/cfw/v1/softwarepools/{poolId}/capacity`
- Descriptions:  
Retrieves the software pool capacity details.

### 30.8.1 Retrieving Software Pool Capacity API Example

A GET on SoftwarePoolCapacity resource is issued to get the Software Pool capacity details. Note that you can also use EMCLI commands to retrieve Software Pool Capacity information and these commands are described in [Section 30.10.2, "Using EMCLI Verbs for SSA Operations"](#)

**Example Request:**

- URL

```
https://
cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/687FE169ED3556CED38D11DC9BDD5CCD/capacity
```

- Headers

Authorization: basic xxxxxxxxx

- Method

GET

**Example Response:**

```
Status 200 OK
Content-Type: application/json
{
 name: "mySchemaPool"
 type: "schaas_pool"
 capacity: {
 items: [1]
 0: {
 entityName: "Oemrep_Database"
 entityType: "oracle_database"
 entityCapacity: {
 items: [4]
 0: {
 resourceType: "Storage Utilization"
 resourceUsed: "64.74"
 resourceThreshold: "100.0"
 resourceUnit: "%"
 }-
 }-
 1: {
 resourceType: "CPU Utilization"
 resourceUsed: "0.0"
 resourceThreshold: "0.0"
 resourceUnit: "%"
 }-
 2: {
 resourceType: "Memory Utilization"
 resourceUsed: "0.0"
 resourceThreshold: "0.0"
 resourceUnit: "%"
 }-
 3: {
 resourceType: "Instances"
 resourceUsed: "0"
 resourceThreshold: "100"
 }-
 }-
 }-
 canonicalLink:
 "https://cloudcompany.com/em/websvcs/restful/extws/cloudservices/admin/cfw/v1/softwarepools/2E7C9C632B18C3FEB4FEA5A614E51BE7/capacity"
}
```

## 30.9 Composite Resource Attributes

Some of the resource entities in this specification contain the following composite resource attributes:

- [ValueDescriptor](#)
- [ValueEntity](#)
- [EntityValueHolder](#)
- [PoolEntityCapacity](#)
- [ResourceTypeCapacity](#)

### 30.9.1 ValueDescriptor

This describes the metadata regarding the type of value that can be set upon an associated object. The following table describes the data model:

**Table 30–13** *ValueDescriptor Data Model*

| Field           | Type                    | Occurs | Description                                                 |
|-----------------|-------------------------|--------|-------------------------------------------------------------|
| name            | String                  | 1      | Name of the datum being described.                          |
| description     | String                  | 1      | Description of the datum being described.                   |
| defaultValue    | String                  | 1      | Default value of the datum being described.                 |
| isRequiredValue | Boolean                 | 1      | True if the value is required.                              |
| possibleValues  | Collection<ValueEntity> | 0..1   | Optional list of valid values of the datum being described. |

### 30.9.2 ValueEntity

This represents a value (string representation) and their description (caption). The following table describes the data model:

**Table 30–14** *ValueEntity Data Model*

| Field       | Type   | Occurs | Description                                              |
|-------------|--------|--------|----------------------------------------------------------|
| value       | String | 1      | String representation of the actual value of the entity. |
| description | String | 0..1   | Caption of the entity that describes it.                 |

### 30.9.3 EntityValueHolder

This represents a simple entity having a name, value and description. The following table describes the data model:

**Table 30–15** *EntityValueHolder Data Model*

| Field       | Type   | Occurs | Description                                              |
|-------------|--------|--------|----------------------------------------------------------|
| name        | String | 1      | Name of the entity                                       |
| value       | String | 1      | String representation of the actual value of the entity. |
| description | String | 0..1   | Caption of the entity that describes it.                 |

### 30.9.4 PoolEntityCapacity

This represents the capacity details of individual entities of the Software Pool. The following table describes the data model:

**Table 30–16 PoolEntityCapacity Data Model**

| Field          | Type                              | Occurs | Description                                                                                             |
|----------------|-----------------------------------|--------|---------------------------------------------------------------------------------------------------------|
| entityName     | String                            | 1      | Name of the entity                                                                                      |
| entityType     | String                            | 1      | Target type of the entity.                                                                              |
| entityCapacity | Collection<Resource TypeCapacity> | 1      | Collection of capacity details of a particular resource type like Storage Utilization, CPU Utilization. |

### 30.9.5 ResourceTypeCapacity

This represents the capacity details of a particular resource type, such as Storage Utilization and CPU Utilization. The following table describes the data model:

**Table 30–17 ResourceTypeCapacity Data Model**

| Field             | Type   | Occurs | Description                                                                             |
|-------------------|--------|--------|-----------------------------------------------------------------------------------------|
| resourceType      | String | 1      | Type of the resource like Storage Utilization.                                          |
| resourceUsed      | String | 1      | String representation of the used resource.                                             |
| resourceThreshold | String | 0..1   | String representation of the resource threshold.                                        |
| resourceUnit      | String | 0..1   | String representation of the unit in which the utilization of the resource is measured. |

## 30.10 EMCLI Verbs for Self Service Applications

The Enterprise Manager Command Line Interface (EMCLI) enables you to access Enterprise Manager Cloud Control functionality from text-based consoles (shells and command windows) for a variety of operating systems. You can call Enterprise Manager functionality using custom scripts, such as SQL\*Plus, OS shell, Perl, or Tcl, thus easily integrating Enterprise Manager functionality with a company's business process.

This section provides details on the EMCLI verbs for PaaS Infrastructure Zone and Software Pool operations for Enterprise Manager Self Service Applications Admin users.

### 30.10.1 Introduction

EMCLI verbs will be supported for the following operations in Enterprise Manager release 12.1.0.5:

- PaaS Infrastructure Zones and Software Pools – Create, Edit, List, Delete
- PaaS Infrastructure Zones and Software Pools – Add and Remove members

All PaaS Infrastructure Zone EMCLI operations are available only to users with a EM\_CLOUD\_ADMINISTRATOR role and Software Pool operations can be performed only by users with a EM\_SSA\_ADMINISTRATOR role. EMCLI verbs will obtain user information from the security context and verify only those users authorized to perform these operations.



### 30.10.2 Using EMCLI Verbs for SSA Operations

The following table provides details of the supported verbs. Note that these verbs are available to SSA Administrator users only and cannot be used with an EM\_SSA\_USER role.

**Table 30–18 Verb Details**

| Verb             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        | Format                                                                                                                                                                                                                                                                                                                                                                                                                                             |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| create_paas_zone | <p>Creates a PaaS Infrastructure Zone.</p> <ul style="list-style-type: none"> <li>name - Name of the PaaS Infrastructure Zone to be created</li> <li>credential - Global Named Credentials to be used for provisioning in this PaaS Infrastructure Zone.</li> <li>hosts - Comma separated list of the Host targets to be added as members of this PaaS Infrastructure Zone. PaaS Infrastructure Zone can contain either hosts or Oracle VM Zones as members.</li> <li>ovm_zones - Comma separated list of the Oracle VM Zone targets to be added as members of this PaaS Infrastructure Zone.</li> <li>roles - Comma separated list of SSA roles that can access this PaaS Infrastructure Zone.</li> <li>description - Description of the PaaS Infrastructure Zone.</li> <li>cpu_utilization - Placement Policy Constraints allow the cloud administrator to set maximum thresholds for any host. Value entered must be between 1 and 100. If not provided, default value is taken to be 80 percent. Parameter is not needed if Oracle VM Zone targets are added as members.</li> <li>memory_utilization - Another Placement Policy Constraint for PaaS Infrastructure Zone. Value entered must be between 1 and 100. If not provided, default value is taken to be 80 percent. Parameter is not needed if Oracle VM Zone targets are added as members.</li> </ul> | <p>emcli create_paas_zone-name=&lt;Name of PaaS Zone&gt;</p> <p>-credential=&lt;Global Named Credential&gt;</p> <p>[-hosts=&lt;Host1,Host2,Host3...&gt;]</p> <p>[-ovm_zones=&lt;OVMZone1,OVMZone2,OVMZone3...&gt;]</p> <p>[-roles=&lt;SsaRole1,SsaRole2,...&gt;]</p> <p>[-description=&lt;Description of PaaS Zone&gt;]</p> <p>[-cpu_utilization=&lt;Value between 1 and 100&gt;]</p> <p>[-memory_utilization=&lt;Value between 1 and 100&gt;]</p> |

Example:

```
emcli create_paas_zone
 -name=MyPaaSZone
 -credential=BBANTHIA
 -hosts=slc03qtn.us.oracle.com
 -roles="SSA_USER_ROLE1,SSA_USER_ROLE2"
 -description="Test PaaS zone"
 -cpu_utilization=40
 -memory_utilization=70
```

Sample Output:

```
PaaS Infrastructure Zone "MyPaaSZone" created
successfully.
```

**Table 30–18 (Cont.) Verb Details**

| Verb                 | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            | Format                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| update_paaS_zone     | <p>Updates a PaaS Infrastructure Zone.</p> <ul style="list-style-type: none"> <li>name - Name of the existing PaaS Infrastructure Zone.</li> <li>credential - Global Named Credentials to be updated.</li> <li>add_hosts - Comma separated list of the Host targets to be added as members of this PaaS Infrastructure Zone. The hosts to be added must not be already added to other existing PaaS Zones.</li> <li>remove_hosts - Comma separated list of the Host targets to be removed as members from this PaaS Infrastructure Zone.</li> <li>add_ovm_zones - Comma separated list of the Oracle VM Zone targets to be added as members of this PaaS Infrastructure Zone. The Oracle VM Zones to be added must not be already added to other existing PaaS Zones.</li> <li>remove_ovm_zones - Comma separated list of the Oracle VM Zone targets to be removed as members from this PaaS Infrastructure Zone.</li> <li>add_roles - Comma separated list of SSA roles to be added that can access this PaaS Infrastructure Zone.</li> <li>remove_roles - Comma separated list of SSA roles to be removed from this PaaS Infrastructure Zone.</li> <li>description - Updated description of the PaaS Infrastructure Zone.</li> <li>cpu_utilization - Placement Policy Constraints allow the cloud administrator to set maximum thresholds for any host. Value entered must be between 1 and 100.</li> <li>memory_utilization - Another Placement Policy Constraint for PaaS Infrastructure Zone. Value entered must be between 1 and 100.</li> </ul> | <p>emcli update_paaS_zone-name=&lt;Name of PaaS Zone&gt;</p> <p>[-description=&lt;Description of PaaS Zone&gt;]</p> <p>[-credential=&lt;Global Named Credential&gt;]</p> <p>[-add_hosts=&lt;Host1,Host2,Host3...&gt;]</p> <p>[-remove_hosts=&lt;Host4,Host5...&gt;]</p> <p>[-add_ovm_zones=&lt;OVMZone1,OVMZone2,OVMZone3...&gt;]</p> <p>[-remove_ovm_zones=&lt;OVMZone4,OVMZone5...&gt;]</p> <p>[-add_roles=&lt;SsaRole1,SsaRole2,...&gt;]</p> <p>[-remove_roles=&lt;SsaRole3,SsaRole4,...&gt;]</p> <p>[-cpu_utilization=&lt;Value between 1 and 100&gt;]</p> <p>[-memory_utilization=&lt;Value between 1 and 100&gt;]</p> |
| delete_paaS_zone     | <p>Deletes a PaaS Infrastructure Zone. PaaS Infrastructure Zone cannot be deleted if there is an existing Software Pool associated with it.</p> <p>name - Name of the existing PaaS Infrastructure Zone to be deleted.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             | <p>emcli delete_paaS_zone -name=&lt;Name of PaaS Zone&gt;</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| get_paaS_zone_detail | <p>Retrieves the PaaS Infrastructure Zone details.</p> <p>name - Details of the existing PaaS Infrastructure Zone, such as Name, Description, Named Credentials, Number of Hosts, Roles, Maximum Memory Allocation (%), and Maximum CPU Utilization (%).</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | <p>emcli get_paaS_zone_detail -name=&lt;Name of PaaS Zone&gt;</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |

**Table 30–18 (Cont.) Verb Details**

| Verb                                                                                                                                                                                                                                                                                                                                                                                                                             | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              | Format                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| create_pool                                                                                                                                                                                                                                                                                                                                                                                                                      | Creates a Software Pool. <ul style="list-style-type: none"> <li>name - Name of the Software Pool to be created</li> <li>target_type - Target type of the Software Pool to be created.</li> <li>paas_zone - Name of PaaS Infrastructure Zone in which the Software Pool is to be created.</li> <li>members - Comma separated list of targets to be added as members of the Software Pool.</li> <li>description - Description of the Software Pool.</li> <li>placement_constraints - Comma separated key value pairs of the placement constraints that allow the self service administrator to set maximum ceilings for resource utilization. This provides protection for the members of the Software Pool in terms of resource consumption. Refer to the get_pool_allowed_placement_constraints verb to get the available placement constraints for a pool target type.</li> <li>member_constraints - Comma separated key value pairs that restricts the addition of member targets to a Software Pool with a set criteria. Refer to the get_pool_allowed_member_constraints verb to get the available member constraints and their possible values for a pool target type.</li> <li>properties - Comma separated key value pairs for additional properties that need to be specified based on the Software Pool target type.</li> </ul> | emcli create_pool<br>-name=<Name of Software Pool><br><br>-target_type=<Target type of Software Pool><br><br>-paas_zone=<Paas Infrastructure Zone of Software Pool><br><br>-members=<Member1, Member2...><br><br>[-description=<Description of Software Pool>]<br><br>[-placement_constraints=<constraint1=value1, constraint2=value2...>]<br><br>[-member_constraints=<constraint1=value1, constraint2=value2>]<br><br>[-properties=<property1=value1, property2=value2>] |
| Example:                                                                                                                                                                                                                                                                                                                                                                                                                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <pre>emcli create_pool -name=MyDbPool -target_type=oracle_cloud_zone -paas_zone=MyPaaSZone -members=OraDb11g_home1_1_slc03qtn -description="Test Database Pool" -member_constraints="CONFIGURATION=oracle_ database,VERSION=11.2.0.3.0,PLATFORM=226" -placement_constraints="MAX_INSTANCES=7" -properties="host_credential_ guid=DBA449B8967AAF77E040F00A73B11F55, root_credential_ guid=DBA449B8967AAF77E040F00A73B11F55"</pre> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Sample Output:                                                                                                                                                                                                                                                                                                                                                                                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Software Pool "MyDbPool" created successfully.                                                                                                                                                                                                                                                                                                                                                                                   |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| Note:                                                                                                                                                                                                                                                                                                                                                                                                                            |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <ul style="list-style-type: none"> <li>This verb is for DB Pool. The target_type for Schema Pool is schaas_pool.</li> <li>Though the properties are not mandated by the verb (because SSA framework owns it), they are needed and without them the pools will not work.</li> <li>Valid property names for Schema Pool are host_credential_guid, database_credential_guid, and gi_credential_guid.</li> </ul>                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

**Table 30–18 (Cont.) Verb Details**

| Verb            | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          | Format                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| update_pool     | <p>Updates a Software Pool.</p> <ul style="list-style-type: none"> <li>name - Name of the existing Software Pool.</li> <li>target_type - Target type of the existing Software Pool. For example: "mwaas_zone" for Middleware Pool, "oracle_cloud_zone" for Database Pool, and "schaas_pool" for Schema Pool.</li> <li>description - Description of the Software Pool.</li> <li>add_members - Comma separated list of targets to be added as members of the Software Pool. The targets to be added must satisfy the member constraints of the Software Pool.</li> <li>remove_members - Member targets to be removed from the Software Pool.</li> <li>placement_constraints - Comma separated key value pairs of the placement constraints that allow the self service administrator to set maximum ceilings for resource utilization. This provides protection for the members of the Software Pool in terms of resource consumption.</li> <li>properties - Comma separated key value pairs for properties that need to be updated based on the Software Pool target type.</li> </ul> | <p>emcli update_pool<br/>-name=&lt;Name of Software Pool&gt;</p> <p>-target_type=&lt;Target type of Software Pool&gt;</p> <p>[-description=&lt;Description of Software Pool&gt;]</p> <p>[-add_members=&lt;Member1, Member2...&gt;]</p> <p>[-remove_members=&lt;Member4, Member5...&gt;]</p> <p>[-placement_constraints=&lt;constraint1=value1,constraint2=value2...&gt;]</p> <p>[-properties=&lt;property1=value1, property2=value2&gt;]</p> |
| delete_pool     | <p>Deletes a Software Pool. Software Pool cannot be deleted if there is an existing Service Template associated with it.</p> <ul style="list-style-type: none"> <li>name - Name of an existing Software Pool.</li> <li>target - Target type of the Software Pool.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         | <p>emcli delete_pool<br/>-name=&lt;Name of Software Pool&gt;</p> <p>-target_type=&lt;Target type of Software Pool&gt;</p>                                                                                                                                                                                                                                                                                                                    |
| get_pool_detail | <p>Retrieves the Software Pool details, such as name, target type, description, PaaS Infrastructure Zone, number of members, placement constraints, and member constraints.</p> <ul style="list-style-type: none"> <li>name - Name of the Software Pool.</li> <li>target - Target type of the Software Pool</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               | <p>emcli get_pool_detail<br/>-name=&lt;Name of Software Pool&gt;</p> <p>-target_type=mwaas_zone</p>                                                                                                                                                                                                                                                                                                                                          |

**Table 30–18 (Cont.) Verb Details**

| Verb                                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                | Format                                                                                                                               |             |               |                             |                                                                                                                                                                                                                                                |
|----------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------|-------------|---------------|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| get_pool_capacity                      | <p>Retrieves the software pool capacity details like CPU utilization, memory allocation, and number of instances per host for the pool entities.</p> <ul style="list-style-type: none"><li>■ name - Name of Software Pool</li><li>■ target_type - Target type of Software Pool</li></ul> <p>Example:</p> <pre>emcli get_pool_capacity       -name=MyDbPool       -target_type=oracle_cloud_zone</pre> <p>Sample Output:</p> <pre>Details of Software Pool "MyDbPool" : Capacity information for Software Pool "MyDbPool" : Capacity information for Entity "slc03qtn.us.oracle.com" of type "host" :       Storage Utilization : 29.0 : 100.0       CPU Utilization : 2.96 : 40.0       Memory Utilization : 27.68 : 70.0       Instances : 2 : 7</pre>                                                                                                                    | <pre>emcli get_pool_capacity       -name=&lt;Name of Software Pool&gt;       -target_type=&lt;Target type of Software Pool&gt;</pre> |             |               |                             |                                                                                                                                                                                                                                                |
| get_pool_allowed_placement_constraints | <p>Retrieves the list of placement constraints for a given pool target type.</p> <p>target_type - Target type of the Software Pool</p> <p>Example:</p> <pre>emcli get_pool_allowed_placement_constraints       -target_type=oracle_cloud_zone</pre> <p>Sample Output:</p> <table><thead><tr><th>Name</th><th>Description</th></tr></thead><tbody><tr><td>MAX_INSTANCES</td><td>Maximum Number of Instances</td></tr></tbody></table>                                                                                                                                                                                                                                                                                                                                                                                                                                       | Name                                                                                                                                 | Description | MAX_INSTANCES | Maximum Number of Instances | <pre>emcli get_pool_allowed_placement_constraints       -target_type=&lt;mwaas_zone&gt;</pre>                                                                                                                                                  |
| Name                                   | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |                                                                                                                                      |             |               |                             |                                                                                                                                                                                                                                                |
| MAX_INSTANCES                          | Maximum Number of Instances                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |                                                                                                                                      |             |               |                             |                                                                                                                                                                                                                                                |
| get_pool_filtered_targets              | <p>Retrieves the filtered targets available for Software Pool creation based on the criteria passed.</p> <ul style="list-style-type: none"><li>■ target_type - Target type of the Software Pool</li><li>■ paas_zone - Name of PaaS Infrastructure Zone within which the filtered targets are to be retrieved</li><li>■ member_constraints - Comma separated key value pairs that restricts the addition of member targets to a Software Pool with a set criteria</li></ul> <p>Example:</p> <pre>emcli get_pool_filtered_targets       -target_type=oracle_cloud_zone       -paas_zone=MyPaaSZone       -member_constraints="CONFIGURATION=oracle_ database,VERSION=11.2.0.3.0,PLATFORM=226"</pre> <p>Sample Output:</p> <table><thead><tr><th>Target Type</th><th>Target Name</th></tr></thead><tbody><tr><td>oracle_home</td><td>OraDb11g_home1</td></tr></tbody></table> | Target Type                                                                                                                          | Target Name | oracle_home   | OraDb11g_home1              | <pre>emcli get_pool_filtered_targets       -target_type=&lt;Target type of Software Pool&gt;       -paas_zone=&lt;Paas Infrastructure Zone of Software Pool&gt;       -member_constraints=&lt;constraint1=value1, constraint2=value2&gt;</pre> |
| Target Type                            | Target Name                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |                                                                                                                                      |             |               |                             |                                                                                                                                                                                                                                                |
| oracle_home                            | OraDb11g_home1                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |                                                                                                                                      |             |               |                             |                                                                                                                                                                                                                                                |

**Table 30–18 (Cont.) Verb Details**

| Verb                                | Description                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Format                                                                                |
|-------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------|
| get_pool_allowed_member_constraints | Retrieves the list of allowed Software Pool member constraint values for creation of Software Pool.<br>target_type - Target type of the Software Pool.<br>Example:<br>emcli get_pool_allowed_member_constraints -target_type=oracle_cloud_zone                                                                                                                                                                                                                                                                                                                                                                                                               | emcli get_pool_allowed_member_constraints -target_type=<Target type of Software Pool> |
| get_named_credential                | Retrieves credential information used to create DBaaS and SaaS Pools.<br><ul style="list-style-type: none"> <li>cred_name - The credential name</li> <li>cred_owner - Owner of the credential</li> </ul> Example:<br>emcli get_named_credential -cred_name=BBANTHIA -cred_owner=ORACLE<br><br>Sample Output:<br>Output:<br>Credential Name:BBANTHIA<br>Credential Owner:ORACLE<br>Credential Type:HostCreds<br>Credential Target Type:host<br>Credential Username:bbanthia<br>Credential Scope:global<br>Credential Guid:DBA449B8967AAF77E040F00A73B11F55<br>Credential Stripe:TARGETS<br>Credential Columns:<br>HostPassword=*****<br>HostUserName=bbanthia | emcli get_named_credential -cred_name=<name> -cred_owner=<owner>                      |

---

## Introduction to Blueprints

This chapter introduces the cloud blueprints (referred to as *blueprints* in this chapter) for the Oracle Private Cloud and describes how to use them and create your own blueprints. It contains the following sections:

- [Introduction to Blueprints](#)
- [Uses of the Blueprint](#)
- [Enterprise Manager Cloud Concepts](#)
- [Blueprint Concepts](#)
- [Deploying a Blueprint](#)
- [Blueprint Deployment Processing](#)
- [Blueprint Examples](#)
- [Conclusion](#)

### 31.1 Introduction to Blueprints

Blueprints are used to describe a desired set of inter-related cloud resources. Like architectural blueprints, they describe what you want including how they are configured to interact with each other, but not how to build them. For instance, a blueprint does not describe the order in which to create the components. Rather, the blueprint orchestration logic figures that out based on inter-resource dependencies.

As an example, suppose you want to create a set of interacting cloud resources such as a WebLogic server instance, an application, and a database. To do so, you must first create the database and WebLogic server instance, deploy the application, and create a JEE datasource that is to be used by the WebLogic server to connect to the database.

You could perform all these operations manually, through the Enterprise Manager Cloud Self Service Portal. You would request creation of the WebLogic server and database and wait for either to complete. Periodically, you would check the status of the creation requests. Once the WebLogic server is created, you could deploy the application. When both the WebLogic server and database are created, you could create the JEE datasource.

Alternatively, you can use a blueprint that describes the four cloud resources to automate the process. To do so, you request instantiation of the blueprint and provide any input parameter values required by the blueprint. The blueprint initiates the creation of the resources and monitors the creation process to ensure that the dependent resources are automatically created as soon as the required resources are created.

The rest of this chapter introduces the blueprint concepts including how to deploy an existing blueprint as well as how to write your own. You can also refer to the [Chapter 32, "Cloud Blueprints and Blueprint Reference"](#) for more details on blueprints.

## 31.2 Uses of the Blueprint

A blueprint can be used to automate the creation of service instances. An EM\_SSA\_USER can use blueprints for various reasons:

- To create an application composed of several service instances and related cloud resources.
- To create such sets of instances several times.
- To facilitate instance creation for other EM\_SSA\_USERS.
- To eliminate the manual interactive steps that would otherwise be needed to create the set of instances
- To create a textual representation that can be reviewed and modified by other users.

To summarize, blueprints allow you to automate the creation of a set of service instances.

For example, the Quality Assurance team in an enterprise needs to allocate and release resources required to test a Web application. Instead of manually creating the service instances using the Enterprise Manager Cloud Self Service application, a blueprint can be used to perform this task. One person authors a blueprint so that all QA engineers can simply invoke the blueprint and enter a few input parameter values, after which the resources are created. Each user can watch as the blueprint processor displays the status for creation of each resource.

Another example illustrates the use of a blueprint to address simplicity and consistency concerns. An IT shop has a service template that accepts 8 input parameters. For a specific group of users, the same set of values should be used for 6 of those 8 parameters. A simple blueprint accepts 2 parameters and uses the template to instantiate the instances with the other 6 parameters consistently defined.

## 31.3 Enterprise Manager Cloud Concepts

The concepts described in this section are Enterprise Manager cloud concepts. They are not introduced as part of blueprints but are concepts used by blueprints. Since they form the basis for blueprints, they are summarized below.

- **Oracle Cloud API:** The Oracle Cloud API (see [Chapter 24, "Introduction to Cloud APIs"](#)) defines a RESTful programming interface to consumers of IaaS, MWaaS, and DBaaS based on Oracle's solution stack. It is the Oracle Cloud API that the blueprint processor uses to create cloud resources based on what's specified by the blueprint.
- **Oracle Cloud Resource Model:** The Oracle Cloud Resource Model (see [Chapter 25, "Cloud Resource Models"](#)) details the types of resources one can manipulate via the Oracle Cloud API and, for each type, both its attributes and the operations that can be performed on instances of a resource type.
- **Cloud Resource Types and Attributes:** The cloud resource model specifies a set of attributes that are common to all cloud resources, such as uri (its URI) and resource\_status (with values such as READY and CREATING). The model also



defines a set of cloud resource types and their attributes. Some resource types discussed later in this document are:

- **DbPlatformInstance:** A database platform instance is created using a template. It has attributes such as zone (location for DBaaS instance) and params,(to specify username and password).
- **JavaPlatformInstance:** A JEE server instance is also created using a template. Similar to DbPlatformInstance, a JavaPlatformInstance has attributes like zone. An example of a MWaaS-specific attribute is `application_instance_deployments`, that identifies all applications deployed to the instance.
- **ApplicationInstanceDeployment:** A resource type that represents an application deployment to a JavaPlatformInstance.
- **Datasource:** A resource type that represents a datasource of a JavaPlatformInstance. It is contained in a JavaPlatformInstance and refers to a DatabaseInstance.
- **Resource Containment:** A cloud resource can contain other resources. For instance, a Datasource of a JavaPlatformInstance is contained in the JavaPlatformInstance, and a template contains all service instances that were created using that template.

Every resource you create will be contained in a parent cloud resource. As part each resource definition in your blueprints, you will identify the resource's container.

## 31.4 Blueprint Concepts

A blueprint generally describes one or more cloud resources to be created. A user *deploys* a blueprint to create the resources described by the blueprint, at which time he provides any input parameter values used by the blueprint.

To create each resource, the blueprint specifies its attribute values, which may be hard-coded, come from blueprint input parameters, or gleaned from other resources. In instantiating all the resources, the blueprint system determines the resource dependencies and uses this to order the resource creation and/or configuration required to properly create the instances described by a blueprint.

- **Input Parameters:** A blueprint can define **input parameters**. Each parameter definition specifies a name, datatype and optional specifications such as default value. A user who deploys a blueprint must specify the value for all parameters not having a default.
- **Resource:** A blueprint **resource** defines how to construct a cloud resource. It specifies a set of attributes and the parent cloud resource that is to contain the newly created resource. Each resource in a blueprint also has a name, which must be unique within a blueprint.
- **Outputs:** If outputs are specified, the specified values are displayed when blueprint deployment completes. Output parameters can be used to display information derived during deployment such as to inform the end user of the URL of a JEE application deployment.
- **Intrinsic Functions:** The blueprint language includes a set of predefined functions, illustrated below.
- **Named Literals:** As a programming convenience, blueprints can include a Data section. This is commonly used to specify symbolic names for literal values.

- **Dependencies:** Blueprint deployment is done by creating resources in parallel when possible, but a resource that depends on another resource can't be created until the latter resource is created. Such dependencies are often implicit, but blueprint processing identifies dependencies and orchestrates the overall steps.

## 31.5 Deploying a Blueprint

### Prerequisites:

- You have the blueprint file.
- You must have installed the blueprint processor software.

A blueprint file is a text file in which the author has used the blueprint language to describe what is required. You deploy the blueprint by running the blueprint processor. For instance, on Windows you would use a command as follows:

```
bp_processor.py myfile.yml -u myname -c https://myhost/em/cloud
```

In this example, the `-u` option specifies the Enterprise Manager user id. If the password is not specified here, you are prompted for it when you execute the blueprint file. The `-c` option is used to specify the cloud URL.

There are numerous other command line options. To view a description of supported options, enter

```
bp_processor.py deploy -h
```

## 31.6 Blueprint Deployment Processing

When a blueprint is deployed, the runtime logic processes the input parameters and orchestrates the creation of resources, doing so in parallel when possible. It also monitors progress and keeps the user informed.

The output you see will depend on the blueprint and your environment. The example below shows the use of a blueprint that deploys a Weblogic server, application, database, and JEE Datasource. The output from running this blueprint should give you a sense for the blueprint processing steps.

```
C:\Users\myname\Dropbox\Code\blueprints>bp_processor.py xyzApp.yml -c
https://... -g deployment_report
...
Cloud user id: ssa_user1
Password:
```

The command you entered above specifies the cloud URL (via the `-c` option) and a directory into which to place the optional deployment report (the `-g` option). You then entered the credentials as prompted to proceed.

### Blueprint Processor - Invocation Summary

```

Cloud URI: https://...
User: ssa_user1
Blueprint file: xyzApp.yml
Timeout: 90 minutes, 0 seconds
Refresh frequency: 15 seconds
Inputs:
Pause points: (none)
Debug logging: False
Instance name: default_instance_name
```

```

Graphical report dir: deployment_report
Versions:
 Blueprint processor: 12.1.0.5, 10-Oct-2012
 Cloud protocol: 10001

```

```
14:24:59 INFO: Connecting to cloud: https://....
```

**Note:** Most of the information in the Invocation Summary reflects default values used because you have not specified the corresponding options.

Because the blueprint you are using defines input parameters, you are then prompted to provide values:

```

Input Parameter Value Entry

Zone to use for db (Zone1):
Password to use for db (welcome1):

```

In this example, you pressed **Enter** to accept the default value for zone, Zone1 and then entered a password.

Once all the necessary information has been provided, the blueprint processor starts creating the resources and monitoring the progress. This is depicted in a “vertical timeline” so you can observe the progress. Each resource to be created is represented by a column. As the state transitions for each resource occur, they are noted in the corresponding column of the vertical timeline.

First, the list of all possible states, with their abbreviations, is printed. The states reflect the processing phases and the outcome of each phase.

```

14:25:03 INFO:
14:25:03 INFO: Resource State Timeline
14:25:03 INFO: -----
14:25:03 INFO: State Key:
14:25:03 INFO: e : Evaluating
14:25:03 INFO: ep: Evaluation pending. (See right side for pendee)
14:25:03 INFO: es: Evaluation succeeded, creation requested
14:25:03 INFO: EF: Evaluation failed
14:25:03 INFO: c : Creating
14:25:03 INFO: CF: Creation failed
14:25:03 INFO: CS: Creation succeeded. State = READY
14:25:03 INFO:

```

To process any resource definition, the first step is to evaluate the expressions of the definition that describe the resource (State Key: e). In some cases, evaluation must be delayed (State Key: ep). Once fully evaluated, a creation request is made (State Key: c) and the processing for that resource is successfully completed (State Key: CS). After the State Key is printed, the resource creation timeline appears:

```

14:25:03 INFO:
14:25:03 INFO: MyApp
14:25:03 INFO: /
14:25:03 INFO: / MyDB
14:25:03 INFO: / / MyDS
14:25:03 INFO: / / / MyWebServer
14:25:03 INFO: / / / /
14:25:03 INFO: -----
14:25:03 INFO: | | | | e |
14:25:04 INFO: | | | | es |

```

```

14:25:10 INFO: | | | | c |
14:25:10 INFO: | | e | | . |
14:25:12 INFO: | | es | | . |
14:25:17 INFO: | | c | | . |
14:25:17 INFO: | | . | e | . |
14:25:17 INFO: | | . | ep | . | Awaiting creation of MyWebServer
14:25:17 INFO: | e | . | | . |
14:25:17 INFO: | ep | . | | . | Awaiting creation of MyWebServer
14:25:33 INFO: | | . | | . |
14:25:50 INFO: | | . | | . |
.....
14:41:18 INFO: | | . | | . |
14:41:33 INFO: | | . | | CS |
14:41:33 INFO: | | . | | ==== |
14:41:44 INFO: | | . | ep | | Awaiting creation of MyDB
14:41:57 INFO: | es | . | | |
14:42:00 INFO: | c | . | | |
14:42:17 INFO: | . | . | | |
....
14:44:44 INFO: | . | . | | |
14:45:16 INFO: | CS | . | | |
14:45:16 INFO: | ==== | . | | |
14:45:32 INFO: | | . | | |
14:45:47 INFO: | | . | | |
14:54:17 INFO: | | . | | |
14:54:17 INFO: | | CS | | |
14:54:17 INFO: | | ==== | | |
14:54:18 INFO: | | | es | |
14:54:19 INFO: | | | c | |
14:54:33 INFO: | | | CS | |
14:54:33 INFO: | | | ==== | |
14:54:33 INFO: -----
14:54:33 INFO:
14:54:33 INFO:

```

You can see that the timeline is vertical and the four resources are MyApp, MyDB, MyDS, and MyWebServer.

The MyWebServer resource definition is evaluated (state key: e) first. When this is successful, the creation process begins (state key: c). Parallely, the MyDB resource is evaluated and the creation process is initiated. Then the MyDS resource is evaluated and the blueprint processor determines that the evaluation cannot be completed until the MyWebServer resource is created. The same process is applicable to the MyApp resource. When the MyDB and MyWebServer resources are successfully created (state key: CS), the creation process for MyApp and MyDS resources can proceed.

Next, the Outputs section of our example blueprint is processed:

```

14:54:33 INFO: Output Processing
14:54:33 INFO: -----
14:54:33 INFO:
14:54:33 INFO: Output values specified: 1
14:54:49 INFO: Value of URL: {u'ms_1': u'http://...}
14:54:49 INFO:

```

In the example above, you can see that the blueprint specifies one output value named "URL" and a value is represented with `https://....`

When all the resources have been successfully created, the blueprint processor summarizes the results. This includes the processing summary for each requested resource as well as the timing information for each resource and the overall run:

```

14:54:49 INFO: Blueprint Processing Summary
14:54:49 INFO: -----
14:54:49 INFO:
14:54:49 INFO: Resource State Summary:
14:54:49 INFO: MyWebServer: READY
14:54:49 INFO: URI: /em/cloud/jaas/javaplatforminstancerequest/163
14:54:49 INFO: Cloud resource state: READY
14:54:49 INFO: Timing info:
14:54:49 INFO: Creation start: 14:25:04
14:54:49 INFO: Creation end: 14:41:33
14:54:49 INFO: Duration: 16 minutes, 29.6 seconds
14:54:49 INFO: MyDB: READY
14:54:49 INFO: URI: /em/cloud/dbaas/dbplatforminstance/byrequest/164
14:54:49 INFO: Cloud resource state: READY
14:54:49 INFO: Timing info:
14:54:49 INFO: Creation start: 14:25:12
14:54:49 INFO: Creation end: 14:54:17
14:54:49 INFO: Duration: 29 minutes, 5.7 seconds
14:54:49 INFO: MyDS: READY
14:54:49 INFO: URI: /em/cloud/jaas/datasourcerequest/QA_app_DS@201
14:54:49 INFO: Cloud resource state: READY
14:54:49 INFO: Timing info:
14:54:49 INFO: Creation start: 14:54:18
14:54:49 INFO: Creation end: 14:54:33
14:54:49 INFO: Duration: 0 minutes, 15.3 seconds
14:54:49 INFO: MyApp: READY
14:54:49 INFO: URI:
/em/cloud/jaas/applicationinstancedeploymentrequest/myApp@181
14:54:49 INFO: Cloud resource state: READY
14:54:49 INFO: Timing info:
14:54:49 INFO: Creation start: 14:41:57
14:54:49 INFO: Creation end: 14:45:16
14:54:49 INFO: Duration: 3 minutes, 18.9 seconds
14:54:49 INFO:
14:54:49 INFO: Timing Summary (seconds):
14:54:49 INFO: Client-side CPU time: 0 minutes, 6.474 seconds
14:54:49 INFO: Elapsed time:
14:54:49 INFO: Processing time: 29 minutes, 58.0 seconds
14:54:49 INFO: Paused time: 0 minutes, 1.9 seconds
14:54:49 INFO: Total elapsed time: 29 minutes, 59.9 seconds
14:54:49 INFO:
14:54:49 INFO: Graphical Report Generation
14:54:49 INFO: -----
14:54:49 INFO:
14:55:37 INFO: Graphical report generated: deployment_report/bp_report.html

C:\Users\myname\Dropbox\Code\blueprints>

```

## 31.7 Blueprint Examples

This section illustrates the use of the blueprint concepts and syntax by guiding the reader through progressively more complex examples of blueprints. If you do not plan to author any blueprints, you can skip this section.

### 31.7.1 Blueprint Structure and Basics

A cloud blueprint specifies a set of desired cloud resources and represents these resources via a text file. Blueprints leverage a standard for easily readable data-structured text called YAML.

YAML is a standard notation, like XML and JSON. As with JSON, YAML is used to represent information via lists, dictionaries, and nesting. These concepts are sufficient to capture all blueprint semantics.

### 31.7.1.1 Simple Blueprint

A blueprint is a document that you can think of as containing sections. The simplest useful blueprint specifies only the **Resources** section and a single resource. In this example, the resource is a database defined by a template.

```
Resources
MyDB:
 Type: application/oracle.com.cloud.common.DbPlatformInstance+json
 Container: ... (refers to db template)
 Properties: ... (provides properties of db)
```

The above blueprint defines one blueprint resource named MyDB. The Type entry specifies the media type for 'database' as defined by the Cloud Resource Model API. The Container entry identifies the parent cloud resource to contain the newly created object. (Per [Chapter 25, "Cloud Resource Models"](#), all cloud resources are created by adding them to existing containers.) Being a database service instance, it will be created via the database template used to create the service. How to specify the container will be shown later, so we just use ellipses here. Similarly, the data required by that template is specified in the Properties entry and shown later.

Now let us add an 'Inputs' section...

```
Inputs:
 DbZone:
 Type: String
 DefaultValue: Zone1
 Prompt: Zone to use for db
 DbPassword:
 Type: String
 DefaultValue: welcome1
 Prompt: Password to use for db
 Sensitive: True
Resources:
 MyDB:
 Container: ...
 Properties: ...
```

When the user requests deployment of this blueprint, he provides a value for the DbZone parameter or takes the default specified by the blueprint (Zone1). The value of user provided input parameters can be used by other parts of the blueprint, in particular to provide attribute values needed to create resources. More on this later.

### 31.7.1.2 Simple Resource: Database Service Instance

In a blueprint, one uses blueprint resources to describe the cloud resources to create. For each blueprint resource, the information required by the Cloud Resource Model is provided.

The following blueprint specifies a single resource to create a database service instance:

```
MyDB:
 Type: application/oracle.com.cloud.common.DbPlatformInstance+json
 Container: ...
 Properties:
 zone: ...
```

```

name: jdbName
params:
 username: app_user
 password: change_me

```

Each resource definition in a blueprint specifies a name, a Container into which to add the resource, and Properties used to specify the characteristics of what to create. In this case:

- The **name** of the blueprint resource definition is MyDB. The name is used in the scope of a blueprint, e.g. to inform the user deploying a blueprint about progress for each resource. In more complex cases, we will see that the name can be referenced elsewhere within a blueprint.
- The **Container** entry specifies the URI of the container to which the new resource will be added. To create a database service instance, we identify the service template that corresponds to the kind of database we want. (We'll see how to do that when we introduce the topic of intrinsic functions.)
- The **Properties** entry specifies values needed to create the resource. In this case, the model requires that we specify zone, name, and params properties. These specify the zone in which the instance is to be created, its name, and a list of name/value pairs required by the selected template.

### 31.7.1.3 Intrinsic Functions

To operate on data, blueprints support the use of intrinsic functions. All function names begin with "f\_" and are invoked with a list of arguments.

Continuing the above example, we use two intrinsic functions, to return the URI of the desired container and the desired zone.

Resources:

MyDB:

Type: application/oracle.com.cloud.common.DbPlatformInstance+json

Container:

**f\_getTemplateURI:**

- Small Database Service Template\_automation\_VIMAL\_si
- dbaas

Properties:

zone:

**f\_getZoneURI:** ...

name: jdbName

params:

username: app\_user

password: ...

As you can see, the `f_getTemplateURI` function takes 2 arguments, the name of the template and its service type. The current supported service types are: dbaas, jaas, and iaas (Database-, Java-, and Infrastructure-as-a-Service). The `f_getZoneURI` is analogous to the `f_getTemplateURI` but for zones.

Other intrinsic functions will be introduced below. The full set of functions is described in [Chapter 32, "Cloud Blueprints and Blueprint Reference"](#).

### 31.7.1.4 Simple Resource with Parameter

To the above example blueprint, we now add the use of 2 parameters. By doing so, the user who deploys the blueprint, can specify which zone and password to use.

Inputs:

```
DbZone:
 Type: String
 DefaultValue: Zone1
 Prompt: Zone to use for db
DbPassword:
 Type: String
 DefaultValue: welcome1
 Prompt: Password to use for db
 Sensitive: True
Resources:
 MyDB:
 Type: application/oracle.com.cloud.common.DbPlatformInstance+json
 Container:
 f_getTemplateURI:
 - Simple DB Template
 - dbaas
 Properties:
 zone:
 f_getZoneURI:
 - f_path:
 - "Inputs.DbZone.Value"
 - dbaas
 params:
 username: app_user
 password:
 f_path:
 - 'Inputs.Password.Value'
```

The Inputs section defines the two input parameters and the values of the parameters are accessed via the `f_path` intrinsic function.

The `f_path` function is used to evaluate path expressions to access any data in your blueprint as well as any cloud resource data to which you have access. In our example, the path expression just uses the dot operator to access nested attributes, first access the Inputs attribute (Inputs section) of the blueprint and within that, the UserId attribute, and within that the Value attribute.

#### 31.7.1.5 Data Section (Named Literals)

Suppose your blueprint creates several databases and suppose that you do not want to prompt the user for the username and password. Furthermore, you want to code your blueprint so that it is easy to change the password later.

In a procedural language, you would use a named literal in order to document the intent and so that you can change it once at the top of your code. Within a blueprint, you do this by using the Data section.

```
Data:
 QADBCreds:
 user: sysman
 password: sysman
Resources:
 MyDB1:
 Type: application/oracle.com.cloud.common.DbPlatformInstance+json
 Container:
 f_getTemplateURI:
 - Small Database Service Template_automation_VIMAL_si
 - dbaas
 Properties:
 zone:
 f_getZoneURI:
```



```

 - f_path:
 - "Inputs.DbZone.Value"
 - dbaas
 params:
 username:
 f_path:
 - "Data.QADBCreds.user"
 password:
 f_path:
 - "Data.QADBCreds.password"
 name: jbName

MyDB2:
 Container:
 ...

```

In the above case, you can see that the Data section takes a YAML structure, which can be traversed via the 'path' function, in the same way shown for Inputs earlier.

### 31.7.2 Putting It All Together – Multiple Interdependent Resources

In this more elaborate example, we show how one might create a database and an application that uses it. To do so, the blueprint specifies four cloud resources:

- Database service instance
- Java service instance
- Datasource of the Java service instance
- Application of the Java service instance

New constructs are highlighted below

```

Inputs:
 DbZone:
 Type: String
 DefaultValue: Zone1
 Prompt: Zone to use for db
 DbPassword:
 Type: String
 DefaultValue: welcome1
 Prompt: Password to use for db
 Sensitive: True
Macros:
 # Return a name with unique (date-time) suffix
 # The one argument is a 'name' string
 f_myDescriptiveName:
 - 1
 - f_concat:
 - arg_1
 - '_'
 - f_path:
 - f_path:
 - 'Info.time_suffix'
Resources:
 MyDB:
 Type: application/oracle.com.cloud.common.DbPlatformInstance+json
 Container:
 f_getTemplateURI:
 - Small Database Service Template_automation_VIMAL_si
 - dbaas

```

```
Properties:
 zone:
 f_getZoneURI:
 - f_path:
 - "Inputs.DbZone.Value"
 - dbaas
 params:
 username: app_user
 password:
 f_path:
 - "Inputs.DbPassword.Value"
 name: jbName
MyWebServer:
 Container:
 f_getTemplateURI:
 - PS4_LowHeapTemplate
 - jaas
 Properties:
 name:
 f_myDescriptiveName:
 - jb_pf
 zone:
 f_getZoneURI:
 - Zone1
 - jaas
MyDS:
 Type: application/oracle.com.cloud.jaas.DataSource
 Container:
 f_getResourceURI:
 - MyWebServer
 Properties:
 name: QA_app_DS
 jndi_name:
 - jndi_1
 - jndi_2
 jdbc_driver: oracle.jdbc.OracleDriver
 database_type: Oracle
 database_connect_string:
 f_concat:
 - 'jdbc:oracle:thin:@'
 - f_getResourceAttr:
 - MyDB
 - connect_string
 username: app_user
 password:
 f_path:
 - "Inputs.DbPassword.Value"
MyApp:
 Type: application/oracle.com.cloud.jaas.ApplicationInstanceDeployment
 Container:
 f_getResourceURI:
 - MyWebServer
 Properties:
 application_instance_component:
 f_getAppCompURI:
 - jbcomponent
 - SSA_USER1
 -
 name: myApp
Outputs:
```

```

URL:
 Description: URL of the deployed app
 Value:
 f_getResourceAttr:
 - MyApp
 - http_application_invocation_url

```

### 31.7.2.1 Macro Section

If you have a sequence of constructs that you tend to repeat, you can use macro expansion to improve the readability of your blueprint. Macros also enable you to encapsulate logic, for instance, you can modify the logic in one place to affect all code that refers to it.

Our example blueprint defines a macro named `f_myDescriptiveName`. It takes one string parameter and appends “\_” as well as a string representation of the current time.

### 31.7.2.2 Attributes of Created Resources (Dependencies)

The key new feature introduced by this example is the ability to refer to attributes of created resources. For instance, the Cloud Resource Model for Datasource defines a `database_connect_string` attribute whose value is required to create a Datasource. The contents of the string will not be available until after the database is created, so it is clearly not something the blueprint author can know in advance. Instead, he uses an intrinsic function to refer to the needed property of the newly created database. (Each `DatabasePlatformInstance` exposes a ‘connect\_string’ property.)

To do this, the Datasource specifies the value of `JDBCConnectString` via the ‘`f_getResourceAttr`’ intrinsic function:

```

MyDS:
 ...
 Properties:
 ...
 database_connect_string:
 f_concat:
 - 'jdbc:oracle:thin:@'
 - f_getResourceAttr:
 - MyDB
 - connect_string

```

In this case, the `f_getResourceAttr` function waits for the `MyDB` resource to be created and then returns the value of its `connect_string` property.

Similarly, the creation of a Datasource is done by adding it to the `JavaPlatformInstance` that is created first, so we need to refer to the `JavaPlatformInstance`’s URI.

```

MyDS:
 ...
 Container:
 f_getResourceURI:
 - MyWebServer

```

In this case, the `f_getResourceURI` function waits for the `MyWebServer` resource to be created and then returns its URI. (In addition to adding a Datasource to `MyWebServer`, we also must add an `ApplicationInstanceDeployment`, so the same approach is used for both.)

**Note:** Though the `f_path` function can be used to achieve the same effect, the `f_getResourceAttr` knows to wait for the resource creation to succeed and its state be `READY` before attempting to get its attribute.

More generally, blueprint resources can refer to other resources and blueprint orchestration accounts for such dependencies, creating resources in parallel when possible.

### 31.7.3 Visual Depiction of Blueprint Processing

The blueprint processor can also generate an HTML report that includes a graphical representation of the blueprint. This may be used to help understand the overall structure of the blueprint and the relationship of blueprint entities. In addition, the report can include the results of deploying the blueprint.

The following example report is for the example blueprint described in [Section 31.6, "Blueprint Deployment Processing"](#).

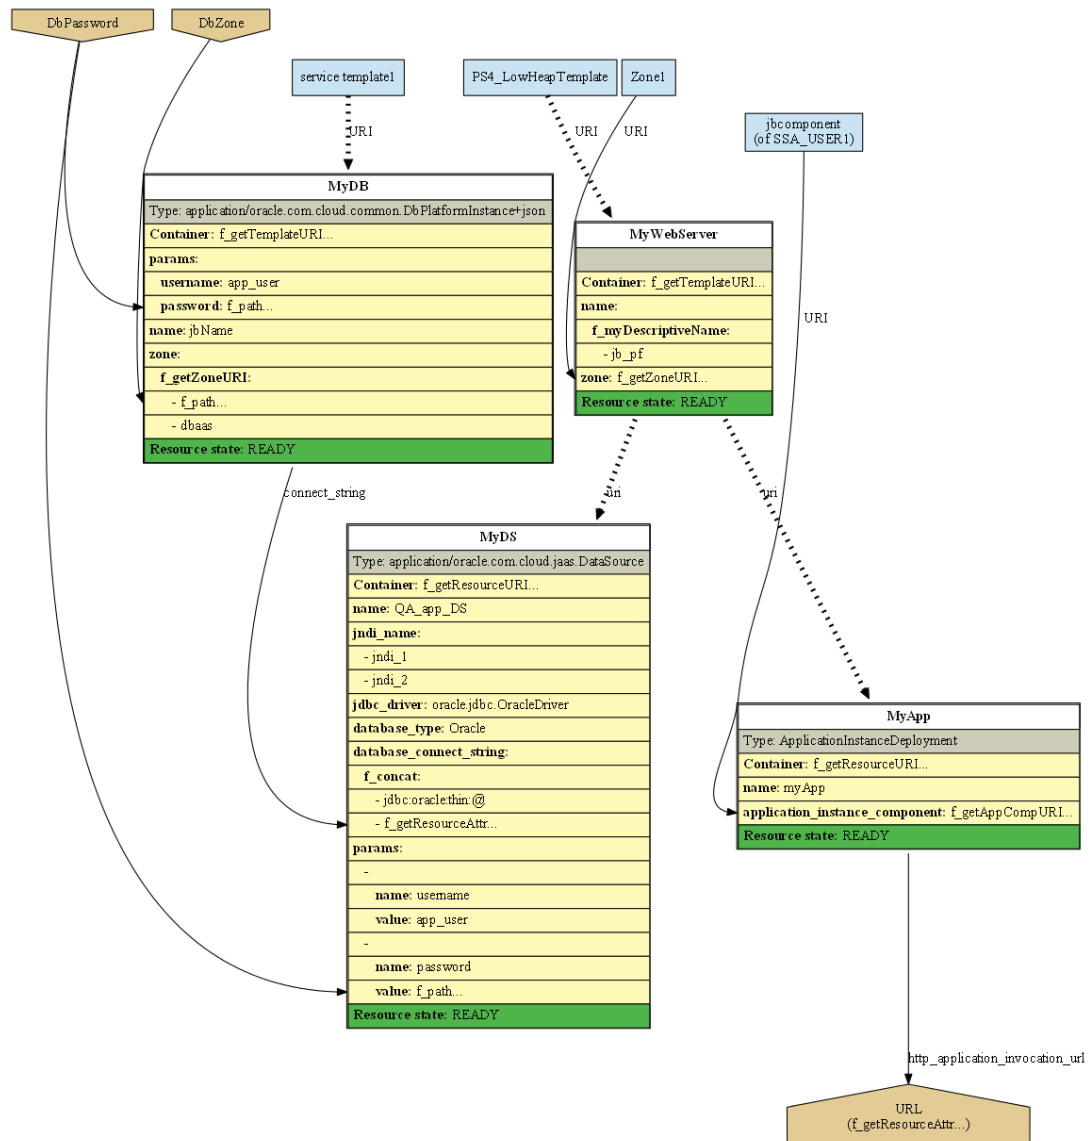
The first part of the report summarizes the run:

**Figure 31–1** *Blueprint Report Title*



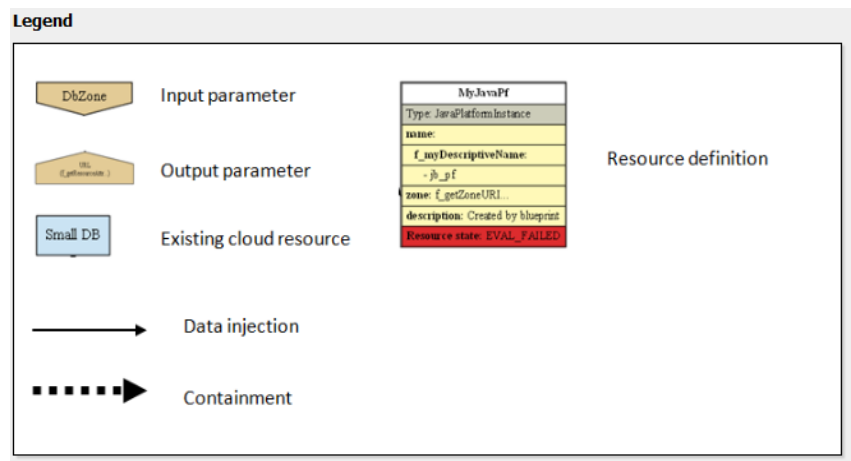
Next is the graphical depiction, in which two types of arcs are used. One depicts containment. For example, the containment of a datasource within a WebLogic server. The other depicts how data is used across the elements of a blueprint. If blueprint deployment is successful, the status of all resources will be Ready (green) and you can click on a resource, which links to another report section where you can view the values of its attributes at the time of creation.

Figure 31-2 Blueprint Graphical Depiction



This is followed by a legend to explain the graphical conventions:

Figure 31–3 *Blueprint Legend*



Finally, each created resource is summarized. (Clicking on a resource definition in the graphical depiction takes you directly to the resource summary for the selected resource.)

Figure 31–4 *Blueprint Created Cloud Resources - I*

**Created Cloud Resources (and their current state)**

**MyApp (Type: ApplicationInstanceDeployment)**

| Attribute | Name                             | Value                                                                                                                                                                                                              |
|-----------|----------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|           | uri                              | /em/cloud/jaas/applicationinstancedeployment/C03DC24FDC8A36E37BEC9229A88CF08B                                                                                                                                      |
|           | contained_in                     | {u'status': u'RUNNING', u'media_type': u'application/oracle.com.cloud.jaas.JavaPlatformInstance+json', u'uri': u'/em/cloud/jaas/javaplatforminstance/67A4827A18B5159E2FDEFF4294F28F4F', u'name': u'jb_pf_14_25_1'} |
|           | context_id                       | C03DC24FDC8A36E37BEC9229A88CF08B                                                                                                                                                                                   |
|           | http_application_invocation_url  | {u'ms_1': u'http://slc01rbw.us.oracle.com:21307/derby'}                                                                                                                                                            |
|           | https_application_invocation_url | {u'ms_1': u'https://slc01rbw.us.oracle.com:32893/derby'}                                                                                                                                                           |
|           | media_type                       | application/oracle.com.cloud.jaas.ApplicationInstanceDeployment+json                                                                                                                                               |
|           | name                             | myApp                                                                                                                                                                                                              |
|           | resource_state                   | {u'state': u'READY'}                                                                                                                                                                                               |
|           | service_family_type              | jaas                                                                                                                                                                                                               |
|           | status                           | RUNNING                                                                                                                                                                                                            |

**MyDB (Type: application/oracle.com.cloud.common.DbPlatformInstance+json)**

| Attribute | Name             | Value                                                                                                                      |
|-----------|------------------|----------------------------------------------------------------------------------------------------------------------------|
|           | uri              | /em/cloud/dbaas/dbplatforminstance/byrequest/164                                                                           |
|           | available_space  | 0.19140625                                                                                                                 |
|           | based_on         | /em/cloud/dbaas/dbplatformtemplate/D0C4E06F54AA4740E040F20A4C1B51A6                                                        |
|           | connect_string   | (DESCRIPTION=(ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=slc01rbw.us.oracle.com)(PORT=1522)))(CONNECT_DATA=(SID=sl54f90d))) |
|           | context_id       | 164                                                                                                                        |
|           | created          | 2012-12-17 11:53:20                                                                                                        |
|           | database_type    | oracle_database                                                                                                            |
|           | db_version       | 11.2.0.1.0                                                                                                                 |
|           | destination_zone | /em/cloud/dbaas/zone/A1B44A4EBCC456312SD9D0A3AAE4FD51                                                                      |
|           | last_backup      | null                                                                                                                       |
|           | load             | 0.0310767354289456                                                                                                         |
|           | master_username  | app_user                                                                                                                   |
|           | media_type       | application/oracle.com.cloud.common.DbPlatformInstance+json                                                                |
|           | name             | [REDACTED]                                                                                                                 |
|           | resource_state   | {u'state': u'READY'}                                                                                                       |
|           | status           | RUNNING                                                                                                                    |
|           | total_sessions   | 31                                                                                                                         |
|           | total_sga        | 1019.41796875                                                                                                              |
|           | up_time          | 0                                                                                                                          |

**Figure 31–5 Blueprint Created Cloud Resources - II**

```

MyDS (Type: application/oracle.com.cloud.jaas.DataSource)

Attribute Name Value
uri /em/cloud/jaas/datasource/QA_app_DS@67A4827A1BB5159E2FDEFF4294F28F4F
contained_in {u'media_type': u'application/oracle.com.cloud.jaas.JavaPlatformInstance+json', u'uri': u'/em/cloud/jaas/javaplatforminstance/67A4827A1BB5159E2FDEFF4294F28F4F', u'name': u'jb_pf_14_25_1', u'context_id': u'67A4827A1BB5159E2FDEFF4294F28F4F'}
context_id QA_app_DS@67A4827A1BB5159E2FDEFF4294F28F4F
database_connect_string jdbc:oracle:thin:@(DESCRIPTION=(ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=██████████)(PORT=1522)))
(CONNECT_DATA=(SID=si54f90d)))
jdbc_driver oracle.jdbc.OracleDriver
jndi_name [u'jndi_1']
media_type application/oracle.com.cloud.jaas.DataSource+json
name QA_app_DS
resource_state (u'state': u'READY')
service_family_type jaas

MyWebServer (Type: application/oracle.com.cloud.jaas.JavaPlatformInstance+json)

Attribute Name Value
uri /em/cloud/jaas/javaplatforminstance/67A4827A1BB5159E2FDEFF4294F28F4F
application_instance_deployments {u'media_type': u'application/oracle.com.cloud.jaas.ApplicationInstanceDeployment+json', u'total': u'1', u'elements': [{u'status': u'RUNNING', u'name': u'myApp', u'service_family_type': u'jaas', u'context_id': u'C03DC24FDC8A36E37BEC9229A88CF08B', u'uri': u'/em/cloud/jaas/applicationinstancedeployment/C03DC24FDC8A36E37BEC9229A88CF08B', u'media_type': u'application/oracle.com.cloud.jaas.ApplicationInstanceDeployment+json'}]}
based_on {u'service_family_type': u'jaas', u'context_id': u'D0B2A18D85E73A7EE040F20A4C1B696Z', u'uri': u'/em/cloud/jaas/javaplatformtemplate/D0B2A18D85E73A7EE040F20A4C1B696Z', u'media_type': u'application/oracle.com.cloud.jaas.JavaPlatformTemplate+json', u'name': u'PS4_LowHeapTemplate'}
context_id 67A4827A1BB5159E2FDEFF4294F28F4F
data_sources {u'media_type': u'application/oracle.com.cloud.jaas.DataSource+json', u'total': u'1', u'elements': [{u'context_id': u'QA_app_DS@67A4827A1BB5159E2FDEFF4294F28F4F', u'uri': u'/em/cloud/jaas/datasource/QA_app_DS@67A4827A1BB5159E2FDEFF4294F28F4F', u'media_type': u'application/oracle.com.cloud.jaas.DataSource+json', u'name': u'QA_app_DS'}]}
datasource_params [{u'sensitive': u'false', u'require': u'true', u'type': u'String', u'name': u'username', u'description': u'TODO description i18n'}, {u'sensitive': u'true', u'require': u'true', u'type': u'String', u'name': u'password', u'description': u'TODO description i18n'}]
jdbc_drivers [u'com.ddtek.jdbc.oracle.OracleDriver', u'oracle.jdbc.xa.client.OracleXADataSource', u'oracle.jdbc.OracleDriver', u'com.ddtek.jdbc.oracle.OracleDataSource']
media_type application/oracle.com.cloud.jaas.JavaPlatformInstance+json
name jb_pf_14_25_1
resource_state (u'state': u'READY')
server_count 1
service_family_type jaas
status RUNNING
zone {u'service_family_type': u'jaas', u'context_id': u'A1B44A4EBCC4563125D9D0A3AAE4FD51', u'uri': u'/em/cloud/jaas/zone/A1B44A4EBCC4563125D9D0A3AAE4FD51', u'media_type': u'application/oracle.com.cloud.jaas.Zone+json', u'name': u'Zone1'}

```

## 31.8 Conclusion

This ends the introduction to cloud blueprints. Many features are only described in [Chapter 32, "Cloud Blueprints and Blueprint Reference"](#) including:

- Blueprint macros
- Path expressions for browsing the Cloud and blueprint
- Setting breakpoints and debugging blueprints
- Hints, tips, and frequently asked questions.





---

## Cloud Blueprints and Blueprint Reference

---

This chapter serves as a reference to the blueprint processing features of Oracle Enterprise Manager Cloud. It presents concepts, describes how to install and run the blueprint processor, and documents the blueprint language. Also included are sections to help in your use of the blueprint processor, such as use of the blueprint debugger, what to do when errors are diagnosed, and frequently asked questions. Before reading this document, you must have read [Chapter 31, "Introduction to Blueprints"](#).

It contains the following sections:

- [Installing the Blueprint Processor](#)
- [Optional Components for Graphical Summary Report](#)
- [Running the Blueprint Processor](#)
- [Blueprint Processing Phases](#)
- [Language Specifics](#)
- [Overview of Blueprint Content](#)
- [Dealing with Errors](#)
- [Simulation Mode](#)
- [Debugging with the Blueprint Processor](#)
- [Tips and Hints](#)

### 32.1 Installing the Blueprint Processor

Installing the blueprint processor consists of two steps:

- Install Python 2.7, if not already present.
- Install the blueprint processor files.

To use the “-g” option and generate graphical depictions of blueprints, additional software must be installed as described in [Section 32.2, "Optional Components for Graphical Summary Report"](#).

Detailed instructions to install the blueprint processor on Linux and Windows is given below.

---

**Note:** Due to schedule constraints, this version of the blueprint processor is not localized and is available only in English.

---

## 32.1.1 Linux (Oracle Linux)

Installing the blueprint processor on Oracle Linux involves the following steps:

- Installing Python
- Installing the Blueprint Processor From Zip File
- Testing the Installation

### 32.1.1.1 Installing Python

- Download a version of Python, version 2.7 or higher (but not 3.x). For instance, use <http://www.python.org/ftp/python/2.7.3/Python-2.7.3.tgz>.
- Untar it, for example by entering the following commands:

```
tar xzf Python-2.7.3.tgz
cd Python-2.7.3
```

- Execute these commands:

```
./configure --prefix=$HOME
make
make install
```

**Note:** During the configure step, you may see output that includes warnings about some modules that could not be created on your platform, but those modules may not be needed; their absence may not adversely affect use of the blueprint processor.

Python build finished, but the necessary bits to build these modules were not found:

```
bsddb185 dl imageop
sunaudiodev
```

To find the necessary bits, look in the `setup.py` in `detect_modules()` for the module's name.

Failed to build these modules:  
sqlite3

- The python interpreter should be installed in your home bin directory. Test it.  
`$HOME/bin/python2.7`
- You should see the Python banner... something like this:

```
Python 2.7.3 (default, May 25 2012, 11:33:27)
[GCC 4.1.2 20080704 (Red Hat 4.1.2-50)] on linux2
Type "help", "copyright", "credits" or "license" for more information.
```

### 32.1.1.2 Installing the Blueprint Processor from Zip File

- Unzip the blueprint distribution file to a directory, e.g. named `bp_installation`
- To run the blueprint processor, you must do the following:

cd to the directory, e.g. `bp_installation`

Enter this command: `python bp_processor.py <filename> <options>`

### 32.1.1.3 Testing the Installation

- Navigate to the directory that contains the blueprint processor files.
- Enter this command to confirm you can run the blueprint processor, in this case just to see the help text:

```
$HOME/bin/python2.7 bp_processor.py helloWorld.yml -h
```

- You should see the **help** text, as follows:

#### Usage:

```
bp_processor.py [options] BlueprintFileName
```

#### Example:

```
bp_processor.py myfile.yml -i "name:Joe" -i "count:2" -u jabauer -c
https://...
```

#### Options:

```
-h, --help show this help message and exit
-c CLOUD_URI, --cloud_uri=CLOUD_URI.
 Ex: https://myhost:4473/em/cloud
-u USER, --user=USER user id
-p PASSWORD, --password=PASSWORD
 password If not provided, you'll be prompted.
-t TIMEOUT, --timeout=TIMEOUT
 timeout (seconds)
-i INPUT_VALUE, --input_value=INPUT_VALUE
 <param name>:<param value> e.g. -i "name:Joe Blogs"
-r REFRESH_FREQUENCY, --refresh_frequency=REFRESH_FREQUENCY
 # of seconds between dots in timeline. Zero: no dots
-n INSTANCE_NAME, --instance_name=INSTANCE_NAME
 name of blueprint instance to create
-d, --debug more logging, including http traffic
-E, --pause_error Drop into debugger if error occurs
-I, --pause_input Pause before processing Input section
-R, --pause_resource Pause before processing Resource section
-O, --pause_output Pause before processing Output section
-T, --pause_terminate
 Pause before termination, but after output processing
 or error
-g GRAPHIC_RESULTS, --graphic_results=GRAPHIC_RESULTS
 Directory for deployed blueprint graphical report
-G GRAPHIC_BLUEPRINT, --graphic_blueprint=GRAPHIC_BLUEPRINT
 Directory for undeployed blueprint graphical report
```

- Enter this command to simulate running a blueprint without having to connect to a cloud environment:

```
$HOME/bin/python2.7 bp_processor.py helloWorld.yml
```

You should see something like this...

```
[jabauer@zzzzzzzz blueprints]$ $HOME/bin/python2.7 bp_processor.py
helloWorld.yml
```

```
Blueprint Processor - Invocation Summary
```

```

Cloud URI: sim
User: None
Blueprint file: helloWorld.yml
Timeout: 90 minutes, 0 seconds
```

```
Ellipses frequency: 15 seconds
Inputs:
Pause points: (none)
Debug logging: False
Instance name: default_instance_name

17:10:31 WARNING: No Resources specified in blueprint. Nothing will be
created

17:10:31 INFO: Output Processing
17:10:31 INFO: -----
17:10:31 INFO:
17:10:31 INFO: Output values specified: 1
17:10:31 INFO: Value of MyMsg: Hello World
17:10:31 INFO:
17:10:31 INFO: Blueprint Processing Summary
17:10:31 INFO: -----
17:10:31 INFO:
17:10:31 INFO: Timing Summary (seconds):
17:10:31 INFO: Client-side CPU time: 0.218
17:10:31 INFO: Elapsed time:
17:10:31 INFO: Processing time: 0.0
17:10:31 INFO: Paused time: 0.0
17:10:31 INFO: Total elapsed time: 0.0
```

## 32.1.2 Windows

Installing the blueprint processor on Windows involves the following steps:

- Installing Python
- Installing the blueprint processor from Zip File
- Testing the Installation

### 32.1.2.1 Installing Python

- Go to <http://www.python.org/download/releases/2.7.3>. You will see that there are several options for installing Python.
- Download the MSI Installer. Run it. Accept the defaults suggested by the installer.
- Note the directory into which Python was installed, typically  
**;%c:\python27**
- Add the Python directory to your path, so you can run it from the command line:
  - Windows XP
    - \* Start -> My Computer -> Properties
    - \* Select the **Advanced** tab
    - \* Click **Environment Variables**
    - \* You can update the path for 'Current User' and, if you have sufficient privileges, for 'System'. The latter is required if you want all users to be able to run Python.
    - \* Add this to the end of the Path string:  
**;%c:\python27**

- Windows 7 (Windows Vista)
  - \* From the **Start** menu, select **My Computer**, then select **Properties**.
  - \* Select the **Advanced System Settings** tab on the left.
  - \* Click **Environment Variables**.
  - \* You can update the path for 'Current User' and, if you have sufficient privileges, for 'System'. The latter is required if you want all users to be able to run Python.
  - \* Add this to the end of the Path string:
 

```
;%c:\python27
```
  - \* Click **OK** as needed (i.e. for two dialog boxes)
- Test the Python installation:
  - Open a new command window. (The new environment variable value will not be visible in previously opened command windows.)
  - Type **python**
  - The Python interpreter should start.
  - Type 'exit()'

### 32.1.2.2 Installing the Blueprint Processor from Zip File

- Unzip the blueprint distribution file to a directory, e.g. named bp\_installation.
- To run the blueprint processor you will...
  - cd to bp\_installation
  - Enter this command: `python bp_processor.py <filename> <options>`
  - Alternatively use: `bp_processor.py <filename> <options>`

### 32.1.2.3 Testing the Installation

- Navigate to the directory that contains the blueprint processor files.
- Enter this command to confirm you can run the blueprint processor, in this case just to see 'help' text:
 

```
python bp_processor.py helloWorld.yml -h
```
- Enter this command to simulate running a blueprint without having to connect to a cloud environment:
 

```
python bp_processor.py helloWorld.yml
```
- You should see something like this, in particular the output value "Hello World":
 

```
C:\...>python bp_processor.py interactiveTests\helloWorld.yml
```

Blueprint Processor - Invocation Summary

```

Cloud URI: sim
User: None
Blueprint file: helloWorld.yml
Timeout: 90 minutes, 0 seconds
Ellipses frequency: 15 seconds
Inputs:
Pause points: (none)
```

```
Debug logging: False
Instance name: default_instance_name

17:10:31 WARNING: No Resources specified in blueprint. Nothing will be created

17:10:31 INFO: Output Processing
17:10:31 INFO: -----
17:10:31 INFO:
17:10:31 INFO: Output values specified: 1
17:10:31 INFO: Value of MyMsg: Hello World
17:10:31 INFO:
17:10:31 INFO: Blueprint Processing Summary
17:10:31 INFO: -----
17:10:31 INFO:
17:10:31 INFO: Timing Summary (seconds):
17:10:31 INFO: Client-side CPU time: 0.218
17:10:31 INFO: Elapsed time:
17:10:31 INFO: Processing time: 0.0
17:10:31 INFO: Paused time: 0.0
17:10:31 INFO: Total elapsed time: 0.0
```

## 32.2 Optional Components for Graphical Summary Report

The blueprint processor can generate a summary report that includes a graphical depiction of the blueprint. (To generate such reports, use the `-g` or `-G` option.) For these options to produce reports, the following third party software must be installed.

- **Install GraphViz:** To install GraphViz, see <http://www.graphviz.org>. Download the software for your platform and follow the instructions.
- **Install Pydot:** To install pydot, see <http://code.google.com/p/pydot/>. Download the software (zip or tar file). The blueprint processor was tested using pydot version 1.0.28.

PyDot can be installed using `setuptools`. For example, `sudo easy_install pydot`

You can also use the `setup.py` script in the zip/tar file. From the directory into which you unzipped or untarred the file, run that script: `python setup.py install`

## 32.3 Running the Blueprint Processor

To deploy a blueprint, you run the blueprint processor and provide the name of the blueprint file plus any desired command options. For a full set of command options, use the command's `-h` option:

### Usage:

```
bp_processor.py [options] BlueprintFileName
```

### Example:

```
bp_processor.py myfile.yml -i "name:Joe" -i "count:2" -u jabauer -c
https://...
```

### Options:

```
-h, --help show this help message and exit
-c CLOUD_URI, --cloud_uri=CLOUD_URI.
 Ex: https://myhost:4473/em/cloud
-u USER, --user=USER user id
-p PASSWORD, --password=PASSWORD
```

```

 password If not provided, you'll be prompted.
-t TIMEOUT, --timeout=TIMEOUT
 timeout (seconds)
-i INPUT_VALUE, --input_value=INPUT_VALUE
 <param name>:<param value> e.g. -i "name:Joe Blogs"
-r REFRESH_FREQUENCY, --refresh_frequency=REFRESH_FREQUENCY
 # of seconds between dots in timeline. Zero: no dots
-n INSTANCE_NAME, --instance_name=INSTANCE_NAME
 name of blueprint instance to create
-d, --debug more logging, including http traffic
-E, --pause_error Drop into debugger if error occurs
-I, --pause_input Pause before processing Input section
-R, --pause_resource Pause before processing Resource section
-O, --pause_output Pause before processing Output section
-T, --pause_terminate Pause before termination, but after output processing
 or error
-g GRAPHIC_RESULTS, --graphic_results=GRAPHIC_RESULTS
 Directory for deployed blueprint graphical report
-G GRAPHIC_BLUEPRINT, --graphic_blueprint=GRAPHIC_BLUEPRINT
 Directory for undeployed blueprint graphical report

```

Options that are not self explanatory are described below:

- **Input\_value:** For each input parameter, you can provide a `-i` or `--input_value` string of the form `<param name>:<value>`. You should use quotation marks around each input parameter specification, e.g. `-i "name:Joe Blogs"`. To provide values for more than one input parameter, use the `-i` command line option more than once. For example, `bp_processor.py _processor.py myfile.yml -i "DbZone:Zone1" -i "DbPassword:myPw"7`
- If an input parameter defined in the blueprint is not provided on the command line, you will be prompted as follows:

`pause_*`: Pause (enter the blueprint debugger), just prior to beginning the Input/Resource/Output processing phases.

In the case of `pause_error`, the debugger is entered in the event of an error. In the case of `pause_terminate`, the debugger is entered prior to terminating the blueprint processor. Use of the debugger commands is described in [Section 32.9, "Debugging with the Blueprint Processor"](#). If no `CLOUD_URI` is specified, the blueprint processing is simulated. See the section titled [Section 32.8, "Simulation Mode"](#).

- **Timeout:** The number of seconds after which the blueprint process must be terminated. (If the processor is waiting for completion of a cloud request, e.g. a POST request, termination occurs when that request completes.)

## 32.4 Blueprint Processing Phases

Blueprint processing is done in phases. Describing the phases, will help you:

- Understand how blueprints are evaluated and resources created.
- Set pause points at phase transitions. (This can help with cloud browsing and debugging as explained in [Section 32.9.2, "Pause Points"](#)).
- Connect to the designated Cloud resource.
- Monitor the process. If creation succeeds, other resources may then become unblocked for creation. If creation fails, terminate the resource creation phase.

The phases are:

- **Initialization:** Parse the blueprint.
- **Input Processing:** If the blueprint defines input parameters and if some were not provided on the command line, the user is prompted for the parameter values.
- **Resource Creation:** For each resource defined in the blueprint, verify if the resource has no dependencies on other resources that have not been created yet, initiate creation.
- **Output Processing:** Evaluate and display any output values. (If the resource creation phase was terminated due to errors, some output values may not be available.)

A detailed description of each phase is given below.

### 32.4.1 Initialization

The blueprint file is read and parsed and a connection to the designated Cloud resource is made. The parsed blueprint content is captured in memory and then augmented. A member named Cloud is added, whose value is the Cloud resource as documented in [Chapter 25, "Cloud Resource Models"](#). Another member named Info, which provides environmental information, is also added. Its value is a set of name/value pairs such as time and a string like '13:02:45' representing the time when processing began. Oracle-provided macros are loaded into the Macros section of the blueprint (except for any whose name conflicts with a blueprint-defined macro).

### 32.4.2 Input Parameter Evaluation

Input parameter processing is then done and the input parameters set. Any parameters specified in the blueprint that were not provided on the command line are prompted for.

Each input parameter value is stored in the Value attribute of the input parameter in the blueprint. The values can be accessed with this path expression: Inputs.<parameterName>.Value. (See [Section 32.6.7.1, "Evaluation Intrinsics"](#) for details.)

### 32.4.3 Resource Creation

To process the Resources section, each resource is effectively processed in parallel. The following is done for each resource:

First, the expression specified for the Container section is evaluated. When successful and when the cloud resource it identifies is in the READY state, all expressions of the resource's Properties subsection are evaluated. (When evaluation of a resource's Properties or Container section cannot be completed and must wait, the resource is marked and the blueprint processor proceeds to process other resources. Periodically, it reattempts to evaluate this and any other marked resources.)

Once all evaluation for a resource definition is complete, the document derived from the Properties section is used to request creation of the cloud resource (i.e. it is POSTed to the container URI). If successful, the URI of the newly created cloud resource is stored in the '\_uri' attribute of the blueprint's resource definition. At this point, that URI represents the resource that is being created.

The blueprint processor then polls the resource being created to track its status. From a state of CREATING, the resource should eventually transition to a success or failure state. If it transitions to a failure state, the blueprint processor diagnoses the situation



and terminates. If creation succeeds and the resource enters the READY state, this may enable evaluation of other resource definitions to proceed.

If the timeout value specified on the command line is exceeded, the blueprint processor is terminated. Any cloud resources whose creation was initiated may continue in the 'creating' state for some time before succeeding or failing.

In addition to timing out, other failures may occur such as an error evaluating an expression or an error code returned from the cloud, e.g. in response to a creation request. In all such cases, the blueprint processor diagnoses the situation and terminates.

### 32.4.4 Outputs

After all resources have been successfully evaluated, the Outputs section is processed, at which time each named output expression is evaluated and their values printed. Then a graphical report that depicts the blueprint and what got created is generated, if it was requested and a summary of blueprint processing is displayed to the user.

## 32.5 Language Specifics

A blueprint is a text file that represents a set of cloud resources that are to be created. The text is formatted in YAML or JSON. Both YAML and JSON are notations for representing data structures of lists and name/value pairs, which can be nested. A blueprint is such a structure. (See [YAML] and [JSON] for concise descriptions.)

You can use JSON or YAML notation for your blueprints, but we recommend using YAML. With YAML you can write blueprints that are more concise and somewhat easier to read. YAML also offers useful capabilities not present in JSON, such as the ability to include comments. For these reasons, we have chosen to use YAML in the examples that follow. (Incidentally, YAML allows the inclusion of JSON notation in YAML documents, i.e. YAML is a superset of JSON.)

## 32.6 Overview of Blueprint Content

At the top level, a blueprint may contain any of the following name/value pairs, which can be viewed as blueprint section types:

- Inputs
- Data
- Macros
- Resources
- Outputs

Order of appearance in the blueprint is not significant. No other sections are allowed and each section can appear only once.

No section is truly required in a blueprint, but you must include a Resources section with at least one resource definition if you intend to create a cloud resource via the blueprint. A section may only appear once in a blueprint.

### 32.6.1 Inputs Section

The Inputs section is used to describe input parameters. Each parameter has a unique name and has the following attributes:

- Type: “String” or “Number”. If “Number” is specified, the input value must be numeric. (Default: “String”.)
- Prompt: A string to be used when prompting. (Default: The parameter name.)
- DefaultValue: The value to use if the user responds to the prompt by simply pressing Enter. (Default: “”.)
- Order: The sequence in which the values need to be entered. If not specified, ordering is arbitrary.
- Sensitive: “True” or “False”. If “True”, what the user enters interactively will not be echoed. (Default: “False”.)
- Value: Set at runtime, using the value provided by the user (or the default value)

Processing input parameters starts with those input parameter values provided on the command line. For any not provided, the user is prompted.

### 32.6.1.1 Example 1

```
Inputs:
 UserId:
 DefaultValue: qa_user
 Prompt: User id
 Order: 1
 Password:
 Sensitive: True
 Order: 2
...
```

In this blueprint snippet, the author has defined 2 input parameters. Both are of type String. A default user id is specified but the user must provide a password. If neither is provided on the command line, the use of Order assures that UserId is requested first.

### 32.6.1.2 Example 2

In this interaction

```
C:\work>bp_processor.py test2.yml -c https://... -u jon -p myPW -t 665
```

```
Blueprint Processor - Invocation Summary
```

```

Cloud URI: https://slc01rbw.us.oracle.com:15430/em/cloud
User: ssa_user1
Blueprint file: xyzApp.yml
Timeout: 90 minutes, 0 seconds
Refresh frequency: 15 seconds
Inputs:
Pause points: Inputs
Debug logging: False
Instance name: default_instance_name
Graphical report dir: deployment_report
Versions:
 Blueprint processor: 12.1.0.5, 10-Oct-2012
 Cloud protocol: 10001
```

```
16:19:18 INFO: Connecting to cloud: https://...
User id (qa_user):
Password:
...
```

The user provided no input parameters as part of the command line, so he is prompted for the two values. He pressed **Enter** for the user id, accepting the default of **qa\_user**. And he entered a password, which was not echoed.

### 32.6.1.3 Example 3

The following interaction is almost the same, except that the user id is provided on the command line:

```
C:\work>bp_processor.py test2.yml -c https://... -u jon -p myPW -t 665 -i
"UserId:joe"
```

Blueprint Processor - Invocation Summary

```

Cloud URI: https://slc01rbw.us.oracle.com:15430/em/cloud
User: ssa_user1
Blueprint file: xyzApp.yml
Timeout: 90 minutes, 0 seconds
Refresh frequency: 15 seconds
Inputs:
Pause points: Inputs
Debug logging: False
Instance name: default_instance_name
Graphical report dir: deployment_report
Versions:
 Blueprint processor: 12.1.0.5, 10-Oct-2012
 Cloud protocol: 10001

16:19:18 INFO: Connecting to cloud: https://...
16:19:19 WARNING: No Resources specified in blueprint. Nothing will be created
16:19:19 INFO: Creating blueprint instance named default_instance_name
Password:
...
```

In this case, the user is only prompted for a password.

Another example can be found at [Section 32.6.7.1.4, "Example 3 - Default Input Parameter Value via Cloud Lookup"](#) which uses an intrinsic function to perform a runtime lookup for the default value.

## 32.6.2 Resources Section

The Resources section is used to describe the cloud resources you wish to create. Each resource description has a unique name and has the following attributes:

- **Container:** The URI of the parent cloud resource.
- **Type:** The media type of the resource, as defined by the Cloud Resource Model. (see [Chapter 25, "Cloud Resource Models"](#)). (If the specified Container is a resource that is a subtype of ServiceTemplate, this value is optional and the default type of the ServiceTemplate is assumed.)
- **Properties:** A set of name/value pairs used to specify values required per the Cloud Resource Model.

For each resource you wish to create, you must know its type, its parent, and the properties you must provide to specify its characteristics. For more details, see the [Chapter 25, "Cloud Resource Models"](#).

For instance, a JavaPlatformInstance must specify:

- A container resource, using the URI of a `JavaPlatformTemplate`. (You choose a template by selecting one that creates an instance most suited to your needs.)
- The properties required by the Cloud Resource Model. For `JavaPlatformInstance`, these are:
  - Name
  - Zone (the URI of a zone)

### Example

In a blueprint, the Resource description of a `JavaPlatformInstance` might look like this:

```
MyJavaServer1:
 Container:
 f_getTemplateURI:
 - Small WLS
 - jaas
 Properties:
 name: Foo
 zone:
 f_getZoneURI:
 - Zone1
 - jaas
```

The above example does not specify **Type** because, this is optional when the container is a subtype of `ServiceTemplate`. Had we wanted to explicitly specify the media type, we could have written:

```
MyJavaServer1:
 Type: application/oracle.com.cloud.jaas.JavaPlatformInstance
 Container:
 ...
```

Our example uses the `f_getTemplateURI` and `f_getZoneURI` intrinsic functions to look up the required URI's. Resource definitions have access to numerous such functions as described in [Section 32.6.7, "Intrinsic Functions"](#). Resource definitions can also use user-defined macros, as described in [Section 32.6.5, "Macros Section"](#).

## 32.6.3 Outputs Section

The Outputs section is used to describe the set of “outputs” of a blueprint. In this release, outputs are just used by the blueprint author to specify information to display at the end of a successful deployment. For instance, he may define an output that displays the URL of an application deployed to a `JavaPlatformInstance`.

Each output description has a unique name and one required attribute named “Value”. That attribute generally specifies a blueprint expression whose value would be of interest to the user who instantiates a blueprint. (One may also include a “Description” attribute for each output definition.)

### Example

The following Outputs section specifies one output to display:

```
Outputs:
 Application_URL:
 Description: URL of the deployed app
 Value:
 f_getResourceAttr:
```

- MyApp
- http\_application\_invocation\_url

In this example, the output definition uses the `f_getResourceAttr` intrinsic function to retrieve the `http_application_invocation_url` attribute of a newly created `ApplicationInstanceDeployment`. (See [Chapter 25, "Cloud Resource Models"](#) for more information on that attribute.)

### 32.6.4 Data Section

The Data section contains arbitrary YAML text. The data defined in this section can be accessed through intrinsic functions and can be used in various ways during blueprint deployment. For instance, the author of a blueprint may wish to use the concept of named literals to improve the readability and maintainability of a blueprint. He may want to use short descriptive names instead of long literal values such as cryptic URI's. He may also anticipate needing to change a value used by a blueprint, in which case it can be specified once in the Data section and referenced throughout the blueprint.

#### Example

The following blueprint excerpts show a Data section that contains an item named `db_conn_str`, a long JDBC connect string that is referenced later in the a `DataSource` resource definition.

```
Data:
 db_conn_str: 'jdbc:oracle:thin:sysman/sysman@hostname.zzz.com:15044:smay16'
Resources:
 ...
 MyDataSource:
 Type: application/oracle.com.cloud.jaas.DataSource
 ...
 Properties:
 name: jbTest
 jdbc_driver: oracle.jdbc.OracleDriver
 database_type: Oracle
 database_connect_string:
 f_path:
 - Data.db_conn_str
 ...
```

### 32.6.5 Macros Section

The notation used in blueprints can be verbose. If you have a sequence of constructs that you tend to repeat, you can use macro expansion to improve the readability of your blueprint. Macros also enable you to encapsulate logic so that changes made one in section get reflected throughout the document.

The Macros section is used to define macros that can be invoked/expanded from elsewhere in the blueprint. A macro invocation may occur wherever a function invocation is allowed, indeed, the notation is identical and one cannot tell from the invocation whether it is an intrinsic function or a macro that is being invoked. (In this way, a macro can be used to override an intrinsic function. Also, some Oracle-provided intrinsic functions are implemented as macros.)

Each macro definition has a unique name and its definition specifies two values:

- The number of arguments it uses.
- The textual representation of the macro expansion.

When a macro is invoked, its textual representation replaces the invocation. Wherever the textual representation specifies a value of `arg_<integer>`, the value of that argument is used instead.

### Example

Consider this somewhat contrived blueprint:

```
Macros:
 # Return a string that describes a resource being created
 # The one argument is a 'name' string
 f_myDescriptiveName:
 - 1
 - f_concat:
 - "Resource "
 - arg_1
 - " created for blueprint instance "
 - f_path:
 - "Info.instance_name"
 - " on "
 - f_path:
 - "Info.date"
Resources:
 MyJavaServer:
 Container:
 f_getTemplateURI:
 - Small WLS
 - jaas
 Properties:
 name:
 f_myDescriptiveName:
 - MyFirstJavaServer
 zone:
 f_getZoneURI:
 - Zone1
 - jaas
Outputs:
 NameOfServer:
 Value:
 f_path:
 - "Resources.MyJavaServer.Properties.name"
```

The Macros section defines one macro named `f_myDescriptiveName`, which takes one string argument and constructs a larger string adding descriptive information.

The macro is then invoked as part of the `MyJavaServer` resource definition. It is invoked with a value of `"MyFirstJavaServer"` and the resource will be created with a `'name'` property whose value is `"Resource MyFirstJavaServer created for blueprint instance myQAInstance on 5/9/2012"`.

Blueprint Processor - Invocation Summary

```

Cloud URI: https://...
User: sysman
...
13:32:22 INFO:
13:32:22 INFO: Output values specified: 1
13:32:22 INFO: Value of NameOfServer: Resource MyFirstJavaServer created for
blueprint instance myQAInstance on 5/9/2012
...
```

## 32.6.6 Expressions

A blueprint uses ‘expressions’ to compute values at deployment time. In many cases, the expressions are simple literal string values such as a `user_name` attribute whose value is `app_user`. In other cases, values are constructed via user-defined macros (see [Section 32.6.5, "Macros Section"](#)) and intrinsic functions (see [Section 32.6.7, "Intrinsic Functions"](#)).

Two intrinsic functions, `f_path` and `f_eval`, are provided to evaluate the two types of expression string:

- Path expression
- Eval (or Blueprint) expression

These expression types are described below. They can be used not only in calls to `f_path` and `f_eval` but also in the blueprint debugger.

### 32.6.6.1 Path Expressions

Path expressions are similar to JSONPath [JSONpath] and XPath expressions and are used to extract values from blueprints and cloud resources as well as to traverse URI's that link you to other resources.

The starting point for evaluating a path expression is generally the in-memory blueprint. Recall that it contains all the information from your blueprint plus these attributes:

- Cloud: the cloud resource, which is defined in [Chapter 25, "Cloud Resource Models"](#) and has attributes that describe the overall cloud as well as those that enable you to traverse to all other cloud resources (to which you have access).
- Info: a section that contains runtime information you may wish to refer to such as the instance name specified when blueprint processing was initiated, the current date, etc.
- Other blueprint-processor-computed values such as the input parameters provided by the end-user and the URI's of resources once they have been created.

All this information can be accessed via path expressions. The rest of this section summarizes syntax and semantics and provides examples. As you have seen in earlier sections, the `f_path` intrinsic function provides one way to evaluate path expressions. In the examples that follow, we use another mechanism, the blueprint debugger, described further in [Section 32.9, "Debugging with the Blueprint Processor"](#). As you will see, you can use this not only to explore the contents of your blueprint runtime data but also to explore the contents of the cloud.

In the following descriptions, `<doc>` refers to the YAML data structure on which the operator acts.

In the example: `Cloud.zones.elements[0]`

The value of the runtime blueprint's `Cloud` attribute is the first `<doc>`. The dot-operator in `zones` is applied to that `<doc>`, yielding a new `<doc>` whose value is that of the `zones` attribute of the first `<doc>`. Numerous examples are provided to illustrate the use of the operators.

### 32.6.6.2 Operator Summary

- **Dot Operator ("Member of"):** `<doc>.<name>`: As seen in previous examples, the dot operator selects a value from a document. If `<doc>` has an attribute named `<name>`, the value of `<name>` is returned. Otherwise, expression evaluation fails.

- **Square Bracket (“List Indexing”):** `<doc>[<integer>]`: If `<doc>` is a list of values, the value of the `<integer>`th element is returned. Otherwise, expression evaluation fails. Indexing is zero-based, i.e. the first element is specified as `“<doc>[0]”`.
- **Dollar Sign (“Literal string prefix”):** `<doc> $ <string>`: To begin a path expression with a string literal, use the `‘$’` prefix, for instance:

```
$"/em/cloud/jaas/zone/A1B44A4EBCC4563125D9D0A3AAE4FD51" ->
```

In the above example, you know the URI of a specific resource and use the arrow operator to view its contents.

This syntax is only useful at the beginning of a path expression, but for consistency with the overall path notation, one can place the `$` operator anywhere in a path expression. Other path expression operators operate on the `<doc>` to the left. The `$` operator simply replaces the left value with the value of the literal string to the right.

In short, the `$` operator returns the literal string value. If the right operand is not a literal string, expression evaluation fails.

- **Arrow (“URI Traversal”):** `<doc> ->`: If `<doc>` is a URI, traverse to the identified resource and return its document. In other words, perform a GET on the URI specified by `<doc>`. If `<doc>` is not a URI or the GET fails, expression evaluation fails.

For instance, in this path expression...

```
Cloud.zones.elements[0].uri->
```

the expression to the left of the arrow operator returns the URI of a zone. The arrow operator is used to traverse to the zone, i.e. it performs a GET on the URI and displays the contents:

```
context_id: A1B44A4EBCC4563125D9D0A3AAE4FD51
description: Zone for Physical Pool
media_type: application/oracle.com.cloud.jaas.Zone+json
name: Zone1
resource_state:
 state: READY
service_family_type: jaas
service_instances:
 elements: []
 media_type: application/oracle.com.cloud.common.ServiceInstance+json
 total: '0'
uri: /em/cloud/jaas/zone/A1B44A4EBCC4563125D9D0A3AAE4FD51
```

In most cases the traversal operator is sufficient, but you can also specify traversal qualifiers. In particular, you can specify a media type and request parameters. These are optional and enclosed in square brackets. Multiple qualifiers can be specified, separated by commas.

To specify a media type, use a string (enclosed in quotes or double quotes).

- To specify a request parameter there are two styles:
  - Identifier
  - Identifier = quotedValue

Overall, there are 3 forms of traversal qualifiers

- *quotedValue*:



media type

- *Identifier:*  
a request parameter that has no value
- *Identifier = quotedValue*  
a request parameter that has a value

For instance, the following path expression traverses to the cloud URI and specifies three qualifiers (a media type and two request parameters).

```
Cloud.uri->
["application/oracle.com.cloud.common.ServiceTemplateFinds+json",
filters='{"filters": {"service_family_type": "jaas"}', name]
```

The first parameter is the media type to be used to request that the cloud process a filter/query as described in [Chapter 25, "Cloud Resource Models"](#). The remaining two are request parameters, also described in [Chapter 25, "Cloud Resource Models"](#). The second is named 'filters' with a value that specifies the filtering to perform (expressed as a JSON string). The last is named 'name' and has no value. This specifies that the name attribute is to be returned.

### 32.6.6.3 Example: Viewing all values of Info

As noted above, the Info section contains environmental information that may be of use when constructing your blueprint. This example shows how to see those values that are currently available in the Info section.

Suppose you are looking for a value you can use to construct a unique name. You run the blueprint processor and enter debug mode. (You can do this by specifying "-I" on the command line.

```
C:\Users\jabauer\Dropbox\Code\blueprints>bp_processor.py helloWorld.yml
-c https://hostname.us.oracle.com:15430/em/cloud -u sysman -p sysman -I
```

...

Blueprint Processor - Invocation Summary

```

Cloud URI: https://hostname.us.oracle.com:15430/em/cloud
User: sysman
Blueprint file: helloWorld.yml
...
```

```
...Pause point, prior to Input processing...
For command info, enter (h)elp
```

```
Paused: Info
date: 1/11/2013
date_suffix: '1_11_2013'
instance_name: default_instance_name
time: '16:47:7'
time_suffix: '16_47_7'
uuid: 81dcacf6895fa4fb881e82d1c16ef7025
```

Here, you see that there are 6 values stored in the Info section. There is one named uuid that is a universally unique hexadecimal string. You may prefer to go with time\_suffix, as being sufficiently unique and more readable.

More values may be added between the time of writing this document and when the blueprint processor ships, so you can use this technique to see what's available in the version you are using.

#### 32.6.6.4 Examples: Viewing Blueprint Values

Suppose you have a blueprint that begins with...

```
Inputs:
 DbPassword:
 Type: String
 DefaultValue: welcome1
 Prompt: Password to use for db
 Sensitive: True
```

Say you run the blueprint processor and specify command line options to pause just before processing the Inputs and Resources section, e.g. you specify "-RI".

```
...Pause point, prior to Input processing...
For command info, enter (h)elp
```

```
Paused: Inputs.DbPassword
DefaultValue: welcome1
Prompt: Password to use for db
Sensitive: true
Type: String
```

At the first pause point above, you enter the expression 'Inputs.DbPassword' and see that it has the attributes you specified in your blueprint. That includes DefaultValue, Prompt, and Sensitive. Note that it does not have an attribute named Value because input processing has not been performed yet.

```
Paused: c
...continuing...
```

```
Input Parameter Value Entry

 Password to use for db (welcome1):
```

```
...Pause point, prior to processing Resources section...
For command info, enter (h)elp
```

```
Paused: Inputs.DbPassword.Value
welcome1
```

```
Paused: Inputs.DbPassword
DefaultValue: welcome1
Prompt: Password to use for db
Sensitive: true
Type: String
Value: welcome1
```

You then enter **c** to continue and are prompted for a password. You enter one, which is not echoed because you specified that it was **Sensitive**.

Then the second (prior to processing the Resources section) pause point is reached. You enter "Inputs.DbPassword.Value" to see the value of your password and then enter "Inputs.DbPassword" to see the value of all attributes for the DbPassword input parameter.

### 32.6.6.5 Examples: Browsing Your Cloud

Path expressions also offer an easy way to explore the resources in the cloud and their attributes. That is because, at the beginning of blueprint processing, the cloud resource (as defined in [Chapter 25, "Cloud Resource Models"](#)) is read and placed into the in-memory blueprint structure. By starting your path expression with **Cloud**, you can browse attributes of the Cloud resource and navigate via URI's to any other resource to which you have access.

To start, we look at the **description** of the cloud to which you connected:

Paused: **Cloud.description**

This represents the Cloud resource of the Oracle Enterprise Manager Cloud Management solution  
Paused

Now let us do something more useful, for example, look at the cloud's zones attribute:

Paused: **Cloud.zones**

elements:

```
- media_type: application/oracle.com.cloud.jaas.Zone+json
 name: Zone1
 service_family_type: jaas
 uri: /em/cloud/jaas/zone/A1B44A4EBCC4563125D9D0A3AAE4FD51
- description: Zone for Physical Pool
 media_type: application/oracle.com.cloud.common.DbZone+json
 name: Zone1
 type: self_service_zone
 uri: /em/cloud/dbaas/zone/A1B44A4EBCC4563125D9D0A3AAE4FD51
- media_type: application/oracle.com.cloud.opc.OpcZone+json
 name: OPC Zone
 service_family_type: opc
 type: opc
 uri: /em/cloud/opc/opczone
media_type: application/oracle.com.cloud.common.Zone+json
total: '3'
```

Here, we see that the 'zones' attribute contains three attributes, elements, media\_type, and total. Their meanings are described in [Chapter 25, "Cloud Resource Models"](#).

We wish to focus on the first zone listed, so we use the square bracket (list indexing) syntax:

Paused: **Cloud.zones.elements[0]**

media\_type: application/oracle.com.cloud.jaas.Zone+json  
name: Zone1  
service\_family\_type: jaas  
uri: /em/cloud/jaas/zone/A1B44A4EBCC4563125D9D0A3AAE4FD51

We can further specify that we wish to focus on the 'uri' attribute by adding another dot-operator:

Paused: **Cloud.zones.elements[0].uri**

/em/cloud/jaas/zone/A1B44A4EBCC4563125D9D0A3AAE4FD51

To view the resource to which that URI refers, we add the arrow (traversal) operator:

Paused: **Cloud.zones.elements[0].uri->**

context\_id: A1B44A4EBCC4563125D9D0A3AAE4FD51  
description: Zone for Physical Pool  
media\_type: application/oracle.com.cloud.jaas.Zone+json  
name: Zone1  
resource\_state:

```
state: READY
service_family_type: jaas
service_instances:
 elements: []
 media_type: application/oracle.com.cloud.common.ServiceInstance+json
 total: '0'
uri: /em/cloud/jaas/zone/A1B44A4EBCC4563125D9D0A3AAE4FD51
```

As you would expect, you can continue to add to your path expressions, for instance you can write "multi-hop" expressions that traverse multiple URIs, e.g. ...

```
Paused: Cloud.service_templates.elements[0].uri->zones.elements[0].name
Zone1
```

```
Paused: Cloud.service_templates.elements[0].uri->zones.elements[0].uri->
context_id: A1B44A4EBCC4563125D9D0A3AAE4FD51
description: Zone for Physical Pool
media_type: application/oracle.com.cloud.jaas.Zone+json
name: Zone1
resource_state:
 state: READY
service_family_type: jaas
service_instances:
 elements: []
 media_type: application/oracle.com.cloud.common.ServiceInstance+json
 total: '0'
uri: /em/cloud/jaas/zone/A1B44A4EBCC4563125D9D0A3AAE4FD51
```

Paused:

In the above example, the `service_templates` attribute of the Cloud is used to identify the first service template. We then traverse its URI to get to the template, where we identify the first zone in its list of supported zones. We then traverse its URI to get to the full definition of the zone.

In most cases the traversal operator is sufficient, but if there is no default media-type defined for the URI, you may need to specify the media-type to be retrieved, as specified in [Chapter 25, "Cloud Resource Models"](#). In the following example, we provide a media type even though it wasn't needed.

```
Paused: Cloud.service_templates.elements[0].uri->
["application/oracle.com.cloud.jaas.JavaPlatformTemplate"]

context_id: D2520A0CFFE348BCE040F20A4C1B2D8F
created: '2013-01-02 09:35:52.0'
default_instance_media_type:
application/oracle.com.cloud.jaas.JavaPlatformInstance+json
...
```

Similarly, you can use the traversal qualifier syntax to specify request parameters as defined in [Chapter 25, "Cloud Resource Models"](#). For example:

```
Paused: Cloud.service_templates.elements[0].uri->[created, resource_state]
created: '2013-01-02 09:35:52.0'
resource_state:
 state: READY
```

### 32.6.6.6 Eval or Blueprint Expressions

An 'eval' (aka 'blueprint') expression is any expression you can include in your blueprint.

**Example: Simple Intrinsic Function Evaluation**

Eval expressions can be evaluated in the debugger via the 'e' or 'eval' command. After entering the command, you enter the lines that comprise the expression followed by an empty line. Suppose you want to experiment with the `f_concat` intrinsic:

```
Paused: e
Eval: f_concat:
Eval: - xxx
Eval: - yyy
Eval:
xxxxyy
Paused:
```

You enter the call to `f_concat` (in 3 lines) and the value is printed after you terminate the expression with an empty line.

This time you nest another call:

```
Paused: e
Eval: f_concat:
Eval: - xxx
Eval: - f_path:
Eval: - 'Inputs.DbPassword.Value'
Eval: - yyy
Eval:
xxxmySecretyyy
Paused:
```

As you can see, this provides a way to experiment with snippets of blueprint.

**Example: Lookup Intrinsic Function**

In this example, suppose your blueprint includes the use of `f_getTemplateURI` to look up a template URI:

```
Resources:
 MyDB:
 Container:
 f_getTemplateURI:
 - Small DB
 - dbaas
```

If you think the wrong URI is being returned, you can check like this...

```
Paused: e
Eval: f_getTemplateURI:
Eval: - template1
Eval: - jaas
Eval:
/em/cloud/jaas/javaplatformtemplate/BFAB458D36BDA87EE040E50A038F6D45
Paused:
```

This shows the URI value returned by `f_getTemplateURI` with the arguments you entered.

**32.6.7 Intrinsic Functions**

Intrinsic functions are functions that blueprints can use to compute/return desired information. For instance, the `f_concat` function returns the concatenation of its string arguments and the `f_getZoneURI` function looks up and returns the URI of a zone. The normal usage of intrinsic functions is to provide values needed as part of the

Resources and Outputs sections, but an intrinsic function can be placed wherever a literal value is allowed. For instance, one can use an intrinsic to derive the DefaultValue used for an Input parameter.

This section describes the currently available intrinsic functions.

### 32.6.7.1 Evaluation Intrinsic

These two intrinsics are used to evaluate expressions (of different types) and return a single value.

#### 32.6.7.1.1 `f_path(pathExpr)`

Apply the pathExpr string to the blueprint document, returning the specified value.

##### Parameters

- **pathExpr:** e.g. "member.subMember..."  
A path expression that describes how to traverse and extract information from the document. See [Section 32.6.6.1, "Path Expressions"](#).

#### 32.6.7.1.2 Example - Value from Data Section

To specify a resource's property value using a literal value defined in the Data section of your blueprint ...

```
...
 params:
 MasterUser:
 f_deref:
 - "Data.QADBCreds.user"
 ...
```

#### 32.6.7.1.3 Example 2 - Value from Inputs Section

To do the same as above, only using an input parameter value ...

```
...
 params:
 MasterUser:
 f_deref:
 - "Inputs.my_param.value"
 ...
```

#### 32.6.7.1.4 Example 3 - Default Input Parameter Value via Cloud Lookup

Intrinsic functions can appear in sections other than the Resources section. For instance, this blueprint shows the use of `f_path` in the Inputs section.

```
Inputs:
 JavaSvcZone:
 DefaultValue:
 f_path:
 - 'Cloud.zones.elements[0].name'
 Type: String
 Prompt: Enter the name of a jaas zone
 ...
```

The blueprint author here wants to provide an arbitrary default zone name, so the path expression selects the first zone that appears in the 'zones' attribute of the cloud.

The default value (shown in the prompt) is computed at runtime:

```
Input Parameter Value Entry
```

```

```

```
Enter the name of a jaas zone (east_coast_zone):
```

### 32.6.7.1.5 `f_eval`(blueprintExpr)

Evaluate `blueprintExpr`, returning the specified value. See [Section 32.6.6.6, "Eval or Blueprint Expressions"](#) for information on blueprint expressions.

#### Parameters

**blueprintExpr:** YAML text to be evaluated as if it appeared in a blueprint.

### 32.6.7.1.6 Example (Contrived)

The `f_eval` intrinsic is used internally by the blueprint processor, and it's unlikely you will need to use it. (You just use the blueprint expression directly.) One reason you might want to use this function is if you have a variable whose value is a blueprint expression in the form of a YAML string.

```
Data:
 demoOfYamlMultilineText: |
 This is a multi-
 line text string which is
 carefully indented. :-)
 myBlueprintExpressionText: |
 f_concat:
 - 'Mister '
 - 'Mxyzptlk'
Outputs:
 demoOfYamlMultilineText:
 Value:
 f_path:
 - 'Data.demoOfYamlMultilineText'
 myBlueprintExpressionText:
 Value:
 f_path:
 - 'Data.myBlueprintExpressionText'
 useOfEvalonExpressionText:
 Value:
 f_eval:
 - f_path:
 - 'Data.myBlueprintExpressionText'
```

Which results in this output:

```
18:49:58 INFO: Output Processing
18:49:58 INFO: -----
18:49:58 INFO:
18:49:58 INFO: Output values specified: 3
18:49:58 INFO: Value of demoOfYamlMultilineText: This is a multi-
line text string which is
carefully indented. :-)
18:49:58 INFO: Value of myBlueprintExpressionText: f_concat:
- 'Mister '
- 'Mxyzptlk'
18:49:58 INFO: Value of useOfEvalonExpressionText: Mister Mxyzptlk
```

### 32.6.7.2 Resource Access Intrinsics

These intrinsics are used to access resource attributes. As part of their operation, unlike `f_path`, they assure that the resource is in the READY state, waiting if needed.

### 32.6.7.2.1 `f_getResourceAttr(bpResName, pathExpr)`

Get the value of a cloud resource attribute after it is READY.

#### Parameters

- **bpResName:** Resource name (specified in blueprint)
- **pathExpr:** Same semantics as used in `f_deref`, only against the document of the cloud resource identified by `bpResName`.
- **Returns:** result of evaluating `pathExpr` of the resource, once it is created and its `resource_state` is READY.

#### Example 1

To add an application to a MW platform, an Application resource can use this Container clause...

```
Container:
 f_getResourceAttr:
 - myJavaPlatform
 - uri
```

#### Example 2

To access the name of the zone in which your MW platform was created, you can write...

```
f_getResourceAttr:
 - myPlatform
 - zone.name
```

### 32.6.7.2.2 `f_getResourceURI(bpResName)`

Get the URI for a blueprint-defined resource. This is just a shorthand for using `f_getResourceAttr` with the specified attribute being 'uri'.

#### Parameters

- **bpResName:** Name used in the blueprint resource definition.
- **Returns:** URI

#### Example

To define a Datasource resource that is to be contained in a JavaPlatformInstance resource created elsewhere in your blueprint...

```
MyDatasource:
 Type: Datasource
 Container:
 f_getResourceURI:
 - MyJavaServer
```

### 32.6.7.3 Lookup Intrinsic

These intrinsic functions search for a template, zone, or application component, returning its URI.

#### 32.6.7.3.1 `f_getTemplateURI(name, type)`

Get the URI for a template, based on its name and type.

#### Parameters

- **name:** Template name



- **type:** A service type name. The current list of allowed values is iaas, jaas, and dbaas
- **Returns:** URI

### Example

A blueprint resource to create a database using template simpleDb, could be written...

```
Container:
 f_getTemplateURI:
 - simpleDb
 - dbaas
```

#### 32.6.7.3.2 f\_getZoneURI(name, type)

Get the URI for a zone, based on its name and type.

Parameters

- **name:** Zone name
- **type:** A service type name. The current list of allowed values is iaas, jaas, and dbaas
- **Returns:** URI

### Example

To get the URI of zone EMEA\_db\_zone...

```
f_getZoneURI:
 - EMEA_db_zone
 - dbaas
```

#### 32.6.7.3.3 f\_getAppCompURI(name, owner, version) ...

Get the URI for an application component, based on its name, owner, and version.

Parameters

- **name:** Application component name
- **owner:** Owner of application component
- **version:** Version of application component. If blank, the latest version is used.
- **Returns:** URI

### Example

To get the URI of the most recent version of the application component jbscomponent, owned by SSA\_USER1 ...

```
f_getAppCompURI:
 - jbscomponent
 - SSA_USER1
 -
```

Note that the third argument is required.

### 32.6.7.4 Debugging Intrinsic

These intrinsic are used to establish breakpoints or printpoints.

#### 32.6.7.4.1 f\_break(expression, [breakpointMessage])

Pause evaluation of the blueprint, print optional message, and enter blueprint debugger.

**Parameters**

- **expression:** any blueprint expression
- **breakpointMessage:** Message to be printed when the intrinsic function is invoked, just prior to entering the debugger.
- **Returns:** Value of expression. This is computed when the 'continue' command is entered.

**Note:** See [Section 32.9, "Debugging with the Blueprint Processor"](#) for examples and other information on how to use breakpoints to help debug your blueprints.

**32.6.7.4.2 f\_print(expression, [printpointMessage])**

Print a line that displays the value of expression. Pause evaluation of the blueprint, print optional message, and enter blueprint debugger.

**Parameters**

- **expression:** any blueprint expression
- **printpointMessage:** Message to be printed when the intrinsic function is invoked, just prior to entering the debugger.
- **Returns:** Value of expression.

**Note:** The line printed looks like this:

```
>>> Printpoint [<printpoint message>]:
Value = <expression>
```

See [Section 32.9, "Debugging with the Blueprint Processor"](#) for examples and other information on how to use breakpoints to help debug your blueprints.

**32.6.7.5 Other Intrinsic****f\_concat(string1, ... stringN):**

Returns the concatenation of the string arguments.

**Parameters**

**string\*:** A string to be concatenated with the other string arguments

**Example**

To set the description of a JavaPlatformInstance to "Created by blueprint FOO on <current date>" ...

```
Resources:
 MyJavaServer:
 ...
 Properties:
 description:
 f_concat:
 - "Created by blueprint FOO on "
 - f_path:
 - "Info.date"
 ...
```

## 32.7 Dealing with Errors

This section illustrates various types of errors you may encounter and, by example, how to interpret/resolve the issues.

### 32.7.1 YAML Syntax Errors

Any syntax errors encountered by the YAML parser are diagnosed by the parser. Consider this blueprint snippet...

```
Example of YAML syntax error
Data:
 userId: Lex
 password: changeMe
...
```

In the above example, the YAML parser would detect an indentation error and diagnose it like this:

```
18:55:59 ERROR: Error loading blueprint YAML:
while parsing a block mapping
 in "<string>", line 2, column 1:
 Data:
 ^
expected <block end>, but found '<block mapping start>'
 in "<string>", line 4, column 3:
 password: changeMe
 ^
```

The second half of the diagnostic, the “expected” part, is usually the most helpful. In this case, it tells you the error was detected at the token ‘password’, what it was expecting, and what it found.

One common error to avoid is the use of tabs in the YAML file. YAML does not allow tab characters.

### 32.7.2 Protocol Version Mismatch

A diagnostic like this...

```
Blueprint Processor - Invocation Summary

Cloud URI: http://hostname.us.oracle.com:4473/em
User: sysman
Blueprint file: examples/evalintrinsic.yml
Timeout: 90 minutes, 0 seconds
Refresh frequency: 15 seconds
Inputs:
Pause points: Inputs, Termination
Debug logging: False
Instance name: default_instance_name
Versions:
 Blueprint processor: 12.1.0.4 May 25
 Cloud protocol: 10001

19:05:06 INFO: Connecting to cloud: http://adc2100705.us.oracle.com:4473/em
19:05:07 ERROR: Cloud protocol version mismatch. Expected 10001. Found None.
```

The diagnostic in bold, indicates that the blueprint processor was able to connect to the site but did not get the expected response. In particular, the site returned no x-specification-version value as part of the HTTP response.

This can happen if you specified an incorrect cloud URI. For instance, in the above example, the URI was not of the form

<https://host:port/em/cloud>

A common mistake is to omit the `"/cloud"`. Also, do not forget to use *https*, not *http*.

If the diagnostic indicates it found a protocol version that is lower than the one expected, that may indicate that you are using a version of the blueprint processor that requires a more recent version of Enterprise Manager.

### 32.7.3 Expression Evaluation Error

As part of blueprint processing, an attempt to evaluate an expression may result in an error. Examples of errors include:

- Passing the wrong number of parameters to an intrinsic function
- Referring to a non-existent intrinsic function
- Referring to a non-existent cloud resource such as a zone or template

When an expression evaluation error occurs, the issue is diagnosed and the expression is displayed. For instance, consider this contrived blueprint:

```
Outputs:
 ExampleValue:
 Value:
 f_concat:
 - MyApp
```

To keep the example short, the blueprint only has an Outputs section and one expression, which is a call to `f_concat`. Notice that `f_concat` requires at least two parameters, but only one is provided. When the blueprint processor is run for this blueprint, this is displayed:

```
15:26:25 INFO: Output Processing
15:26:25 INFO: -----
15:26:25 INFO:
15:26:25 INFO: Output values specified: 1
15:26:25 ERROR: Value of ExampleValue: Expression could not be evaluated.
Error is...
15:26:25 ERROR: Function/macro concat requires parameter count between 2 and 99,
not 1
15:26:25 ERROR: Expression being evaluated at the time:
15:26:25 ERROR: {'f_concat': ['MyApp']}
15:26:25 INFO:
```

The blueprint processor attempts to print the value of `'ExampleError'` when the error occurs. After displaying the diagnostic, the expression being evaluated at the time of the error is displayed (in JSON notation).

Expressions are generally nested, and the error may occur within a subexpression. In that case, the diagnostic includes an expression stack, so that you can see the specific expression in error as well as the outer context. For instance, consider this contrived blueprint:

```
Outputs:
 ExampleValue:
```

```
Value:
 f_concat:
 - aaa
 - bbb
 - f_concat:
 - ccc
```

Notice that the expression involves two uses of `f_concat`. The outer use is correct, but there is an error with the inner use. When the blueprint processor is run for this blueprint, this is displayed:

```
15:42:50 INFO: Output Processing
15:42:50 INFO: -----
15:42:50 INFO:
15:42:50 INFO: Output values specified: 1
15:42:50 ERROR: Value of ExampleValue: Expression could not be evaluated.
Error is...
15:42:50 ERROR: Function/macro concat requires parameter count between 2 and 99,
not 1
15:42:50 ERROR: Expression evaluation stack follows (with failed expression at
bottom) ...
15:42:50 ERROR: -----
15:42:50 ERROR: | Expr: {'f_concat': ['aaa', 'bbb', {'f_concat': ['ccc']}]}
15:42:50 ERROR: | Expr: {'f_concat': ['ccc']}
15:42:50 ERROR: -----
15:42:50 INFO:
```

The expression evaluation stack shows the outer expression at the top and the expression with the error at the bottom. This stack only has two levels but in general there are many levels and each level of evaluation is shown. (The next example illustrates a multi-level expression evaluation stack.)

When a macro is evaluated, its definition is expanded and this expansion is shown in the expression evaluation stack. For instance, consider this contrived blueprint:

```
Outputs:
 ExampleValue:
 Value:
 f_getZoneURI:
 - myZone
 - jaas
```

In this case, the expression is a call to `f_getZoneURI` and the error is that the zone 'myZone' does not exist. When the blueprint processor is run for this blueprint, this is displayed:

```
15:57:17 INFO: Output Processing
15:57:17 INFO: -----
15:57:17 INFO:
15:57:17 INFO: Output values specified: 1
15:57:18 ERROR: Value of ExampleValue: Expression could not be evaluated.
Error is...
15:57:18 ERROR: Name not found: myZone
15:57:18 ERROR: find_one for predicate {'f_EQ': [{'f_pathc': ['name']], 'myZone'}}
failed
15:57:18 ERROR: Expression evaluation stack follows (with failed expression at
bottom) ...
15:57:18 ERROR: -----
15:57:18 ERROR: | Expr: {'f_getZoneURI': ['myZone', 'jaas']}
15:57:18 ERROR: | Expr: {'f_pathc': [{'f_findByName': [{'f_pathc': [{'f_findByName': [{'f_path': ['Cloud.service_family_types.elements']], 'jaas']}]}, 'jaas']}]}
```

```

 '.uri-.zones.elements']], 'myZone']], '.uri']]
15:57:18 ERROR: | Expr: {'f_findByName': [{'f_pathc': [{'f_findByName': [{'f_path': ['Cloud.service_family_types.elements']], 'jaas']],
 '.uri->.zones.elements']], 'myZone']]
15:57:18 ERROR: | Expr: {'f_findOne': [{'f_pathc': [{'f_findByName': [{'f_path':
 ['Cloud.service_family_types.elements']], 'jaas']], '.uri->.zones.elements']],
 {'f_EQ': [{'f_pathc': ['name']], 'myZone']], {'f_concat': ['Name not found: ',
 'myZone']}]}}
15:57:18 ERROR: -----
15:57:18 INFO:

```

When the error is detected, a diagnostic is displayed: “Name not found: myZone”. In this case, the error should be clear and you need not bother reading further.

But to illustrate how macro expansion and multi-level nested expressions are shown in the expression evaluation stack, we continue walking through the example. The top line of the stack shows the outer expression, which is what was specified in the blueprint as the expression for ‘Value:’. Because the `f_getZoneURI` intrinsic is implemented as a macro, line 2 of the expression evaluation stack shows the expression after macro expansion. (It is long, so line wrapping is needed.)

Line 3 shows the subset of line 2 that was being evaluated when the error occurred and line 4 shows the same expression after the `f_findByName` macro was expanded.

---

**Note:** Intrinsic macros use some internal functions, which is why you see names like `f_pathc`.

---

### 32.7.4 Cloud Resource Creation Error

In the previous example, the failure to create a resource was detected some time after the request was accepted. In some cases, the request to create may fail immediately.

If an error occurs while attempting to create a resource, you will see a diagnostic that identifies the resource, the error code, and some diagnostic text. In the following example, the attempt to create a resource name `MyJavaPf` failed with an HTTP code of 500. Reading further, you can see diagnostic text like, “cannot process request for ...”, “Unable to start the Instance deployment”, and “stack\_trace\_cause”: “java.lang.IllegalArgumentException: Unable to service executable from service template...”

```

14:43:37 INFO: MyJavaPf
14:43:37 INFO: /
14:43:37 INFO: / MyDatasource
14:43:37 INFO: / /
14:43:37 INFO: -----
14:43:37 INFO: | e | |
14:43:38 INFO: | es | |
14:43:42 ERROR: Failure creating resource MyJavaPf: 500
{
 "messages" :
 [
 {
 "date" : "2012-05-22T18:43:42+0000" ,
 "text" : "cannot process request for
oracle.sysman.emInternalSDK.ssa.cloudapi.ResourceInteraction@767d6a37 on
/em/cloud/jaas/
javaplatformtemplate/C086733BCCF2A4F3E040F10A716049A8" ,
 "hint" : " Unable to start the Instance deployment" ,
 "stack_trace_cause" : "java.lang.IllegalArgumentException: Unable to

```

```

 service executable from service template : C086733BCCF2A
4F3E040F10A716049A8\n\tat
 oracle.sysman.ssa.mwaas.model.util.remoteop.DPSubmissionHelper.
_createRequestMWaaSSetup(DPSubmissionHelper.
java:359)\n\tat oracle.sysman.ssa.mwaas.model.util.remoteop.DPSubmissionHelper.
_submitMWaaSSetupServiceRequest(DPSubmissionHelper.
java:616)\n\tat
oracle.sysman.ssa.mwaas.model.util.remoteop.DPSubmissionHelper.submitMWaaSSetupSe
rviceRequest(DPSubmissionHelper.
java:712)\n\tat
oracle.sysman.ssa.cloudapi.jaas.JavaPlatformInstance.GenerateJavaPlatformInstance
(JavaPlatformInstance.java:369)\n\tat
oracle.sysman.ssa.cloudapi.jaas.JavaPlatformTemplate.processRequest(JavaPlatformTe
mplate.java:128)\n\tat oracle.sysman.ssa.cloudapi.jaas.J
aaSServiceProvider.processRequest(JaaSServiceProvider.java:520)\n\tat
 oracle.sysman.emInternalSDK.ssa.cloudapi.EMCloudServlet.perform
(EMCloudServlet.java:226)\n\tat
oracle.sysman.emInternalSDK.ssa.cloudapi.EMCloudServlet.performPost
(EMCloudServlet.j" ,
 "stack_trace" :
 "oracle.sysman.emInternalSDK.ssa.cloudapi.CloudServiceException: Unable to start
the Instance deployment\n\t
at
oracle.sysman.ssa.cloudapi.jaas.JavaPlatformInstance.GenerateJavaPlatformInstance
(JavaPlatformInstance.java:373)\n\tat oracle.sysm
an.ssa.cloudapi.jaas.JavaPlatformTemplate.processRequest(JavaPlatformTemplate.
java:128)\n\tat oracle.sysman.ssa.cloudapi.jaas.JaaSer
viceProvider.processRequest(JaaSServiceProvider.java:520)\n\tat
 oracle.sysman.emInternalSDK.ssa.cloudapi.EMCloudServlet.
perform(EMCloudServlet.java:226)\n\tat
oracle.sysman.emInternalSDK.ssa.cloudapi.EMCloudServlet.performPost
EMCloudServlet.java:363)\n\tat oracle.
sysman.emInternalSDK.ssa.cloudapi.rest.AbstractRestServlet.doPost(AbstractRestServ
let.java:134)\n\tat
javax.servlet.http.HttpServlet.service(HttpServlet.java:727)\n\tat
 javax.servlet.http.HttpServlet.service(HttpServlet.java:820)\n\tat
weblogic.servlet.internal.
StubSecurityHelper$ServletServiceAction.run(StubSecurityHelper.java:227)\n\tat
weblogic.servlet.internal.StubSecurityHe
 }
]
}

14:43:42 INFO: | CF | |
14:43:42 INFO: -----
14:43:42 INFO:
14:43:42 ERROR: Create of resource MyJavaPf failed

```

In this example, the HTTP code is 500. Any code that begins with a 5 indicates that the cloud server encountered an unexpected exception. This could be due to an environment issue or even a bug in the server software. Since a 5xx code reflects a server error, you should contact the self service administrator.

In other cases, you may see a 4xx error code, which is returned when the client seems to have erred. In such cases, you should check the 'hint' and 'message' information for clues as to what went wrong, because you may be able to correct an error you made.

The blueprint processor lists the diagnostic information it receives, but for security reasons, the cloud server may not provide sufficient information to diagnose the issue.

In this case, you should contact the self service administrator, who in turn can often diagnose the issue by reviewing the log files for the cloud request.

## 32.8 Simulation Mode

The blueprint processor simulation mode can be used to aid in developing and testing blueprints. In this mode, the requests normally sent to the cloud server are simulated as well as the results returned by the server. Otherwise the blueprint processing logic is the same. To run the blueprint processor in this mode, you simply do not specify a cloud URI, i.e. do not use the `-c` option on the command line.

One benefit of simulation mode is the speed with which you can run a blueprint and try variations. Normal running of blueprints involves cloud requests for which the processing may be quite time consuming. When in simulation mode, the default behavior is that requests to create each resource consume 2 seconds and then succeed.

Another benefit is the ability to test various possibilities. For each resource, you can specify the simulated processing time as well as whether the request succeeds or fails. To do this for a given resource, use the Simulation attribute when defining a resource. For example:

```
MyJavaServer1:
 Container:
 f_getTemplateURI:
 - Small WLS
 - jaas
 Properties:
 destination_zone:
 f_getZoneURI:
 - MyZone
 - jaas
 params:
 user: app_user
 password: pw_you_should_change
 Simulation:
 delay: 3
 result: f
```

In the above example, creation of MyJavaServer1 will fail after 3 seconds.

## 32.9 Debugging with the Blueprint Processor

In addition to running the blueprint processor such that it deploys the blueprint and runs to completion, there are mechanisms you can use to debug blueprints. These are akin to mechanisms you may have used for debugging other applications, like print statements and the use of a debugger to interactively display values used by your application.

To enter the debugger at a particular point of execution, you can use either of two mechanisms. The simpler approach, which will usually be sufficient, is to use command line options that cause execution to pause between processing phases. These are called “pause points” and are described below. The alternative approach enables you to break at a more specific point, such as just prior to evaluating an expression for a specific resource’s property. To do this, you edit the blueprint to include a breakpoint.



When either a pause point or breakpoint is reached, control is transferred to the “debugger”. In the debugger, you enter various commands to display contents of the blueprint as well as that of the cloud to which you are connected.

### 32.9.1 Printing Intermediate Results

The essence of blueprint processing is to evaluate expressions and create resources once all required expressions have been evaluated. At any point during evaluation of an expression, you may wish to see some intermediate results to confirm the value is what you expected. To do so, you use the intrinsic function `f_print`.

Wherever an expression can appear in a blueprint, you simply nest it in a call to `print`. Optionally, you can include a second text message argument. When `print` is encountered, the text message and expression value are printed.

#### Examples

This (contrived) example shows how you plan to use a lookup table to access a template name for use in a call to `f_getTemplateURI`.

```
Data:
 MyTemplates:
 - {name: DbTemplate, type: dbaas}
 - {name: MWTemplate, type: jaas}
 ...
Resources:
 MyDB:
 Container:
 f_getTemplateURI:
 - f_path:
 - 'Data.MyTemplates[0].name'
 - dbaas
 Properties:
 ...
```

Suppose the code is not behaving as you intend, and you want to view the intermediate results before passing the name to `f_getTemplateURI`. Wrap the expression in a call to `f_print` like this:

```
Data:
 MyTemplates:
 - {name: DbTemplate, type: dbaas}
 - {name: MWTemplate, type: jaas}
 ...
Resources:
 MyDB:
 Container:
 f_getTemplateURI:
 - f_print:
 - f_path:
 - 'Data.MyTemplates[0].name'
 - dbaas
 Properties:
 ...
```

At runtime, the value of the expression is printed:

```
16:52:05 INFO:
16:52:05 INFO: Resource State Timeline
16:52:05 INFO: -----
...
16:52:05 INFO:
```

```

16:52:05 INFO: MyDB
16:52:05 INFO: /
16:52:05 INFO: -----
16:52:05 INFO: | e |
>>> Print-point:
 Value = DbTemplate
...

```

A print point message can also be provided, which is useful when you have multiple print points in your blueprint:

```

Data:
 MyTemplates:
 - {name: DbTemplate, type: dbaas}
 - {name: MWTemplate, type: jaas}
...
Resources:
 MyDB:
 Container:
 f_getTemplateURI:
 - f_print:
 - f_path:
 - 'Data.MyTemplates[0].name'
 - My printpoint for template name
 - dbaas
 Properties:
...

```

At runtime, the value of the expression is printed:

```

16:52:05 INFO:
16:52:05 INFO: Resource State Timeline
16:52:05 INFO: -----
...
16:52:05 INFO:
16:52:05 INFO: MyDB
16:52:05 INFO: /
16:52:05 INFO: -----
16:52:05 INFO: | e |
>>> Print-point: My printpoint for template name
 Value = DbTemplate
...

```

## 32.9.2 Pause Points

The easiest way to specify points at which to enter the debugger by using **pause points**. These are specified via command line options as described in [Section 32.3, "Running the Blueprint Processor"](#). You can specify that the blueprint processor pause and the debugger entered at any of these points:

- Prior to evaluating the Input section and prompting for Input parameters
- Prior to evaluating the Resources section
- Prior to evaluating the Output section
- Prior to termination, but after output processing or detecting an error that will terminate processing
- When an error is encountered, just prior to termination

Once the debugger is entered, you can use the commands described in [Section 32.9.4, "Debugger Commands"](#) below. To continue blueprint execution, enter the **continue** command.

#### Example 1

Suppose you just want to browse the cloud resources at your server. Specify the **-I** option, which drops you into the debugger before attempting any blueprint processing

#### Example 2

It is often useful to specify the **'-E'** (or **'—error\_debug'**) option, which drops you in the debugger if an error is encountered. (Otherwise, execution simply terminates.)

### 32.9.3 Breakpoints

Breakpoints are defined via the `f_breakpoint` intrinsic function as described in [Section 32.6.7.4, "Debugging Intrinsics"](#). Whenever evaluation encounters an `f_breakpoint` invocation, the optional text string parameter value is printed and the debugger entered.

### 32.9.4 Debugger Commands

When you enter the debugger, you will see a prompt of **"Paused:"**. At this prompt, there are several commands you can use including **"help"** or **"h"**. For example:

```
Paused: h
Commands are...
 p[ath] <path expression>: evaluate path expr
 e[val]: read & evaluate blueprint expression
 c[ontinue]: continue blueprint instantiation
 x[it]: exit blueprint processor
 h[elp]: (this command)
If first token isn't a command, the line is treated as a path expression
Paused:
```

#### 32.9.4.1 Path Command

The **path** command is used to evaluate arbitrary path expressions as described earlier in [Section 32.6.6.1, "Path Expressions"](#). (Unlike other debugger commands, the **"p"** or **"path"** keyword is not required.)

#### Example: Viewing values in blueprint

You can view values in your blueprint such as the value of an input parameter:

```
C:\bp> bp_processor.py -c https://...:15430/em/cloud -u sysman -p sysman -R
xyzApp.yml
...
```

#### Blueprint Processor - Invocation Summary

```

Cloud URI: https://hostname.us.oracle.com:15430/em/cloud
User: sysman
Blueprint file: xyzApp.yml
Timeout: 90 minutes, 0 seconds
Refresh frequency: 15 seconds
Inputs:
Pause points: Resources
Debug logging: False
Instance name: default_instance_name
Graphical report dir:
```

```

Versions:
 Blueprint processor: 12.1.0.5, 10-Oct-2012
 Cloud protocol: 10001

18:28:14 INFO: Connecting to cloud: https://hostname.us.oracle.com:15430/em/cloud

Input Parameter Value Entry

 Zone to use for db (Zone1):
 Password to use for db (welcome1):

...Pause point, prior to processing Resources section...
For command info, enter (h)elp

Paused: path Inputs
DbPassword:
 DefaultValue: welcome1
 Prompt: Password to use for db
 Sensitive: true
 Type: String
 Value: mySecret
DbZone:
 DefaultValue: Zone1
 Prompt: Zone to use for db
 Type: String
 Value: Zone1

Paused: Inputs.DbPassword.Value
mySecret

Paused:
```

In the above example, the simple test blueprint specifies two input parameters, DbZone and DbPassword.

- When prompted, you accepted the default, for the first parameter by pressing **Enter**. For the second, you entered your password.
- Notice that your command line options included “-R”, which tells the blueprint processor to pause just prior to evaluating the Resources section of your blueprint. The **Paused** prompt appears, and you enter the **path** command with the path expression **Inputs**. The value of the Inputs section of the blueprint is then printed, namely the 2 input parameters and their values. The values include both those provided by the blueprint and the current runtime values, in this case **Zone1** and **mySecret**.
- You then simply entered a path expression. The **path** command is assumed if no explicit command is entered.

#### 32.9.4.2 Continue Command

The continue command is used to resume blueprint processing.

Example

Continuing the previous example...

```

Paused: Inputs.MyNum
{Sensitive: true, Type: Number, Value: '123'}

Paused: continue
```

### 32.9.4.3 Exit Command

The **exit** command terminates the blueprint processor.

### 32.9.4.4 Eval Command

The **eval** command is used to evaluate any expression you can include in your blueprint.

Example

Suppose you are debugging your blueprint and it appears to be failing when looking up a template by name. You can use the **eval** command to evaluate expressions that appear in your blueprint.

First you try executing the expression of interest as it appears in your blueprint:

```
Paused: e
 Eval: f_getTemplateURI:
 Eval: - JaaS Template
 Eval: - jaas
 Eval:
18:05:01 ERROR: Name not found: JaaS Template
Expression evaluation stack follows (with failed expression & diagnostic at
bottom) ...

| Expr: {'f_getTemplateURI': ['JaaS Template', 'jaas']}
| Expr: {'f_path': [{'f_findByName': [{'f_path': [{'f_findByName': [{'f_path':
['Cloud.service_types.elements']}, 'jaas']], '.uri->
.service_templates.elements']}, 'JaaS Template']], '.uri']}
| Expr: {'f_findByName': [{'f_path': [{'f_findByName': [{'f_path':
['Cloud.service_types.elements']}, 'jaas']], '.uri->.service
_templates.elements']}, 'JaaS Template']}
| Expr: {'f_findOne': [{'f_path': [{'f_findByName': [{'f_path': ['Cloud.service
types.elements']}, 'jaas']], '.uri->.service
_templates.elements']}, {'f_EQ': [{'f_path': ['name']}, 'JaaS Template']},
{'f_concat': ['Name not found: ', 'JaaS Template']}]}}
| End of stack for error message: Name not found: JaaS Template

Paused:
```

In the above example, you first enter the **eval** or **e** command. Then you enter the expression. Note that indentation is significant, as it always is in YAML.

You see the same diagnostic you got when processing your blueprint, but now you can experiment with other values. Eventually, you realize that the template was created with two spaces in the name. You try with that name. It works, and the result of the expression evaluation, in this case a URI, is displayed.

```
Paused: e
 Eval: f_getTemplateURI:
 Eval: - JaaS Template
 Eval: - jaas
 Eval:
/em/cloud/jaas/javaplatformtemplate/C086733BCCF2A4F3E040F10A716049A8
Paused:
```

## 32.10 Tips and Hints

This section lists a few tips and hints that may be useful while working with blueprints

### 32.10.1 Editing YAML - Notepad ++ Example

YAML documents use indentation to denote containment semantics. This may affect your choice of editor or editing options. For instance, YAML does not allow tabs, so you should disable any editor options that cause automatic tab insertion.

As an example, suppose you use Notepad++ (<http://notepad-plus-plus.org>) . You would set the "replace by spaces" setting in Preferences -> Language Menu/Tab Settings. Better still, if your file has a suffix of ".yaml", Notepad++ sets options to be suitable for YAML. For instance, it colorizes the text based on YAML syntax. If you do not use ".yaml", you can manually set the language to YAML. (Settings -> Preferences -> Language Menu -> ...)

### 32.10.2 YAML Duplicate and Name / Value Pairs

YAML requires that name/value pairs at the same level use unique names. Any duplicates override earlier occurrences.

For instance, these two blueprints are equivalent:

```
Data:
 Password: doNotChangeMe
 UserId: QA_user
 Password: changeMe
Resources:
 ...

Data:
 UserId: QA_user
 Password: changeMe
Resources
 ...
```

### 32.10.3 Explicit Dependencies

Suppose your blueprint defines two resources X and Y. If creation of X depends on the successful creation of Y, there will usually be a data-dependency between the two. However, if that's not the case, you can include anywhere in the definition of ResourceX an expression like ...

```
f_getResourceAttr:
- ResourceY
- uri
```

### 32.10.4 Hint: Use '-T' Option

While developing and testing your blueprint, it's a good idea to use the -T command line option. This tells the blueprint processor to drop you into the debugger prior to terminating for any reason. Should any results be unexpected, whether an outright error or just unexpected output, you can use the debugger to investigate.

### 32.10.5 Help Forums

Other questions? Post them at  
<https://forums.oracle.com/forums/forum.jspa?forumID=220>.

---

---

# Index

## A

---

- accessing the security class, 11-12
- acknowledging OVM events, 4-41
- activating and deactivating policies, 5-7
- adding
  - data sources to WebLogic domain, 17-8
  - hosts to PaaS infrastructure, 9-2
  - middleware pools to chargeback application, 16-16
  - Oracle homes to middleware pool, 16-4
  - Oracle VM Manager, 4-6
  - Oracle VM Manager, picture of, 4-7
- adding target
  - Chargeback, 22-15
- Administration Group feature, 1-4
- administrative operations, performing on storage repository, 4-27
- APIs
  - error response messages, 24-8
- Application Home page, viewing, 17-10
- assemblies, cloud setup, 4-29
- assembly
  - binary file prerequisite, 8-21
  - components, creating, 8-21, 20-2
  - definition of, 8-21, 8-28
  - deploying, 8-28
  - instance, definition of, 8-28
- Assembly Deployment
  - Deployment Configuration, picture of, 8-31
  - Instance Details, picture of, 8-28
  - Network and Agent, picture of, 8-30
  - Review page, picture of, 8-36
- assembly instance
  - specifying public and private networks, 8-29
- assembly network, definition of, 8-29
- AssemblyInstance, definition of, 26-3
- assigning
  - users, 3-10
  - virtual drives, .iso files, 4-29
- authentication
  - protocol, 3-13
  - schemes, 3-13
- Automatic Storage Management (ASM), 12-28, 12-29

## B

---

- BI Publisher
  - integrate with Chargeback, 22-24
  - integrate with Consolidation Planner, 23-12
- Blueprints
  - About, 31-1
  - Deploy, 31-4
  - Graphical Representation, 31-14
  - Uses, 31-2
- business hierarchy
  - Chargeback
    - Chargeback
  - business hierarchy, 22-2
  - import from LDAP server, 22-14
- business hierarchy. See also cost center

## C

---

- Cache-Control, 24-6
- catalog archives, importing, 3-7
- charge plan
  - assign to target, 22-19
- charge trends
  - Chargeback, 22-21
- Chargeback
  - add target, 22-15
  - application, adding middleware pools to, 16-16
  - assign charge plan to target, 22-19
  - assign cost center to target, 22-18
  - BI Publisher, 22-24
  - cloud computing, 22-2
  - configuring, 16-16
  - cost center, 22-2
  - currency, 22-19
  - details, viewing, 7-5
  - report generation, 22-23
  - self service applications, 22-4
  - setting up, 21-6
  - shared resources, 22-1
  - target resource usage, 22-22
  - universal metrics, 22-2
  - universal rate adjustment, 22-11
  - usage and charge trends, 22-21
  - usage metering, 22-2
- Chargeback metrics

- configuration-based, 22-3
- usage-based, 22-3
- Chargeback target
  - dedicated, 22-3
  - shared, 22-3
- cloning
  - guest virtual machine, 8-14
  - middleware home from gold image, 15-7
- cloud
  - APIs
    - common behaviors, 24-4
  - computing, 24-1
  - Database Self Service Portal, 12-1
  - defining performance policies, 5-4
  - definition of, 1-8
  - infrastructure, 4-1
  - installing plug-ins, 3-5
  - management
    - faster deployment, 1-1
    - provisioning, 1-2
    - quality of service, 1-1
    - resource elasticity, 1-2
  - Oracle Management Agent and Cloud Control, 9-2
  - policies
    - definition of, 5-1
    - monitoring, 5-2
    - policy groups, 5-9
    - types of, 5-1
    - viewing, 5-2
    - viewing for a target, 5-3
  - resource data models, 25-3
- cloud computing
  - Chargeback, 22-2
- Cloud Control
  - and Management Agent, 9-2
- Cloud Policy Home page, picture of, 5-2
- Cloud Self Service Portal Home page, 7-1
- Cluster Database Home page
  - QoS status on, 13-6
  - viewing, 13-5
- collating targets into groups for better manageability, 1-4
- common behaviors, cloud APIs, 24-4
- common resource attributes, APIs, 24-7
- configuring
  - chargeback, 16-16
  - chargeback, sample scenario, 16-16
  - database connectivity, 17-8
  - database request settings, 12-5
  - deployment procedure, 15-11
  - guest virtual machine, 8-16
  - listener running from same Oracle Home, 11-4
  - listener running from SIHA Oracle Home, 11-5
  - middleware home gold image, 15-10
  - privilege delegation settings, 3-13
  - request settings, 16-8
  - self service login page, 3-14
  - Self Service Portal, 6-1
  - Server Load Balancer (SLB), 3-17
  - software library user, 4-39
  - storage servers, 11-6
  - upload file location, Software Library, 3-2
  - VLAN groups, 4-16
  - WebLogic Domain Provisioning Profile, 15-3
  - YUM repository, 4-40
- connecting to Management Agent, 11-15
- consolidation planner
  - using, 23-4
- consolidation constraints, 23-3
- Consolidation Planner
  - BI Publisher, 23-12
- consolidation planner
  - overview, 23-1
- Consolidation Planner, Enterprise Manager, 1-3
- consolidation project
  - creating, 23-5
  - definition, 23-2
- consolidation scenario
  - custom, 23-7
  - definition, 23-2
  - evaluating, 23-7, 23-11
  - pre-configured, 23-6
- Contractual Service Level Agreements (SLAs), 1-4
- cost center
  - assign to target, 22-18
  - Chargeback, 22-2
  - create, 22-13
- Cost Center target property, 22-12
- Create Data Source page, picture of, 17-8
- Create Database Pool page, picture of, 12-4
- Create Performance Policy page, picture of, 5-5
- Create Schedule Policy page, picture of, 5-6
- Create Storage Repository pop-up, picture of, 4-27
- Create Template Upload Files page, 8-23
- Create Virtual Disk Upload Files page, 8-26
- Create Zone page, 4-37
- creating
  - assembly components, 8-21
  - automated DRS and DPM policies, 5-15
  - custom SSA user role, 3-12
  - data source, 17-8
  - database pool, 12-4
  - database pool for schema as service, 12-29
  - database provisioning profile, 12-8
  - database service template, 12-21
  - deployment plans, 20-3
  - guest virtual machine, 8-11
  - ISO component, 8-25
  - logical networks, 4-17
  - middleware home, 15-7
  - middleware pool, 16-3
  - multiple virtual LAN (VLANs), 4-16
  - named credentials, 11-3
  - network profile, 4-19
  - networks, Configure Ports and VLAN Interfaces page, 4-18
  - networks, prerequisites for, 4-17
  - Oracle Middleware Home Gold Image, 15-9
  - PaaS Infrastructure Zone, 9-2, 9-4



- Paas pool in PaaS infrastructure, 9-7
- provisioning profiles, 15-2
- PXE bootable guest virtual machines, 8-43
- repository export, prerequisites, 4-31
- schema on a database, 12-29
- service template, 16-10
- service template for schema as service, 12-31
- service template using DBCA template profile, 12-27
- service template using RMAN backup profile, 12-25
- snap clone of large database, 12-2
- storage repository, 4-26
- storage servers, illustration of, 4-20
- template components, 8-23
- test administrators, 20-2
- users, 3-12
- virtual disk component, 8-26
- virtual machines, 8-28
- virtual server pool, 4-33
  - Virtual IP address, 4-36
- virtual server pools, 4-34
- VLAN group, 4-16
- WebLogic domain provisioning profile, 15-3
- zones, 4-37, 8-6
  - zones, illustration of, 4-37
- creating a cost center, 22-13
- creating a universal charge plan, 22-9
- creating an extended charge plan, 22-9
- credentials
  - Management Agent, 11-15
  - setting for OMS requirement, 8-22
  - setting up for provisioning, 11-3
  - types of, 11-3
- currency
  - Chargeback, 22-19
- custom SSA user role, creating, 3-12
- customized
  - Post-Login page, picture of, 3-16
  - SSA Login page, picture of, 3-16
- customizing
  - deployment procedure, 15-11
  - storage agent proxy, 11-9

## D

---

- data sharing, 22-25
- data source
  - creating, 17-8
  - editing, 17-9
- database
  - as a service, creating pool for, 12-4
  - configuring request settings, 12-5
  - connectivity
    - configuring, 17-8
  - creating
    - pool for schema as service, 12-29
    - schema on, 12-29
    - service template, 12-21
    - service template for schema as service, 12-31

- deploying, 11-4
- pool, creating, 12-4
- provisioning, 12-2
  - profile, 12-8
  - using DBCA templates, 12-3
  - using RMAN backup, 12-3
  - using snapshots, 12-2
- provisioning profile, 12-8
  - creating using DBCA template, 12-16
  - creating using existing RMAN backup, 12-14
  - creating using export schema objects, 12-18
  - creating using RMAN backup, 12-11
  - creating using snapshots, 12-9
- requesting, 13-3
- self service portal, 13-1
- services, definition of, 13-4
- starting, 13-5
- stopping, 13-5
  - using self service portal, 13-1
- Database as a Service (DBaaS)
  - service model, 1-6
  - setting up and using, 2-4
- Database Cloud Home page, 10-1
- Database Cloud Self Service Portal, 12-1
- Database Instance Home page, viewing, 13-4
- Database Pool Home page, viewing, 14-1
- Database Request Dashboard page, 10-3
- Database Self Service Portal, 13-1
- Database Service Home page, viewing, 13-4
- Database-as-a-Service (DBaaS) cloud, 1-3
- databases software
  - provisioning, 11-3
- DBCA template profile, creating service template using, 12-27
- dedicated target
  - Chargeback, 22-3
- default SDK, overriding, 11-10
- defining
  - performance policies, 5-4
  - policy groups, 5-9
  - quotas, 16-9
  - request purging policy, 4-40
- roles
  - EM\_CLOUD\_ADMINISTRATOR, 3-10
  - EM\_SSA\_ADMINISTRATOR, 3-11
  - EM\_SSA\_USER, 3-11
  - schedule-based policy, 5-6
  - storage QoS, 4-38
- deleting
  - deployed application, 17-8
  - PaaS Infrastructure Zone, 9-7
  - zones, 8-7
- Deploy Application page, picture of, 17-7
- deploying
  - an assembly, 8-28
  - database, 11-4
  - ISO image, 8-41
  - middleware service, 17-4
  - plug-ins, 3-9
  - preconfigured Java EE application, 17-7

- templates, 8-37
- virtual machine
  - infrastructure requirements, 8-20
  - provisioning, 8-20
- deployment
  - information
    - viewing, 21-14
  - plans, creating, 20-3
  - procedure, configuring and saving, 15-11
- Deployment Configuration page, 8-38
- disabling snap clone, 11-17
- discovering
  - Oracle VM Manager, 4-9
  - Oracle VM Server, 4-13
  - virtual servers, illustration of, 4-14
- Distributed Power Management (DPM)
  - policies, 5-14
- downloading plug-ins, 3-6
  - offline mode, 3-7
  - online mode, 3-6
- DRS and DPM policies, creating automated
  - policies, 5-15
- Dynamic Resource Scheduling (DRS) policies, 5-13

## E

---

- editing
  - data source, 17-9
  - existing MWaaS service template, 16-12
  - guest virtual machine, 8-16
  - PaaS Infrastructure Zone, 9-6
  - storage ceiling, 11-19
  - virtual server, 8-9
  - virtual server pools, 4-36
  - zones, 8-6
- EM\_INFRASTRUCTURE\_ADMIN privilege, 3-5
- EM\_SSA\_ADMINISTRATOR role, 12-1
- EM\_SSA\_ADMINISTRATOR, and provisioning
  - database software, 11-3
- EMCLI, setting up, 3-5
- enabling
  - middleware as a service, illustration of, 15-1, 15-12
  - snap clone, 11-17
- enabling metric collection, 22-16
- Enterprise Manager
  - Incident Management system, 1-4
- Enterprise Manager Consolidation Planner, 1-3
- Enterprise Manager Store, available functional
  - updates, 3-3
- Enterprise User Security (EUS), 3-13
- Enterprise user security based authentication, 3-13
- error response messages, APIs, 24-8
- extended charge plan
  - Chargeback
    - Chargeback
      - extended charge plan, 22-2
    - create, 22-9
    - update, 22-12

## F

---

- faster deployment, enabling, 1-1
- file systems, illustration of, 4-20

## G

---

- general privileges, granting, 11-12
- generating reports
  - Chargeback, 22-23
- granting
  - general privileges, 11-12
  - privileges for provisioning, 11-13
  - roles, 11-13
  - target privileges, 11-13
- guest virtual machine
  - cloning, 8-14
  - creating, 8-11
  - definition of, 8-12
  - deploying, 8-19
  - deployment options, 8-20
  - editing, 8-16
  - enabling high availability, 8-33
  - infrastructure requirements, 8-20
  - instances
    - modifying resource configuration
      - parameters, 8-32
  - prerequisites for creating, 8-11
  - provisioning using Oracle VM templates, 8-37
  - provisioning using OVA, 8-28
  - restarting, 8-13
  - restricting amount of physical CPU, 8-33
  - saving as template, 8-18
  - specifying priority to access physical CPUs, 8-33
  - starting, 8-13
  - stopping, 8-13
- Guest Virtual Machine Home page, 8-12

## H

---

- HTTP Status Codes, 24-6
- hypervisor, 8-9

## I

---

- IaaS
  - templates, 3-1
  - virtual assemblies, repository for, 3-1
- importing
  - catalog archives, 3-7
  - plug-in archives, 3-8
- Incident Management system, 1-4
- Infrastructure as a Service (IaaS)
  - service model, 1-5
- Infrastructure as a Service (PaaS)
  - setting up and using, 2-1
- Infrastructure Cloud Home page
  - picture of, 4-7, 8-2
  - viewing, 8-1
- Infrastructure-as-a-Service (IaaS) cloud, 1-3
- installing

- Management Agent on an unmanaged server, 9-2
- Oracle Management Agents on unmanaged hosts, 9-2
- plug-ins, 3-5
- Intelligent Platform Management Interface (IPMI), 8-10
- ISO component, creating, 8-25
- ISO Deployment page, picture of, 8-41

## J

---

- Java EE component, uploading to Software Library, 17-9
- JavaScript Object Notation, 24-4

## L

---

- LDAP server
  - import business hierarchy, 22-14
- logical
  - networks, creating, 4-17
  - pools, 1-6
  - zones, 1-6

## M

---

- machine sizes, setting up, 4-39
- maintaining virtual servers, 8-11
- maintenance mode, virtual servers, 8-11
- MANAGE\_SELF\_UPDATE privileges, 3-4
- Management Agent credentials, 11-15
- managing
  - storage access privileges, 11-19
  - storage allocation, 11-19
  - storage servers, advisory information, 11-5
  - virtualization targets
    - OVM Manager Home page, 8-3
  - zones, 8-5
- Maximum Transfer Unit size, setting, 4-16
- metering and chargeback features, 1-5
- Microsoft Active Directory based authentication, 3-13
- middleware
  - creating middleware home, 15-7
  - deployment of from gold image, 15-7
  - enabling as a service (MWaaS), 15-1, 15-11
  - gold image, 15-7
  - viewing, 10-1, 10-3
  - viewing Middleware Pool Home page, 18-1
- Middleware and Database Cloud Home page, picture of, 10-2
- Middleware and Database Request Dashboard page, picture of, 10-4
- Middleware as a Service (MWaaS)
  - enabling, 15-1, 15-11
  - service model, 1-6
  - setting up and using, 2-8
- Middleware Home Gold Image
  - component, 15-7
  - creating, 15-9
- middleware pool

- adding Oracle homes to, 16-4
- creating, 16-3
- definition of, 1-9, 16-3
- getting detailed information on, 18-1
- Middleware Pool Home page, viewing, 18-1
- Middleware Self Service Portal, purpose of, 17-1
- Middleware Service Home page
  - picture of, 17-6
  - purpose of, 17-6
  - viewing, 17-6
- middleware services
  - Middleware Self Service Portal, 17-1
  - requesting, 17-4
  - scaling up or scaling down, 17-8
- Middleware-as-a-Service (MWaaS) cloud, 1-3
- modifying configuration of guest virtual machine, 8-16
- monitoring cloud policies, 5-2
- MTU field, setting, 4-16
- MWaaS Application Home page, picture of, 17-10
- MWaaS Self Service Portal
  - setting up, 16-1
  - viewing, 17-1
- MWaaS service template, editing, 16-12
- My Oracle Support, Cloud management capabilities, 1-5

## N

---

- named credentials, creating, 11-3
- NetApp hardware licenses, 11-8
- NetApp storage server, setting up, 11-8
- NetApp Storage System, 12-2
- network bootable (PXE boot) virtual machine, 8-43
- network profile
  - DNS servers, 4-20
  - domain name, 4-20
  - gateway, 4-20
  - Netmask, 4-20
  - setting up, 4-19
  - specifying IP addresses, 4-20
- Network Profile page, picture of, 4-19
- network roles
  - cluster heartbeat, 4-18
  - live migrate, 4-18
  - server management, 4-17
  - storage, 4-18
  - Virtual Machine, 4-18
- network type, setting up, 4-38
- NetworkInterface, 26-7
- networks, setting up, 4-14
- New Service Request page, picture of, 17-4

## O

---

- OMS
  - Agent File System, 8-25
    - location, 8-27
  - setting credentials, 8-22
- OMS Shared File System location, 8-25, 8-27

- Oracle Access Manager (OAM) SSO, 3-13
- Oracle Cloud
  - logical pools, 1-6
  - logical zones, 1-6
- Oracle Internet Directory (OID) based
  - authentication, 3-13
- Oracle Listener, setting up for database hosts, 11-4
- Oracle Management Service (OMS), routing SSA
  - requests, 3-17
- Oracle Software Library
  - importing plug-in archives to, 3-8
  - setting up, 3-1
- Oracle Virtual Assembly Builder (OVAB), packaging
  - multi-tier platforms, 1-3
- Oracle VM Manager
  - adding, 4-6
  - registering, 4-6, 4-7
- Oracle VM Network, 4-14
  - Virtual Area Network (VLAN) Group, 4-16
- Oracle VM Server
  - creating virtual server pools, 4-34
  - discovering, 4-13
  - file systems, 4-20
  - patching, 4-40
  - storage servers, setting up, 4-20
  - virtual server, 4-13
  - virtual server pool, 4-33
  - virtualization software (hypervisor), 4-13
- out-of-the-box provisioning profiles, using, 15-5
- overriding
  - default SDK, 11-10
  - third-party server components, 11-11
- OVIM events, acknowledging, 4-41
- OVIM Manager
  - Home page, 8-3
  - Home page, picture of, 8-3
- OVIM Manager targets, synchronizing, 4-12
- OVIM Managers, list of, 8-5

## P

- PaaS database provisioning profiles, repository
  - for, 3-1
- PaaS infrastructure
  - adding hosts, 9-2
  - creating PaaS pool, 9-7
  - database cloud home page, 10-1
  - database request dashboard page, 10-3
  - Paas pool, picture of, 9-7
  - viewing middleware, 10-1, 10-3
- PaaS Infrastructure Zone, 1-9
  - creating, 9-2
  - creating middleware home, 15-7
  - Database Pool Home page, 14-1
  - definition of, 16-16
  - deleting, 9-7
  - editing, 9-6
  - enabling or setting up MWaaS, 16-16
  - middleware pool, 16-3
  - picture of, 9-3

- prerequisites, 9-3
- Roles page, 9-6
- Targets page, 9-5
- PaaS middleware deployment procedures, repository
  - for, 3-1
- patching, Oracle VM server, 4-40
- performance policies
  - defining, 5-4
- Platform as a Service (PaaS)
  - service model, 1-6
  - setting up
    - and using, 2-3
- plug-ins
  - archives, importing, 3-8
  - definition of, 3-5
  - deploying, 3-9
  - deploying in graphical mode, 3-10
  - deploying in silent mode, 3-10
  - downloading in offline mode, 3-7
  - downloading in online mode, 3-6
  - downloading to Software Library, 3-6
  - installing, 3-5
  - online and offline mode, 3-6
- policies
  - activating and deactivating, 5-7
  - defining groups, 5-9
  - Distributed Power Management (DPM), 5-14
  - Dynamic Resource Scheduling (DRS), 5-13
  - schedule-based, 5-6
  - viewing, 7-6
- policy group
  - defining, 5-7
  - definition of, 5-7
  - examples of, 5-11
- PowerBroker, 3-13
- Preboot Execution Environment (PXE), 8-43
- preferences, viewing, 7-6
- prerequisites
  - assembly binary file, 8-21
  - configuring YUM repository, 4-40
  - creating a zone, 4-37
  - creating an assembly component, 8-21
  - creating database provisioning profile, 12-9
  - creating middleware home, 15-7
  - creating middleware home gold image, 15-9
  - creating networks, 4-17
  - creating PaaS Infrastructure Zone, 9-3
  - creating PXE bootable guest virtual machines, 8-43
  - creating repository export, 4-31
  - provisioning guest virtual machine, 8-41
  - provisioning guest virtual machine using ISO image, 8-41
  - provisioning guest virtual machine using templates, 8-37
  - provisioning guest virtual machines, 8-20
  - registering storage servers, 11-6
  - Self Service Portal, 6-1
  - Self Update feature, 3-4
  - Testing as a Service, 21-2

- virtual server pools, 4-33
- WebLogic domain provisioning profile, 15-3
- Presented Servers page, picture of, 4-28
- privilege delegation settings, configuring, 3-13
- privilege delegation tools, Sudo and PowerBroker, 3-13
- privileges
  - and roles, granting, 11-12
  - EM\_INFRASTRUCTURE\_ADMIN, 3-5
  - for provisioning, granting, 11-13
  - MANAGE\_SELF\_UPDATE, 3-4
  - managing storage access, 11-19
  - NetApp hardware users, 11-7
  - versus permissions, 11-8
  - VIEW\_SELF\_UPDATE, 3-4
- Provision Oracle Database deployment procedure, 11-4
- provisioning
  - creating database profile, 12-8
  - credentials, setting up, 11-3
  - database, 12-2
  - database instance within PaaS Infrastructure Zone, 14-1
  - database software, 11-3
  - database using DBCA templates, 12-3
  - database using RMAN backup, 12-3
  - database using snapshots, 12-2
  - guest virtual machine
    - using ISO image, 8-41
  - guest virtual machine using ISO image, 8-41
  - guest virtual machine using OVA, 8-28
  - guest virtual machine using templates
    - prerequisites, 8-37
  - guest virtual machines, 8-20
  - middleware home deployment, 15-7
  - profiles, creating, 15-2
  - using out-of-the-box profiles, 15-5
  - virtualized systems, 8-19
- PXE bootable guest virtual machines, creating, 8-43

## Q

- QoS
  - Priority Class, 8-34
  - setting up, 4-37
  - status on Cluster Database Home page, 13-6
- quality of service, increasing, 1-1
- quota
  - definition of, 12-6, 16-9
  - sample validation log, 12-7
  - setting up, 12-6, 16-9

## R

- redeploying preconfigured Java EE application, 17-7
- redirecting SSA requests, 3-17
- rediscovering virtual server, 4-14
- reference gold images, 3-1
- Referenced File Locations, Software Library, 3-2
- registering

- Oracle VM Manager, 4-6
- storage servers
  - advisory information, 11-5
  - overview of, 11-6
  - prerequisites, 11-6
- Repository-based authentication, 3-13
- Representational State Transfer API, 24-1
- request purging policy, defining, 4-40
- request settings
  - configuring, 16-8
  - for database, configuring, 12-5
- requesting
  - middleware service, 17-4
  - or reserving servers, 7-7
  - server resources, 7-1
- resource attributes, APIs, 24-7
- resource elasticity, providing, 1-2
- restarting guest virtual machine, 8-13
- RESTARTING, VM data model, 26-6
- RESTful API, 24-1
- RESTful APIs, 24-4
- RESUMING, VM data model, 26-6
- RMAN backup profile, creating service template using, 12-25
- roles
  - defining, 3-10
  - definition of, 3-10
  - EM\_CLOUD\_ADMINISTRATOR, 3-10, 9-2
  - EM\_SSA\_ADMINISTRATOR, 3-11
    - provisioning database software, 11-3
  - EM\_SSA\_USER, 3-11
    - granting, 11-13
- roles and privileges, granting, 11-12

## S

- sample quota validation log, 12-7
- saving
  - deployment procedure, 15-11
  - guest virtual machine, 8-18
- Scalability Group, 26-4
- schedule-based policy, 5-6
  - examples of, 5-9
- schema
  - as a service, creating database pool for, 12-29
  - as a service, creating service template for, 12-31
  - creating on a database, 12-29
  - defining one or more, 14-2
  - requesting, 13-3
- Schema Pool Home page, viewing, 14-2
- Schema Service Template, selecting, 13-3
- security class, accessing, 11-12
- security, Enterprise User Security (EUS), 3-13
- self service
  - database portal, 13-1
  - login page, configuring, 3-14
- self service applications
  - Chargeback, 22-4
- Self Service Portal, 6-1
  - chargeback, illustration of, 6-6

- configuring, 6-1
- defining mapping between roles and other entities, 6-3
- defining set of templates or assemblies, 6-4
- Home page, picture of, 7-1
- machine sizes, 6-1
- prerequisites, 6-1
- purpose of, 3-11, 6-1, 7-1
- requesting or reserving servers, 7-7
- requesting server resources, 7-1
- setting up, 12-1
- setting up Chargeback Plan, 6-6
- viewing storage, 7-4
- Self Update
  - downloading assemblies and templates, 4-32
  - prerequisites, 3-4
- Self Update Console page, picture of, 4-33, 15-5
- self-service users, accessing services without IT intervention, 1-10
- server data model, 26-5
- Server Load Balancer (SLB), configuring, 3-17
- server pool, definition of, 1-9
- servers, requesting or reserving, 7-7
- service instance, deploying preconfigured Java EE application to, 17-7
- service template, 25-3
  - creating, 12-21, 16-10
  - creating using DBCA template profile, 12-27
  - creating using RMAN backup profile, 12-25
  - creating using snap clone profile, 12-21
  - definition of, 12-8, 12-21, 16-10
  - editing MWaaS, 16-12
- service, scaling up or scaling down, 17-8
- setting
  - credentials for OMS, 8-22
  - Maximum Transfer Unit size, 4-16
  - MTU field, 4-16
- setting up, 6-1
  - and using DBaaS, 2-4
  - and using IaaS, 2-1
  - and using MWaaS, 2-8
  - and using TaaS, 2-9
  - Chargeback Plan for Self Service Portal, 6-6
  - Chargeback Service, 21-6
  - database as a service, illustration of, 12-1
  - listener for database zone hosts, 11-4
  - machine sizes, 4-39
  - MWaaS Self Service Portal, 16-1
  - NetApp storage server, 11-8
  - network profile, 4-19
  - networks, 4-14
    - network profile, 4-19
    - network type, 4-19
  - provisioning credentials, 11-3
  - QoS, 4-37
  - QoS for network, 4-38
  - quotas, 12-6, 16-9
  - Self Service Portal, 6-1, 12-1
  - Software Library, 3-1
  - storage servers, 4-20
  - Sun ZFS storage server, 11-9
  - Testing as a Service, 20-1
  - virtualization infrastructure, 4-2
  - virtualized systems, 4-1
- setting up and using PaaS, 2-3
- shared resources
  - Chargeback, 22-1
- shared target
  - Chargeback, 22-3
- sharing Chargeback data, 22-25
- snap clone
  - creating service template, 12-21
  - enabling or disabling, 11-17
  - of large database, creating, 12-2
- Software Library
  - configuring upload file location, 3-2
  - creating virtualization components in, 8-21
  - Referenced File Locations, 3-2
  - setting up, 3-1
  - Upload File Locations, 3-2
  - uploading files from storage location, 8-25
  - uploading Java EE component to, 17-9
  - user, configuring, 4-39
- Software Library page
  - accessing, 3-1
  - picture of, 3-2
- software pool, definition of, 1-9
- SSA requests, redirecting, 3-17
- SSA User Portal page, picture of, 1-11
- SSA user role, creating custom role, 3-12
- SSO-based authentication, 3-13
- STARTED, VM data model, 26-6
- starting
  - database, 13-5
  - guest virtual machine, 8-13
  - virtual server, 8-10
- STARTING, VM data model, 26-6
- STOPPED, VM data model, 26-6
- stopping
  - database, 13-5
  - guest virtual machine, 8-13
  - virtual server, 8-10
- STOPPING, VM data model, 26-6
- storage agent proxy, customizing, 11-9
- storage allocation, managing, 11-19
- storage ceiling
  - definition of, 11-19
  - editing, 11-19
- storage entity, definition of, 1-9
- storage location types, 3-2
- storage pool, definition of, 1-9
- storage QoS, defining, 4-38
- Storage Quality of Service (QoS), setting up, 4-37
- storage repository
  - assemblies, illustration of, 4-29
  - creating, 4-26
  - ISOs, 4-29
  - Oracle Virtual Assemblies, 4-29
  - Oracle VM templates, 4-29
  - performing administrative operations, 4-27

- presenting to server pool, 4-28
- virtual disks, 4-29
- Virtual Disks, illustration of, 4-31
- VM files, 4-29
- VM Templates, illustration of, 4-30
- Storage Repository page, picture of, 4-26
- Storage Server page, picture of, 4-21
- storage servers
  - configuring, 11-6
  - deregistering, 11-18
  - managing, 11-18
  - managing access privileges, 11-19
  - registering, 11-13
  - selecting type of storage array, 4-21
  - synchronizing, 11-16
- storage vendor SDK, uploading, 11-10
- Sudo, 3-13
- Sun ZFS
  - hardware licenses, 11-9
  - storage server, setting up, 11-9
- Super Administrator roles, tasks to perform, 2-1
- SUSPENDED, VM data model, 26-6
- SUSPENDING, VM data model, 26-6
- synchronizing OVM Manager targets, 4-12

## T

---

- target privileges, granting, 11-13
- target property, for cost center, 22-12
- target resource usage
  - Chargeback, 22-22
- template components
  - creating, 8-23
  - virtualization types, 8-23
- Template Deployment - Deployment Configuration
  - page, 8-39
- template deployment, selecting target and
  - source, 8-38
- templates
  - deploying, 8-37
  - provisioning guest virtual machine, 8-37
- test
  - activity and usage information, 21-13
  - administrators, 19-2
    - creating, 20-2
    - tasks, 21-3
  - assets, 21-6
    - creating, 21-7
    - verifying, 21-9
  - creating, 21-10
  - designers, 19-2
  - environments, 21-6
    - creating, 21-7, 21-8
    - using, 21-9
    - verifying, 21-9
  - sharing, 21-13
  - testers, 19-3
  - trials
    - comparing results, 21-12
    - creating and running, 21-11

- monitoring, 21-12
- testing
  - home page, 21-14
  - setting up
    - applications, 21-3
    - custom test driver types, 21-4
    - Oracle Load Testing Test Drivers, 21-4
    - quotas, 21-5
- Testing as a Service
  - overview, 19-1
  - prerequisites, 21-2
  - setting up and using, 2-9
- third-party licenses, acquiring, 11-9
- third-party server components, overriding, 11-11
- Tier Deployment Configuration - Miscellaneous,
  - picture of, 8-35
- Tier Deployment Configuration Product
  - Configuration, picture of, 8-35
- topology of a zone, 4-37
- Transport Protocol, 24-4

## U

---

- undeploying an application, 17-8
- universal charge plan
  - Chargeback
    - Chargeback
      - universal charge plan, 22-2
        - create, 22-9
- universal metrics
  - Chargeback, 22-2
- updates available, Enterprise Manager Store, 3-3
- updating an extended charge plan, 22-12
- upgrading
  - virtual server, 4-41, 8-10
- Upload File Locations, Software Library, 3-2
- Upload to Software Library page, picture of, 17-9
- uploading
  - Java EE component to Software Library, 17-9
  - storage vendor SDK, 11-10
- URI Space, 24-4
- usage metering
  - Chargeback, 22-2
- usage trends
  - Chargeback, 22-21
- users
  - assigning, 3-10
  - creating, 3-12

## V

---

- VIEW\_SELF\_UPDATE privileges, 3-4
- viewing
  - Application Home page, 17-10
  - chargeback details, 7-5
  - cloud policies, 5-2
  - cloud policies for a target, 5-3
  - Cluster Database Home page, 13-5
  - Database Instance Home page, 13-4
  - Database Pool Home page, 14-1

- Database Service Home page, 13-4
- Infrastructure Cloud Home page, 8-1
- MWaaS Self Service Portal, 17-1
- OVIM Manager home page, 8-3
- policies, 7-6
- preferences, 7-6
- Schema Pool Home page, 14-2
- storage, Self Service Portal, 7-4
- Zone Home page, 8-5
- virtual
  - appliance images, 3-1
  - disk component, creating, 8-26
  - machine provisioning, 8-20
- virtual drives, assigning using .iso files, 4-29
- virtual machine
  - creating, 8-28
  - creating assembly component, 8-21
  - editing guest machine, 8-16
- Virtual Network Interface Card Manager, 4-15
- Virtual Network Interfaces (VNICs), 4-15
- virtual server
  - definition, 8-9
  - editing, 8-9
  - maintaining, 8-11
  - Oracle VM Server, 4-13
  - rediscovering, 4-14
  - starting, 8-10
  - stopping, 8-10
  - upgrading, 4-41, 8-10
- Virtual Server Home page
  - picture of, 8-9
  - viewing, 8-9
- Virtual Server Pool Home page, picture of, 8-8
- Virtual Server Pool Home page, viewing, 8-7
- virtual server pools
  - creating or editing, 4-33
  - definition, 8-7
  - editing, 4-36
  - prerequisites, 4-33
  - zone, creating, 4-37
- virtualization
  - infrastructure, 4-1
  - infrastructure, setting up, 4-2
  - managing targets, 8-3
  - software (hypervisor), 4-13
  - targets, viewing members, 8-5
  - types, template components, 8-23
- virtualization targets, managing members, 8-5
- virtualized systems
  - provisioning, 8-19
  - setting up, 4-1
- VLAN group
  - configuring, 4-16
  - creating, 4-16
- VNet, definition of, 26-7

## W

---

- WebLogic Domain Provisioning Profile
  - configuring, 15-3

- definition of, 15-2
- WebLogic JDBC data sources, configuring an application to, 17-8

## Y

---

- YUM repository
  - configuring, 4-40
  - upgrading virtual server, 8-10

## Z

---

- ZFS Storage Appliance, 12-2
- Zone Home page, 8-5
  - picture of, 8-6
- zones
  - and server pools, 1-8
  - Create Zone page, 4-37
  - creating, 4-37, 8-6
  - definition, 4-37
  - definition of, 1-8, 25-5
  - deleting, 8-7
  - editing, 8-6
  - managing, 8-5
  - prerequisites to creating, 4-37
  - topology of, 4-37