

Oracle® Enterprise Manager

Cloud Administration Guide

12c Release 2 (12.1.0.2)

E28814-03

October 2012

E28814-03

Copyright © 2011, 2012 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	xi
Audience	xi
Supported Plug-ins	xi
Documentation Accessibility	xi
Related Documents	xii
Conventions	xii

Part I Cloud Management with Enterprise Manager

1 Cloud Management - An Overview

1.1	Introduction to Cloud Management	1-1
1.2	Oracle Cloud Solution	1-1
1.3	Cloud Management Lifecycle	1-3
1.3.1	Planning	1-4
1.3.2	Setting Up	1-4
1.3.3	Building	1-5
1.3.4	Testing and Deploying	1-5
1.3.5	Monitoring and Managing	1-5
1.3.6	Metering, Charging, and Optimization	1-6
1.4	Cloud Resource Models	1-6
1.4.1	IaaS Cloud Anatomy	1-6
1.4.2	DBaaS and MWaaS Cloud Anatomy	1-7
1.5	Users and Roles	1-8
1.6	Enterprise Manager Console and the Self Service Portal	1-8
1.7	Cloud APIs	1-9

2 Getting Started with Cloud Management

2.1	Installing Virtualization and the Cloud Plug-ins	2-1
2.2	Deploying the Cloud Plug-ins	2-2
2.2.1	Deploying the Cloud Plug-ins to the Management Server	2-2
2.2.2	Deploying the Enterprise Manager for Oracle Virtualization Plug-in to the Management Agent 2-3	
2.3	Setting Up the Software Library	2-3
2.4	Using Self Update to Download Oracle VM Templates and Oracle Virtual Assemblies	2-5
2.5	Managing Users and Roles	2-6

2.5.1	Creating a Role	2-7
2.5.2	Creating a User and Assigning Roles	2-7
2.6	LDAP Authentication.....	2-8
2.7	Customizing the Self Service Login Page.....	2-8
2.7.1	Configuring the Self Service Login Page.....	2-8
2.7.2	Switching Back to the Enterprise Manager Login Page.....	2-9
2.7.3	Routing SSA Requests to a Specific OMS Pool.....	2-9

3 Using Consolidation Planner

3.1	Overview of Consolidation Planner.....	3-1
3.1.1	Key Concepts.....	3-2
3.2	Consolidation Constraints	3-3
3.2.1	Source Server Constraints.....	3-3
3.2.2	Destination Server Constraints	3-3
3.3	Using Consolidation Planner	3-4
3.3.1	Creating a Consolidation Project.....	3-4
3.3.2	Using a Pre-configured Consolidation Scenario	3-6
3.3.3	Creating a Custom Consolidation Scenario.....	3-7
3.3.4	Evaluating Consolidation Scenarios	3-9

4 Chargeback Administration

4.1	Overview of Chargeback	4-1
4.1.1	Why Implement Chargeback	4-1
4.1.2	Enterprise Manager Chargeback.....	4-2
4.1.3	Understanding Targets and Metrics	4-2
4.1.4	About Chargeback, Self Service Applications, and Zones	4-4
4.2	Setting Up Chargeback.....	4-5
4.2.1	Defining Charge Plans	4-7
4.2.1.1	Defining the Universal Charge Plan.....	4-7
4.2.1.2	Defining an Extended Charge Plan.....	4-8
4.2.2	Creating Cost Centers	4-11
4.2.2.1	Setting the Cost Center Property on the Target Home Page.....	4-11
4.2.2.2	Creating the Cost Center in Chargeback.....	4-12
4.2.2.3	Importing a Business Hierarchy from an LDAP Server	4-13
4.2.3	Adding Targets for Chargeback	4-14
4.2.4	Assigning Cost Centers to Targets	4-15
4.2.5	Assigning Charge Plans to Targets	4-16
4.2.6	Configuring Chargeback Settings	4-16
4.3	Accessing Chargeback Information	4-17
4.3.1	Following Usage and Charge Trends	4-18
4.3.2	Viewing a Target's Collected Metric Data	4-19
4.3.3	Generating and Distributing Chargeback Reports	4-20
4.3.4	Viewing Chargeback Information in the Self Service Portal.....	4-22
4.3.5	Sharing Chargeback Data with Other Applications.....	4-22

Part II Infrastructure as a Service

5 Setting Up Cloud Infrastructure

5.1	Getting Started.....	5-1
5.2	Registering the Oracle VM Manager.....	5-2
5.2.1	Registering Oracle VM Manager 3.1.x With Enterprise Manager	5-4
5.2.2	Upgrading from Oracle VM Manager 3.0.x	5-5
5.2.3	Discovering and Promoting the Oracle VM Manager (Optional)	5-5
5.3	Synchronizing the OVM Manager Targets (Optional)	5-6
5.4	Discovering an Oracle VM Server	5-6
5.5	Rediscovering a Virtual Server	5-7
5.6	Setting Up Networks	5-7
5.6.1	Generating MAC Addresses	5-8
5.6.2	Configuring a VLAN Group	5-9
5.6.3	Creating Networks	5-10
5.7	Creating a Network Profile	5-12
5.8	Registering Storage Servers	5-13
5.9	Creating a Storage Repository	5-16
5.9.1	Performing Administrative Operations on a Storage Repository	5-17
5.10	Presenting the Storage Repository	5-17
5.11	Importing Assembly, Template, and Virtual Disk Resources into a Repository.....	5-18
5.11.1	Creating a Repository Export (Optional)	5-21
5.12	Creating a Virtual Server Pool	5-21
5.12.1	Editing a Virtual Server Pool	5-24
5.13	Creating a Zone	5-24
5.14	Setting Up the Storage Quality of Service (QoS)	5-25
5.14.1	Setting Up the Network Type	5-26
5.14.2	Setting Up the Machine Sizes.....	5-26
5.14.3	Configuring the Software Library User.....	5-27
5.14.4	Defining the Request Archiving Policy	5-27
5.15	Patching the Oracle VM Server	5-27
5.15.1	Configuring the YUM Repository	5-27
5.15.2	Upgrading the Virtual Server	5-28
5.16	Acknowledging OVM Events	5-28

6 Setting Up and Using the Self Service Application for IaaS

6.1	Setting Up the Self Service Portal	6-1
6.2	Using the Self Service Portal.....	6-6
6.2.1	Viewing My Requests	6-7
6.2.2	Viewing My Servers	6-7
6.2.2.1	Server Details Page.....	6-8
6.2.2.2	Scaling Up a Tier Instance	6-8
6.2.2.3	Scaling Down a Tier Instance.....	6-9
6.2.3	Viewing Storage	6-9
6.2.4	Viewing Chargeback Details.....	6-10
6.2.5	My Library	6-10
6.2.6	Viewing Policies.....	6-11
6.2.7	Viewing My Preferences.....	6-11

6.3	Requesting Servers.....	6-12
-----	-------------------------	------

7 Administering and Monitoring the Cloud for IaaS

7.1	Viewing the Infrastructure Cloud Home Page.....	7-1
7.2	Viewing the OVM Manager Home Page.....	7-3
7.3	Viewing and Managing Members.....	7-5
7.4	Managing Zones.....	7-6
7.4.1	Viewing the Zone Home Page.....	7-6
7.4.2	Creating or Editing a Zone.....	7-7
7.5	Deleting a Zone.....	7-8
7.6	Virtual Server Pool Home Page.....	7-8
7.7	Managing Virtual Servers.....	7-9
7.7.1	Virtual Server Home Page.....	7-10
7.7.2	Editing a Virtual Server.....	7-11
7.7.3	Upgrading a Virtual Server.....	7-11
7.7.4	Starting and Stopping a Virtual Server.....	7-11
7.7.5	Performing Maintenance on the Virtual Server.....	7-12
7.8	Managing Guest Virtual Machines.....	7-12
7.8.1	Creating a Guest Virtual Machine.....	7-12
7.8.2	Guest Virtual Machine Home Page.....	7-13
7.8.3	Live Migrating a Guest Virtual Machine.....	7-13
7.8.4	Starting and Shutting Down a Guest Virtual Machine.....	7-14
7.8.4.1	Starting a Guest Virtual Machine.....	7-14
7.8.4.2	Restarting a Guest Virtual Machine.....	7-14
7.8.4.3	Stopping a Guest Virtual Machine.....	7-14
7.8.5	Cloning a Guest Virtual Machine.....	7-14
7.8.6	Editing a Guest Virtual Machine.....	7-16
7.8.7	Saving a Guest Virtual Machine as a Template.....	7-18
7.9	Deploying Guest Virtual Machines.....	7-19
7.9.1	Getting Started.....	7-20
7.9.2	Deployment Options.....	7-20
7.9.3	Prerequisites.....	7-21
7.9.3.1	Creating the Software Library Components.....	7-21
7.9.3.1.1	Creating an Assembly Component.....	7-21
7.9.3.1.2	Creating a Template Component.....	7-23
7.9.3.1.3	Creating an ISO Component.....	7-24
7.9.3.1.4	Creating a Virtual Disk.....	7-26
7.9.4	Provisioning Guest Virtual Machines Using Oracle Virtual Assemblies (OVA)....	7-27
7.9.5	Provisioning a Guest Virtual Machine Using Oracle VM Templates.....	7-36
7.9.6	Provisioning a Guest Virtual Machine Using an ISO Image.....	7-40
7.9.7	Creating PXE Bootable Guest Virtual Machines.....	7-42
7.10	Bare Metal Provisioning.....	7-44

8 Managing Cloud Policies

8.1	Managing Cloud Policies.....	8-1
8.2	Viewing Cloud Policies.....	8-2
8.3	Viewing Policies for a Target.....	8-2

8.4	Defining a Performance Policy	8-3
8.5	Defining a Schedule Based Policy	8-4
8.5.1	Example 1	8-5
8.5.1.1	Policy for Shutting Down Guest Virtual Machine	8-5
8.5.1.2	Policy for Starting Up Guest Virtual Machine	8-6
8.5.1.3	Creating a Policy Group (Start Up and Shut Down)	8-7
8.6	Activating and Deactivating Policies	8-7
8.7	Defining Policy Groups	8-8
8.7.1	Example 2	8-8
8.7.1.1	Policy to Scale Down a Tier Instance	8-8
8.7.1.2	Policy to Scale Up a Tier Instance	8-10
8.7.1.3	Creating a Policy Group (Scale Up and Scale Down)	8-10
8.8	Managing DRS and DPM Policies	8-10
8.8.1	Creating Automated DRS and DPM Policies	8-12

Part III Platform as a Service

9 About Platform as a Service

9.1	About Platform as a Service (PaaS)	9-1
9.2	Platform as a Service with Enterprise Manager	9-1

10 Using the Self Service Framework

10.1	Creating a PaaS Infrastructure Zone	10-1
10.1.1	Editing a PaaS Infrastructure Zone	10-4
10.1.2	Deleting a PaaS Infrastructure Zone	10-4
10.2	Viewing the Middleware and Database Cloud Home Page	10-4
10.3	Viewing the Middleware and Database Request Dashboard Page	10-6
10.4	Viewing All Requests Page	10-7
10.5	Viewing the Service Instances Page	10-8
10.6	Viewing the Service Template Details Page	10-9
10.7	Viewing the PaaS Infrastructure Zones Page	10-9
10.8	PaaS Infrastructure Zone Home Page	10-10

Part IV Database as a Service

11 About Database as a Service (DBaaS)

11.1	About Database as a Service (DBaaS)	11-1
11.2	Database as a Service in Enterprise Manager	11-2

12 Setting Up Database as a Service

12.1	Getting Started	12-1
12.2	Defining Roles and Assigning Users	12-2
12.2.1	Creating a Custom Role for Self Service Application Users	12-3
12.2.2	Creating a User and Assigning Roles	12-4
12.3	Adding Hosts	12-4

12.4	Configuring Privilege Delegation Settings.....	12-5
12.5	Configuring the Software Library	12-6
12.6	Setting Up Credentials for Provisioning	12-7
12.7	Provisioning Database Software.....	12-7
12.8	Creating Oracle Home and Configuring the Oracle Listener.....	12-8

13 Setting Up and Using the Self Service Application for DBaaS

13.1	Customizing and Saving the Database Deployment Procedure.....	13-1
13.2	Setting Up the Self Service Portal	13-3
13.3	Using the Database Self Service Portal	13-7
13.4	Requesting Databases.....	13-9
13.5	Troubleshooting Tips for Database Request Failures.....	13-11
13.5.1	Viewing the Output Log of a Request	13-11
13.5.2	Initialization Step Errors.....	13-13
13.5.3	User Quota Validation Errors	13-14
13.5.4	Placement Logic Failures.....	13-14
13.5.5	Troubleshooting Failures in Other Steps.....	13-16

Part V Middleware as a Service

14 About Middleware as a Service (MWaaS)

14.1	About Middleware as a Service	14-1
------	-------------------------------------	------

15 Setting Up Middleware as a Service

15.1	Getting Started.....	15-1
15.2	Defining Roles and Assigning Users.....	15-3
15.2.1	Creating a Custom Role for Self Service Application Users.....	15-3
15.2.2	Creating a User and Assigning Roles	15-4
15.3	Adding Hosts.....	15-4
15.4	Configuring Privilege Delegation Settings.....	15-5
15.5	Configuring the Software Library	15-6
15.6	Loading Provisioning Profiles to the Software Library.....	15-7
15.7	Creating a Middleware Home	15-8
15.8	Configuring and Saving the Deployment Procedure	15-9

16 Setting Up the Middleware Self Service Portal

16.1	Setting Up the Self Service Portal	16-1
16.1.1	Sample Scenario	16-5
16.2	Viewing the Middleware Pool Home Page.....	16-5

17 Using the MWaaS Self Service Portal

17.1	Using the MWaaS Self Service Portal.....	17-1
17.2	Requesting a Service.....	17-3
17.3	Viewing the Middleware Service Home Page.....	17-4
17.3.1	Deploying and Redeploying an Application.....	17-5

17.3.1.1	Redeploying an Application	17-6
17.3.1.2	Deleting an Application.....	17-6
17.3.2	Scaling Up and Scaling Down a Service.....	17-6
17.3.3	Creating a Data Source.....	17-6
17.3.3.1	Editing a Data Source.....	17-7
17.4	Uploading a Java EE Component to the Software Library	17-7
17.5	Viewing the Application Home Page	17-8

Part VI Cloud APIs

18 Introduction to Cloud APIs

18.1	Introduction	18-1
18.2	When to Use Cloud APIs	18-2
18.3	Common Behaviors	18-2
18.3.1	Transport Protocol.....	18-2
18.3.2	URI Space	18-2
18.3.3	Media Types	18-2
18.3.4	Request Headers	18-3
18.3.5	Response Headers.....	18-3
18.3.6	HTTP Status Codes.....	18-4
18.3.7	Common Resource Attributes.....	18-5
18.3.7.1	Resource State	18-5
18.3.8	Collection	18-6
18.3.9	Error Response Message	18-6

19 Cloud Resource Models

19.1	Enterprise Manager Cloud Resource Model 10001.....	19-1
19.2	Introduction	19-2
19.3	Resource Data Models.....	19-3
19.3.1	Cloud	19-3
19.3.2	ServiceTemplate.....	19-4
19.3.3	Zone	19-5
19.3.4	Service Family Type	19-6
19.3.5	Service Instance Type	19-7
19.3.6	Metric.....	19-7
19.3.7	Service Instance	19-8
19.4	Resource Models for Infrastructure as a Service	19-9
19.4.1	IaaS Zone [application/oracle.com.cloud.iaas.Zone].....	19-9
19.4.2	AssemblyInstance	19-10
19.4.3	Scalability Group	19-11
19.4.4	VM.....	19-12
19.4.5	VNet.....	19-13
19.4.6	NetworkInterface.....	19-14
19.4.7	VMTemplate.....	19-14
19.4.8	AssemblyTemplate	19-15

20 Cloud API Examples

20.1	Introspecting Resources	20-1
20.1.1	Cloud Resource	20-1
20.1.2	Service Family Type Resource	20-5
20.1.3	Service Instance Type resource	20-8
20.1.4	Zone	20-9
20.1.5	Service Template Resource	20-12
20.1.6	Metric resource	20-15
20.2	Creating a Service Instance	20-19
20.2.1	VM Creation	20-19
20.2.1.1	POST to the "iaas" Zone	20-19
20.2.1.2	POST to the VMTemplate	20-20
20.2.1.3	Polling the VM Creation	20-22
20.2.2	Assembly Instance Creation	20-24
20.3	Updating a Service Instance	20-39
20.3.1	Updating a VM Resource	20-39
20.3.1.1	Change VM Status	20-39
20.3.1.2	Changing VM Sizes	20-43
20.4	Deleting a Service Instance	20-44
20.5	Listing Service Instances	20-45
20.6	Version 10000 Support	20-46

21 Database as a Service APIs

21.1	Operations on ServiceTemplate Resources	21-1
21.1.1	Retrieving Service Template Details	21-1
21.1.2	Deleting a Service Template	21-2
21.2	Operations on DbPlatformInstance Resources	21-3
21.2.1	Creating a New DbPlatformInstance	21-3
21.2.2	Retrieving all the DbPlatformInstances in a Virtual Data Center	21-4
21.2.3	Deleting a DbPlatformInstance	21-4
21.3	Operations on DbPlatformTemplate Resources	21-5
21.3.1	Retrieving all the DbPlatformTemplates in the Cloud	21-5
21.3.2	Retrieving Information about a DbPlatformTemplate	21-6

22 Chargeback and Metering EMCLI Verbs

22.1	When to Use Chargeback APIs	22-1
22.2	Retreiving Metering Data from Enterprise Manager	22-1
22.2.1	get_metering_data Output	22-3
22.2.1.1	Examples	22-4

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 cloud 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 (12.1.0.4)
- Enterprise Manager for Chargeback and Capacity Planning (12.1.0.3)
- Enterprise Manager for Oracle Virtualization (12.1.0.3)

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*

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.
monospace	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

Part I

Cloud Management with Enterprise Manager

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 1, "Cloud Management - An Overview"](#)
- [Chapter 2, "Getting Started with Cloud Management"](#)
- [Chapter 3, "Using Consolidation Planner"](#)
- [Chapter 4, "Chargeback Administration"](#)

Cloud Management - An Overview

This chapter introduces Cloud Computing and the various service offerings and components available in the Oracle Cloud platform. It also describes the Consolidation Planner, various life cycle management capabilities of Oracle Enterprise Manager including resource management, metering, and Chargeback. It contains the following sections:

- [Introduction to Cloud Management](#)
- [Oracle Cloud Solution](#)
- [Cloud Management Lifecycle](#)
- [Cloud Resource Models](#)
- [Users and Roles](#)
- [Enterprise Manager Console and the Self Service Portal](#)
- [Cloud APIs](#)

1.1 Introduction to Cloud Management

Enterprises and Cloud Service Providers can use Oracle Enterprise Manager to build and operate their Cloud services. The Enterprise Manager sophisticated features allow administrators to manage the virtual and physical resources in their data center efficiently. The functionality provided by Enterprise Manager spans the entire Cloud lifecycle and allows you to setup and manage any type of Cloud service.

You can create and manage the Oracle VM based virtualized environment. You can define policies for allocating and balancing resources, power optimization, deployment, and management of the Cloud services.

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. These features are explained in detail in the following sections.

1.2 Oracle Cloud Solution

Enterprises need to support hundreds or even thousands of applications to meet growing business demands and 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; servers, storage and databases deployed as Infrastructure-as-a-Service or Database-as-a-Service or Platform-as-a-Service on Cloud offer compelling advantages in cost, quality of service, and agility.

- **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.

The supported Cloud platform services are:

- **Infrastructure as a Service (IaaS) Service Model:** The Oracle Cloud Management solution includes IaaS. Administrators can build this service entirely on Oracle hardware and software components or use a combination of different vendor hardware with Oracle Software. Oracle IaaS offering consists of Oracle Solaris, Oracle Enterprise Linux and Oracle VM for virtualization, Sun SPARC and x86 servers, and Sun storage.

This 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 to be deployed on the requested Guest VMs. Using pre-packaged assemblies of OS, DB 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. Users can also run chargeback reports to review the usage and chargeback on the resources consumed.

- **Database as a Service (DBaaS) Service Model:** DBaaS extends the Oracle Cloud management solution by automating the lifecycle of a database and allowing users to request database services through the Self Service Portal. With this solution, IT Managers no longer have to perform mundane administrative tasks for provisioning databases. Database users can get instantaneous access to new database services through the Self Service Portal. DBaaS is implemented through several options:
 - **Virtual Machine Based:** The database is deployed as a part of Virtual Assembly or Template and several virtual machines share the same physical server. This offers the maximum level of isolation (at OS level).
 - **Shared Cluster:** The database is deployed on existing clusterware. Typically the Grid infrastructure (Clusterware, ASM) and Database software is pre-installed. The cloud service entails the deployment of databases on top of that infrastructure.

- Shared Installation: The database is deployed as a single instance database on an existing installation.
- Shared Database: 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 uses a different service while accessing the database.

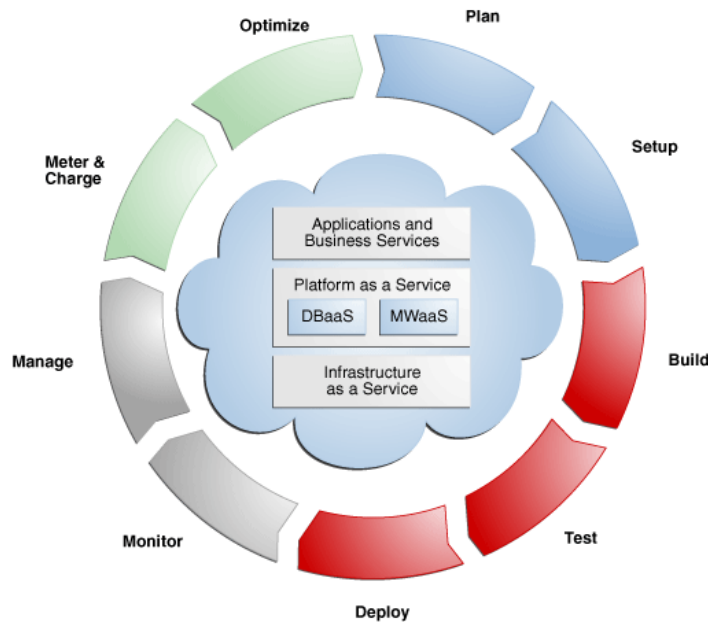
Like in IaaS, the users are allowed a few administrative tasks such as start up, shutdown, backup, and recover databases. Chargeback reports are also made available for Self Service users.

- **Middleware as a Service (MWaaS):** The Oracle MWaaS Platform is built on top of an Oracle IaaS offering. The MWaaS solution includes technologies like Oracle Database with Real Application Clusters and Oracle application grid including WebLogic Server, Coherence in-memory data grid and the JRockit JVM. On top of this foundation of clustered middleware and database technologies, the Oracle MWaaS Platform also includes components such as Oracle SOA Suite, Oracle BPM Suite, Oracle Identity Management and Oracle WebCenter. Users can request deployment of the platform using Self Service and the platform is deployed as a part of Virtual Assembly or Template and several virtual machines share the same physical server. This offers the maximum level of isolation (at OS level).

1.3 Cloud Management Lifecycle

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

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

Figure 1–1 Cloud Lifecycle

1.3.1 Planning

Using Enterprise Manager, you can setup a Cloud with brand new hardware, new software, and even a new data center. It allows you to transform existing data centers into a Cloud environment. Before setting up a Cloud, you need to plan the infrastructure requirements such as the physical and virtual networks, storage arrays, applications and so on. The auto discovery feature in Enterprise Manager helps discover and baseline all physical and virtual IT assets in an enterprise.

The Enterprise Manager Consolidation Planner is a powerful tool that helps administrators plan the Cloud architecture. It allows administrators to identify source and destination targets and applicable technical and functional constraints such as where the application can reside, and so on. Administrators 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).

1.3.2 Setting Up

Enterprise Manager can be used to model infrastructure-as-a-service (IaaS), middleware-as-a-service (MWaaS), and database-as-a-service (DBaaS) clouds. It offers capabilities such as bare metal provisioning of Hypervisor, setting up of server, storage pools and grouping them 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 the Oracle Virtual Assembly Builder (OVAB) to help package a multi-tier platform into a single metadata driven Cloud service. 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 comes with a role-driven access control. Integration with LDAP allows Enterprise Manager to inherit enterprise roles. The resource limits are implemented with quotas that are tailored for the specific service type (IaaS, MWaaS, and DBaaS). This prevents unauthorized usage of a service while also preventing a few users from using majority of the resources in the Cloud.

1.3.3 Building

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.

1.3.4 Testing and Deploying

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, they 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, one can set a policy to elastically scale out a Web Server if the processor load on existing Web Servers exceeds a certain threshold value.

1.3.5 Monitoring and Managing

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 comes with an in-built Incident Management system that can manage by exceptions. Administrators can review, suppress, escalate and remediate the events as needed, and also integrate it with ticketing systems.

Enterprise Manager has the ability to define contractual Service Level Agreements (SLAs) that 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.

Enterprise Manager also provides user experience management and business transaction management. The Enterprise Manager Configuration Management capabilities are optimized for Cloud environments. It can monitor vast numbers of configurations continuously, discover changes, measure drifts, pin-point configuration errors, as well as offer insight into system topologies, within 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 and DBaaS Home pages in Enterprise Manager allow Cloud administrators to get a summary view of the requests, the general state of the service such as zones, pools, servers, and databases. The IaaS Home page also provides a peek into popular policies and summary of incidents.

1.3.6 Metering, Charging, and Optimization

The Metering and Chargeback features in Enterprise Manager support basic metrics like CPU, memory, and storage usage, and offer pricing models based on application usage, database usage, and Middleware-level metrics. Administrators can also extend the pricing models to account for fixed costs, configurations, administrative expenses, people costs, energy utilization or a combination of these.

These capabilities enable enterprises to account for actual usage versus representative usage. Cloud Management also entails an ongoing optimization of resources as well 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. The tuning capabilities in the operating system, database and middleware layers aid in continuous optimization and subsequent improvement.

1.4 Cloud Resource Models

This section describes the resource models followed by Oracle Cloud Services such as IaaS, MWaaS, and DBaaS.

- IaaS Cloud Anatomy
- DBaaS and MWaaS Cloud Anatomy

1.4.1 IaaS Cloud Anatomy

The IaaS Cloud Anatomy consists of the following:

- **Server:** A physical computing system. Typically, a server runs a hypervisor and is part of a server pool.

- **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 be just be a 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:** 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:** 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.
- **Zone:** A Cloud could consist of one or more zones. A zone is a logical grouping of resources (for example, servers and storage) to facilitate self-service provisioning and administrative purposes and it could typically consist of hundreds to thousands of servers. A cloud consumer (self-service user) could, for instance, specify that his request needs to be provisioned within a given zone.

Zones are non-overlapping - meaning that a resource can only belong to one zone. However, resources within a zone may be accessible from another zone. For example, a VM in Zone 1 could interact with a VM in another zone. A zone consists of a set of server pools or just a loose grouping of individual machines. 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.

- **Cloud:** 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. A very large company or a company with complex security requirements may have more than one cloud.

1.4.2 DBaaS and MWaaS Cloud Anatomy

The DBaaS and MWaaS Cloud Anatomy consist of the following:

- **PaaS Infrastructure Zone:** 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 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:** For DBaaS, you need to set up Database Pools. A Database Pool is a collection of database homes. For MWaaS, you need to set up Middleware Pools. A Middleware Pool is a collection of middleware homes. A Database or Middleware 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.

1.5 Users and Roles

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. To set up a Private Cloud in Enterprise Manager, you need to define the following roles and create users for each of these roles:

- **Cloud Self Service User:** This user takes advantage of the cloud platform service to request resources, deploy applications, use them for a specific period (which could be very short or very long), manage, and monitor some aspects of the lifecycle of the application. For example, this user could provision the application, patch it, monitor the performance of the application deployed in the cloud, and finally create backups of specific artifacts from the Cloud. In production deployments, multiple people may perform these functions. For production applications deployed in the Cloud, these functions are fulfilled by “professional IT staff” depending upon the application. This user is interested in viewing periodic reports about usage and costs related to chargebacks.
- **Self-Service Administrator:** This user is responsible for setting up the service catalog with various services such as IaaS, DBaaS, or MWaaS. This user can define allowable VM sizes, define deployment procedures for database or J2EE Application provisioning, assign quotas to users and roles, define access boundaries, setup chargeback plans and maintenance levels.

The Self Service Administrator also requests the Cloud Administrator for additional memory, storage resources, and so on. This user could represent a department or an organization, as typically defined in an LDAP directory, to which a zone has been allocated. This user is interested in viewing periodic reports about usage, costs related to chargebacks, and outage reports with appropriate metrics. This user is also interested in negotiating with the Cloud Infrastructure Administrator on planned outages and other changes.

- **Cloud Administrator:** This user is responsible for putting together the overall cloud infrastructure, the physical and logical cloud partitioning. This user is also responsible for the provisioning of physical resources like bare metal hardware and hypervisor, configuring storage arrays and VLAN, setting up monitoring, configuration, administration, and de-provisioning of the physical infrastructure. This user typically is responsible for standardization of what infrastructure to procure that would be used in the cloud. This user specifies what core resource metrics are measured and how they are used to chargeback the different cloud users.

1.6 Enterprise Manager Console and the Self Service Portal

Typically, Cloud Administrators and Self Service Administrators use the Enterprise Manager Cloud Control Console to setup, monitor, and manage Cloud Services. Each service is managed using an exclusive page for that service. For example, IaaS, DBaaS and MWaaS have their own interfaces that can be accessed directly from the Cloud Summary page or from the Enterprise Manager menu. The Enterprise Manager Cloud Summary console 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).

The Self Service Portal is the Home Page for the Self Service User. It is a customizable page that enterprises can utilize to launch their cloud services. It allows the Enterprise name to be on this portal. Each service is separated with a tab and users who have the necessary privileges can navigate between Services pages. 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.

1.7 Cloud APIs

Enterprise Manager provides a Web Services based RESTful (Representational State Transfer) Cloud API to allow enterprises to customize the Cloud according to their business processes and economic models. Customers with unique requirements for self service can create those services without administrator involvement. They can browse templates containing metadata of a service, deploy them into a cloud, and perform operations such as backup, enable or disable templates, and so on.

Getting Started with Cloud Management

This chapter describes how to start using the Enterprise Manager Cloud Management solution. It includes the following sections:

- [Installing Virtualization and the Cloud Plug-ins](#)
- [Deploying the Cloud Plug-ins](#)
- [Setting Up the Software Library](#)
- [Using Self Update to Download Oracle VM Templates and Oracle Virtual Assemblies](#)
- [Managing Users and Roles](#)
- [LDAP Authentication](#)
- [Customizing the Self Service Login Page](#)

2.1 Installing Virtualization and the Cloud Plug-ins

A plug-in is an Enterprise Manager Server component (module) that can be plugged into an existing Enterprise Manager Server to extend its managing and monitoring capabilities. The Enterprise Manager Server is the component of Enterprise Manager where the centralized part of the Enterprise Manager application logic executes. This includes the Oracle Management Service (OMS) and the Management Repository. On both the OMS and Management Repository sides, Enterprise Manager Server may be considered as a combination of the platform and a set of plug-ins installed on top of the platform. To extend its capabilities of monitoring and managing, plug-ins are used on both the Management Server and the Management Agent.

Before you start monitoring virtualization targets with Enterprise Manager, the cloud plug-ins must be installed. You can install the cloud plug-ins in two ways:

- While installing a new Enterprise Manager system, you can select the plug-ins that need to be installed from the software kit. The plug-ins required to monitor virtualization targets are:
 - Enterprise Manager for Oracle Virtualization
 - Enterprise Manager for Oracle Cloud Application
 - Enterprise Manager for Chargeback and Capacity Planning

These selected plug-ins are installed as part of the Enterprise Manager installation. For more details, see the Enterprise Manager Basic Installation and Configuration Guide.

- If these plug-ins are not selected during Enterprise Manager installation, you can deploy the necessary plug-ins to the Management Service and the Management Agent. See [Section 2.2, "Deploying the Cloud Plug-ins"](#) for details.

2.2 Deploying the Cloud Plug-ins

Plug-ins must first be downloaded to the OMS using Self Update before they can be deployed. To enable monitoring of virtualization targets, you must download the plug-ins and deploy them. The following plug-ins must be deployed on the Management Server:

- Enterprise Manager for Oracle Virtualization
- Enterprise Manager for Oracle Cloud Application
- Enterprise Manager for Chargeback and Capacity Planning

Note:

- After these plug-ins have been deployed to the Management Server, the Enterprise Manager for Oracle Virtualization plug-in must be downloaded to the Management Agent.
 - If you have installed the Enterprise Manager for Oracle Virtualization 12.1.0.2 and the Enterprise Manager for Oracle Cloud Application 12.1.0.2 plug-ins, and need to do an upgrade, you must install the 12.1.0.3 version of both the plug-ins. You cannot upgrade one plug-in without upgrading the other.
-
-

2.2.1 Deploying the Cloud Plug-ins to the Management Server

To deploy the cloud plug-ins to the Management Server, follow these steps:

1. Log in to Enterprise Manager as a user with the `EM_CLOUD_ADMINISTRATOR` role.
2. Click **Deploy** and choose Oracle Management Server from the **Deploy** list.

Be sure that dependent plug-ins are deployed and that all existing Management Agents are compatible with the version of the specified plug-in. Enterprise Manager prompts for credentials if the Management Agent is not available.

Note: The plug-ins must be deployed on the Management Server (OMS) before they are deployed on the Management Agents. In a multi OMS setup, the plug-ins must be deployed on all the OMSes.

3. Specify the required details on the Deploy Plug-in dialog box.
4. In the Version of Plug-in to Deploy section, select the Plug-in version from the Plug-in list. The Target Type information appears in the table. Enter the Repository sys Password, then click **Continue**.
5. Proceed through the steps in the Deploy Plug-in dialog box.
6. Click **Deploy** to deploy the selected plug-in on all Enterprise Manager servers. Enterprise Manager displays a page that allows you to monitor the deployment status. Enterprise Manager deploys the selected plug-in on all Enterprise Manager Servers.

2.2.2 Deploying the Enterprise Manager for Oracle Virtualization Plug-in to the Management Agent

To deploy the Enterprise Manager for Oracle Virtualization plug-in to the Management Agent, 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**, then select **Plug-ins** to navigate to the Plug-ins page.

Figure 2–1 Plug-ins Page

The screenshot shows the 'Plug-ins' page in Oracle Enterprise Manager. At the top, it says 'Page Refreshed Jul 24, 2012 5:38:13 PM UTC'. Below that, a table lists various plug-ins. The 'Oracle Virtualization' plug-in is selected and highlighted in blue. Below the table, there is a section for 'Oracle Virtualization' with details:

Name	Latest Available	Latest Downloaded	On Management Server	Management Agent with Plug-in	Description
EMC CLARION Array	12.1.0.2.0	12.1.0.2.0	12.1.0.2.0	0	EMC CLARION Array Monitoring .
EMC Celerra Storage	12.1.0.2.0	12.1.0.2.0	12.1.0.2.0	0	EMC Celerra monitoring including reports
EMC Symmetrix Array	12.1.0.2.0	12.1.0.2.0	12.1.0.2.0	0	EMC Symmetrix Array Monitoring .
Oracle Cloud Application	12.1.0.4.0	12.1.0.4.0	12.1.0.4.0	0	Plug-in to enable cloud self service framework and portal in Oracle Enterprise Manager.
Oracle Ops Center Infrastructure	12.1.0.1.0			0	Ops Center Infrastructure stack management plug-in.
Oracle Virtualization	12.1.0.3.0	12.1.0.3.0	12.1.0.3.0	1	Plug-in to enable Oracle VM virtualization management capabilities in Oracle Enterprise Manager.
Sun ZFS Storage Appliances	12.1.0.1.0			0	Management Plug-in for the Sun ZFS Storage Appliance Family
Exalogic Elastic Cloud Infrastructure	12.1.0.1.0	12.1.0.1.0	12.1.0.1.0	0	Elastic Cloud Infrastructure

Below the table, the 'Oracle Virtualization' section shows:

- Recent Deployment Activities
- General
- Plug-in ID: oracle.sysman.vt
- Vendor: oracle
- Versions Downloaded: 12.1.0.3.0, 12.1.0.2.0, 12.1.0.1.0
- Supported versions on Management Agent: All older released versions
- Version on Management Server: 12.1.0.3.0
- Latest Available Version: 12.1.0.3.0
- Description: Plug-in to enable Oracle VM virtualization management capabilities in Oracle Enterprise Manager.

3. Expand the Servers, Storage, and Network folder and select the Enterprise Manager for Oracle Virtualization plug-in.
4. Click **Deploy On** and click **Management Agent**. The Deploy Plug-in On Management Agent: General dialog box appears.
5. Click **Add** in the Selected Management Agent section to add the agent on which the plug-in is to be deployed. When you click **Add**, Enterprise Manager displays the Search and Select dialog box where you can select the agents to add. Click **Select** to return to the Deploy Plug-in dialog box. Only agents running the operating systems supported by the selected plug-in can be selected.
6. Click **Continue**. The Deploy Plug-in On Management Agent: Pre-requisite Checks dialog box appears. Click **Next** to view the Review page.
7. Click **Deploy** to deploy the plug-in on the selected Management Agents. Click **Show Status**. Enterprise Manager displays a page that monitors the deployment status and begins the deployment process with the Install option. Deployment occurs in parallel on all selected agents.

2.3 Setting Up the Software Library

Oracle Software Library (Software Library) is a repository that stores software patches, virtual appliance images, reference gold images, application software and their

associated directive scripts. It allows maintaining versions, maturity levels, and states of entities. The software entities can be automatically mass-deployed to provision software, software updates, and servers using Enterprise Manager Cloud Control in a reliable and repeatable manner. These provisioning operations, which are unattended and can be scheduled, lead to substantial cost savings.

Besides acting as a repository for certified software entities, the Software Library is a logical interface between the deployment models and the automation framework required to perform a large number of patching and provisioning tasks. You can pick any of the entities supplied by Oracle, 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.

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

Figure 2–2 Software Library Page

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 **Actions** menu, click **Administration**. The Software Library: Administration page appears where you can select the storage location. 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:** 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.

- **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 be HTTP, NFS, and Agent. For more details on configuring the software library storage, see the *Configuring Software Library* section in the Enterprise Manager Administration Guide.

After the Software Library storage has been configured, you can define virtualization entities such as assemblies, templates, ISO images, and virtual disks. See [Section 7.9.3.1, "Creating the Software Library Components"](#) for details. For more details, see the *Configuring Software Library* chapter in the Enterprise Manager Administration Guide.

2.4 Using Self Update to Download Oracle VM Templates and Oracle Virtual Assemblies

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 Cloud Control to obtain information about available updates.

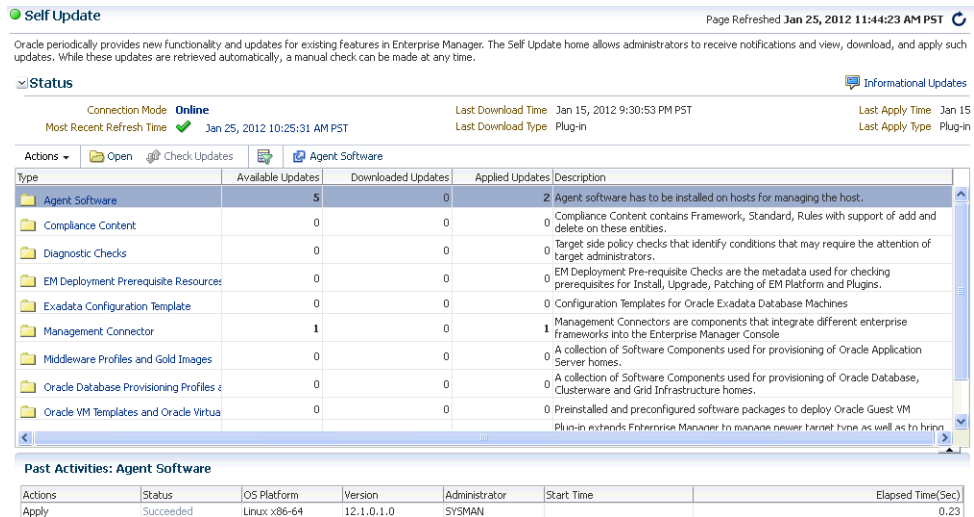
Before you use the Self Update feature, the following prerequisites must be met:

- My Oracle Support credentials must be configured. This allows you to download entities from the My Oracle Support site.
- The Software Library (also known as the local store) must be configured. Updates are downloaded to this local store before being deployed into Cloud Control.

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 2-3 Self Update Console



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 that 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.

2.5 Managing Users and Roles

System security is a major concern of any corporation. Giving the same level of access to all systems to all administrators is dangerous, but individually granting access to tens, hundreds, or even thousands of targets to every new member of the group is time consuming. With the Enterprise Manager administrator privileges and roles feature, this task can be performed within seconds, instead of hours.

One Super Administrator, *SYSMAN*, is created by default during the installation of Oracle Enterprise Manager. The *SYSMAN* Super Administrator then creates other administrator accounts for daily administration work. The *SYSMAN* account should only be used to perform infrequent systemwide, global configuration tasks.

The Super Administrator divides workload among his administrators by filtering target access, or filtering access to management task, or both through the roles, other resource privileges, and Target Privileges he grants them. For example, he can allow some administrators to view any target and to add any target in the enterprise and other administrators to only perform specific operations such as maintaining and cloning on a target for which they are responsible.

To set up a Private Cloud in Enterprise Manager, you need to use the following roles:

- **EM_INFRASTRUCTURE_ADMINISTRATOR:** Users with this role have privileges to manage the Enterprise Manager infrastructure such as managing the plug-in lifecycle, self update feature, and so on.
- **EM_CLOUD_ADMINISTRATOR:** Users with this role are responsible for setting up and managing the cloud infrastructure. This role is responsible for deploying the cloud infrastructure (servers, pools, zones) and infrastructure cloud operations for performance and configuration management. The **EM_CLOUD_ADMINISTRATOR** must have the following privileges:
 - View privileges on all the cloud targets.
 - Operator privileges on the targets (zones, pools) on which the assemblies and templates are being deployed.

These privileges are propagated to the Oracle VM Manager targets and allow the **EM_CLOUD_ADMINISTRATOR** to deploy assemblies and templates.

- **EM_SSA_ADMINISTRATOR:** Users have the privilege to define the self service role. They can define quotas and constraints for the self service users and grant them access privileges. Users with this role have provisioning designer privileges, can create and view patch plans, and support the plug-in lifecycle on the Management Agent.
- **EM_SSA_USER:** Users with this role can access the Self Service Portal.

2.5.1 Creating a Role

In Enterprise Manager, Super Administrators can create, grant, and revoke roles. To create a **CLOUD_ADMINISTRATOR**, 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 for the role and click **Next**.
5. From the list of Available Roles, select the **EM_CLOUD_ADMINISTRATOR** role and move 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 **CLOUD_ADMINISTRATOR** role.

Note: Follow these steps to create roles but ensure that you select the relevant roles during the creation process.

2.5.2 Creating a User and Assigning Roles

To create a user called `cloud_admin` and grant the **CLOUD_ADMINISTRATOR** role privileges, follow these steps:

1. Log into 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 role and click **Next**.
5. From the list of Available Roles, select the EM_CLOUD_ADMINISTRATOR 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. Review all the changes and click **Finish** to create the cloud_admin user.

Note: Follow these steps to create other users. For SSA users, the EM_USER and PUBLIC roles must be removed.

2.6 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, see the *Enterprise User Security* chapter in the Enterprise Manager Administration Guide.

2.7 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](#)

2.7.1 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 properties. If these properties are not set, the default Enterprise Manager 12c logo and copyright message is displayed.

```
$ORACLE_HOME/bin/emctl set property -name oracle.sysman.ssa.logon.show_cloud_provider_brand -value true
$ORACLE_HOME/bin/emctl set property -name oracle.sysman.ssa.logon.show_cloud_tenant_brand -value true
$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"
$ORACLE_HOME/bin/emctl set property -name oracle.sysman.ssa.logon.show_disclaimer_text -value true
$ORACLE_HOME/bin/emctl set property -name oracle.sysman.ssa.logon.disclaimer_text -value "Customer Specified Disclaimer Text"
$ORACLE_HOME/bin/emctl set property -name oracle.sysman.ssa.logon.show_em_branding_text -value true
```

- 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.

2.7.2 Switching Back to the Enterprise Manager Login Page

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

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

2.7.3 Routing SSA Requests to a Specific OMS Pool

If several SSA users are connecting to the OMS, you can configure the Server Load Balancer (SLB) to redirect all SSA requests to a specific OMS pool. This makes the other OMS pool available for administration usage. To redirect SSA requests, you must set 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.

Using Consolidation Planner

This chapter covers usage 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](#)

3.1 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 SPECint[®]_rate_base2006 benchmark, published by Standard Performance Evaluation Corporation (SPEC), for the different classes of hardware.

The Consolidation Planner feature enables you to match managed servers you want to consolidate with the generic physical machines, Oracle Exadata database machines, 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.

3.1.1 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 or an integrated system such as the Oracle Exadata Database Machine.

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

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(s) 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. You can change the values in a scenario, then re-run it to view the new results.

3.2 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](#)

3.2.1 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
- Nodes of an Oracle Cluster

3.2.2 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.

3.3 Using Consolidation Planner

The steps in the consolidation planning process are:

1. Create a consolidation project. See [Section 3.3.1, "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 3.3.2, "Using a Pre-configured Consolidation Scenario"](#).
 - Create a custom consolidation scenario. See [Section 3.3.3, "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 3.3.4, "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 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.

This section includes the following:

- [Creating a Consolidation Project](#)
- [Using a Pre-configured Consolidation Scenario](#)
- [Creating a Custom Consolidation Scenario](#)
- [Evaluating Consolidation Scenarios](#)

3.3.1 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 **Create Project**.
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).

Note that for both types, only source servers running on an operating system listed in the pull down menu can be consolidated.

5. 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.
6. 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** to view a list of existing Oracle Virtual Machine destination servers to consolidate the source servers to. Select the servers you want to add, then click **Select**.
If you are consolidating to an Oracle Virtual Machine server pool, the first virtual server pool is selected by default to provide target server candidates.
 - If you are consolidating from physical servers to physical servers (P2P), click **Add Existing Database Machines** to view a list of Oracle Exadata Database Machine servers to consolidate the source servers to. Select the servers you want to add, then click **Select**.
7. 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 the 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 the collect data. The default maximum value is 90 days.
- Specify when to begin the data collection process.
- Optionally select **Continue Data Collection Over the Maximum Days** to purge the oldest day's data when data for a new day is added.

8. 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.

9. 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.

3.3.2 Using a Pre-configured Consolidation Scenario

When creating a consolidation project, you can optionally choose to generate 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 the scenario(s) 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 ration 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. 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. For CPU capacity, provide the estimated SPEC[®] rate if available; otherwise, the Consolidation Planner will determine requirements based on the CPU configurations provided earlier in the wizard.
The Consolidation Planner will determine how many destination servers are required as part of the consolidation scenario generation.
 4. Click **Use Existing Servers** to use the destination candidates specified in the Consolidation Scope section of the consolidation project.
Note that you can edit the CPU (SPEC[®] rate) and disk I/O fields.
 5. Click **OK** when finished.

3.3.3 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. The 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.
 - **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 ration 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 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, database nodes in an Oracle RAC cluster should not be consolidated on the same destination server. 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. Click **Use New 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.

If you are defining the scenario for a P2V project, fill in the data for each server. For CPU capacity, provide the estimated SPEC[®] rate if available; otherwise, click the **List of Values** icon next to the CPU capacity input field, then select a server configuration that most closely matches your needs.

If you are defining the scenario for a P2P project, in addition to the same new server option as a P2V project, you can also select a pre-configured Oracle Exadata database machine as the planned target server. Click the **List of Values** icon next to the read-only Configuration input field, then select one of the pre-configured Oracle Exadata database machines.

The Consolidation Planner will determine how many target servers are required as part of the consolidation results.

13. Click **Use Existing Servers** to specify a set of existing managed servers to use as targets.

These are the servers you specified when defining the scope for the consolidation project. The Consolidation Planner will determine the available hardware resources based on collected usage data.

By default, the consolidation process will try to use as few target servers as possible. If you prefer, you can choose to balance the source load across all destinations.

14. Enter percentages for **Maximum Allowed Resource Utilization on Destination Servers**. The defaults are 75 for CPU and 90 for memory, respectively.

15. Click **Next** to map the source servers to the destination servers.

Next, 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:

16. Click a source server in the list.
17. 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.
18. Click **Next** to review the consolidation scenario.
19. Click **Submit** to save the new scenario.

3.3.4 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 re-run existing scenarios with new inputs, then compare the results again. 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
- 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. Click on any metric or on a tab at the bottom of the page to view additional details. Key data includes:
 - **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, the Consolidation Planner will try to “fit” source servers into as few destination servers as possible.
- **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

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](#)

4.1 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](#)

4.1.1 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.

4.1.2 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.

4.1.3 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, and database. 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.

4.1.4 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.

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 Enterprise Manager for Oracle Cloud plug-in)
- Oracle VM

Within Chargeback, the respective zone hierarchies appear as follows on the **Targets** tab, under the Servers, Storage and Network category:

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 Single Instance Database Cluster Database Database Instance1 Database Instance2 Oracle Composite Application WebLogic Domain WebLogic Cluster Oracle WebLogic Server1 Oracle WebLogic Server2 Oracle WebLogic Server
Oracle Exalogic	Oracle Exalogic Zone Oracle Assembly Instance Oracle Tier Instance Oracle VM Guest1 Oracle VM Guest2
Oracle VM	Oracle VM Zone Oracle VM Server Pool Oracle VM Guest1 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 4.2.4](#) 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 4.2.5](#) to learn about assigning charge plans.

4.2 Setting Up Chargeback

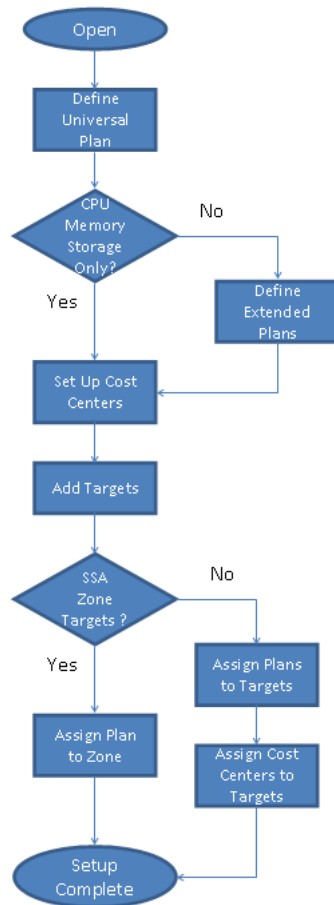
This section describes the following tasks, which must be completed to implement and support Chargeback services:

- [Defining 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 4–1](#) shows a simple workflow of steps required to set up Chargeback.

Figure 4–1 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.

4.2.1 Defining 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 plan uses three basic metrics against which to compute resource consumption: CPU, memory, and storage.

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 usage 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 a universal charge plan for general use. You can then indicate in an extended charge plan how the universal plan rates apply.

This section covers the following topics:

- [Defining the Universal Charge Plan](#)
- [Defining an Extended Charge Plan](#)

4.2.1.1 Defining the Universal 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 day per GB of memory, and a dollar a day 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.

1. From the **Enterprise** menu, select **Chargeback**.
2. Select the **Charge Plans** tab.
3. Set the rates and charge frequency for the three basic metrics.

- a. Select the default CPU row and click **Edit**.

In the dialog that opens, set the rate and charge frequency and then click **OK**.

- b. Set rates and charge frequency for memory and storage allocation in similar fashion.

- c. Make additional entries for different CPU architectures, as required. Click the **Add** button.

In the dialog that opens, select the architecture from the drop-down list, set the rate and charge frequency, and then click **OK**.

The rates set here are effective from the beginning of the current month, and have no impact on any previous month.

You can also initiate this action (manage universal plan) from the **Home** tab.

Updating the Universal Charge Plan

Periodically, it may be desirable to adjust the universal plan rates.

1. From the **Enterprise** menu, select **Chargeback**.
2. Select the **Charge Plans** tab
3. Change charge rates and frequencies as appropriate. Select the table row of the base metric you want to change and click **Edit**. Make the changes and click **OK**.

You can view rates set for previous months, but can only modify rates when the current plan version (date range) is selected.

4.2.1.2 Defining an Extended Charge Plan

An extended charge plan consists of groups of target-type specific charge items. Target-specific charge templates are available for host, virtual machine, database and WebLogic Server target types.

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 a \$100 a month for each host, regardless of consumption.

Conditions introduce considerable flexibility when creating extended charge plans. With conditions, you set charges based on properties. 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.

Similarly, you can place a condition on a 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 \\%).

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.

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.

An extended charge plan allows you to customize how charges accrue.

1. From the **Enterprise** menu, select **Chargeback**.
2. Select the **Charge Plans** tab.
3. In the **Extended Charge Plan** region, click **Create**, type a name for the new charge plan, and then click **OK**.

4. On the Edit Extended Charge Plan page, click **Add**. Complete the dialog that opens as follows:
 - Select a target type from the drop-down list.
 - For Database **Instance** or **Oracle WebLogic Server**, indicate whether it is shared.
 - Select whether to apply a condition. Default means match on target type unconditionally

The target type selected drives the list of conditions available. 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. For the sample extended plan, for example, three conditions are set for target type VM Guest, where the specified value for machine size can be small, medium, or large. To compute charges, Chargeback tests the conditions against the target in question and calculates accordingly. You might, for example, set a condition of high availability equals yes for a VM Guest target type.

Note that you cannot set conditions for a shared WebLogic Server.

Figure 4–2 shows a condition set on a VM Guest target type for high availability.

Figure 4–2 *Setting a Target Type Condition*

The screenshot shows a dialog box titled "Add Target Type Specific Charges". It has a close button in the top right corner. The "Target Type" is set to "Oracle VM Guest" in a dropdown menu. Below this, there are two radio buttons: "Default (Applies when no Condition is matched)" which is unselected, and "Use Condition" which is selected. Under "Use Condition", there are three fields: "Condition" set to "High Availability", "Condition Operator" set to "Equals", and "Condition Value" set to "Y". There is a search icon to the right of the "Condition Value" field. At the bottom of the dialog are "OK" and "Cancel" buttons.

5. Add charge items based on target type for billing purposes. In the **Charge Items** tab, click **Add**. Complete the dialog that opens as follows:
 - Select a charge item from the drop-down list. The list of items derives from the target type. The unit of measure to which the charge applies appears as well, where applicable (MB, GB, and so forth). For a shared database, charges are per service. For a shared WebLogic Server, charges are per application.
 - Select the charge type; that is, whether the charge is usage-based, config-based, or a flat rate. These choices vary, depending on the charge item selected.
 - Specify the rate at which to charge for the item.
 - Select from the drop-down list the frequency at which to apply the charges. It's not unusual to have different frequencies for charge items within the same target type, depending on whether the item is a config or metric, and the charge is a flat rate or usage-based.
 - Select whether to apply a condition; that is, something that must happen, a threshold reached for example, to generate charges. Default means apply unconditionally.

Figure 4-3 shows a condition set on a charge item within the VM Guest target type that sets an additional flat rate charge when the IP address type is static.

Figure 4-3 Setting a Condition on a Charge Item

Click **OK**.

Repeating the step to add an additional charge item of a base rate of a dollar a day, the extended charge plan might be summarized so far as follows: When the target type is a high availability VM Guest, charge a flat base rate of a dollar day per target instance. In addition, if the target has a static IP address, charge an additional flat rate of two dollars a day, as shown in Figure 4-4.

Figure 4-4 Extended Plan Showing Customized Charge Items

Item Name	Charge Type	Rate (\$)	Condition	Per Unit	Per Time Period
Base Charge	Flat	1		Target Instance	Day
IP Address Type	Flat	2	IP Address Type = 'Static'	Target Instance	Day

- Optional. Make rate adjustments to the 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.

For a shared database, only the CPU rate can be adjusted. No adjustments can be made for a shared WebLogic Server.

- Click **Save** (top right of the page) to create the new extended charge plan.

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.

You can also initiate this action (manage extended plans) from the **Home** tab.

Updating an Extended Charge Plan

Over time and given analysis, you will want to make changes to extended charge plans.

1. From the **Enterprise** menu, select **Chargeback**.
2. Select the **Charge Plans** tab.
3. In the extended charge plan region, select in the drop-down list the plan you want to update and click **Edit**.
4. On the Edit Extended Charge Plan page, make the necessary edits to target type-specific charges and charge items by selecting a row in the respective table and clicking the **Edit** or **Remove** button, as appropriate. Also make any desired adjustments to universal rates.

Note that you can only modify rates when the current plan version (date range) is selected, such that the date on which you make the changes falls between the start and end dates of the plan version.

5. Click **Save** to complete the plan update.

4.2.2 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](#)

4.2.2.1 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 from which it can subsequently be reassigned to another 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 from which it can subsequently be reassigned to another 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.

4.2.2.2 Creating the Cost Center in Chargeback

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.

4.2.2.3 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.

4.2.3 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.

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_TARGET` role allows you to see Chargeback data related to a specific target. The `VIEW_ANY_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.
3. In the **Add Targets** dialog, select the targets to manage, and then click **OK**. If there are no targets in the list, click the **Target Selector** button to search for targets to include.

Select the targets in the search results list and click **Select**.

4. In the **Add Targets** dialog, indicate whether the targets are to be shared across different cost centers. In particular, databases, RACs, and WebLogic Servers can be shared. The check box is disabled if a target is ineligible for sharing.

To share a database instance or a RAC, Enterprise Manager metric collection must be enabled. If you add a target before enabling metric collection, you will have to remove it and add it back in shared mode after enabling metric collection.

Click **OK**. Selected targets appear in the **Target Name** column, nested under the major target type categories (Databases, Middleware, and so forth). Selecting a zone automatically includes all targets within the zone.

You can also initiate this action (add targets) from the **Home** tab.

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.

A RAC can be shared by instance or by service. When shared by instance, each RAC member appears under the RAC node in the tree hierarchy in dedicated mode and so can be assigned to different charge plans and different cost centers. When shared by service, RAC services appear under the RAC node in the tree hierarchy. Services can be assigned to different cost centers, but must all be assigned to the same charge plan.

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.

Configuring Enterprise Manager Metric Collection for Shared Database Targets

For shared databases and RACs, metric data 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.

For a RAC, do this for each cluster member.

4.2.4 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.

1. From the **Enterprise** menu, select **Chargeback**.
2. Select the **Targets** tab. Targets without an explicitly assigned cost center should appear at the top. When charges are calculated, unassigned targets are assigned to the default cost center.
3. Select the target (or 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.

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.

4.2.5 Assigning Charge Plans to Targets

The universal charge plan is assigned automatically when a chargeable target instance (host, database, WebLogic Server, VM Guest) that has no parent is added for Chargeback. There is otherwise no automatic assignment so 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 this process. 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.

1. From the **Enterprise** menu, select **Chargeback**.
2. Select the **Targets** tab. A hierarchical list of targets appears showing assigned cost centers and extended charge plans.
3. Select the target or 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.

4.2.6 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 GB * \$2 * 5)
- A hundred dollars a day for storage (100 GB * \$1)
- A base charge of 24 dollars a day (24 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).

4.3 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](#)

4.3.1 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 4–5 displays a usage trend showing CPU, disk, and memory utilization over the past week, aggregated daily.

Figure 4–5 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 4–6 displays a charge trend showing last week’s charges for the VM Guest target type, aggregated daily.

Figure 4–6 Sample Charge Trend Graph


4.3.2 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 4–7 shows a chart tracking instance uptime for the selected target over a span of five days.

Figure 4–7 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.

4.3.3 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.

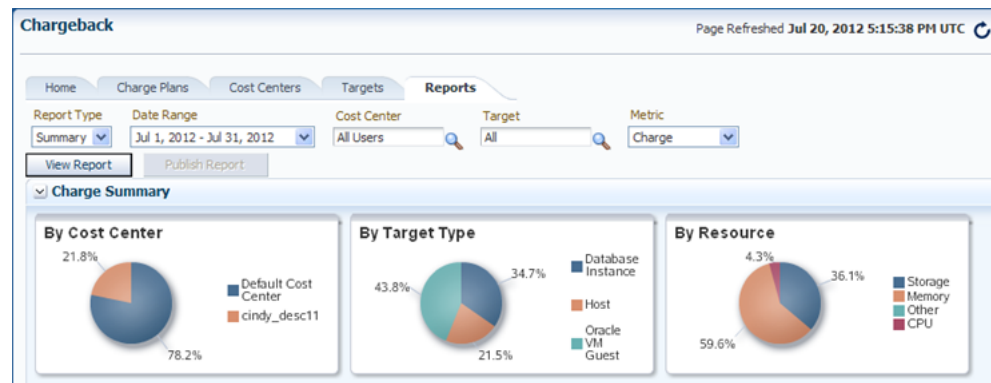
- Choose between summary and trend report types. A summary report presents a pie-chart breakdown, while a trend report uses a y-axis multiple bar chart to denote usage trends.
- Select specific cost centers or report on all users.
- Select specific targets or target types or report on all targets within 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 4–8](#) displays a summary report showing charges for the current reporting cycle for all cost centers and target types, with a breakdown by resource.

Figure 4–8 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 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.

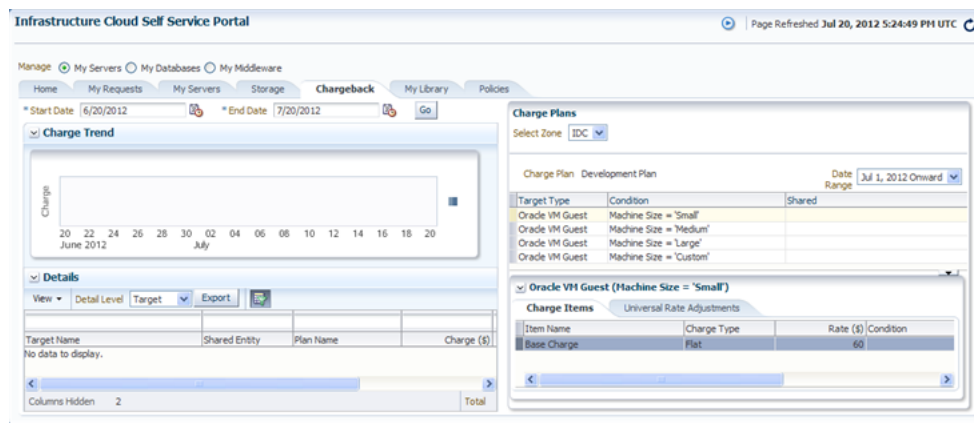
4.3.4 Viewing Chargeback Information in the Self Service Portal

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 4–9 shows a sample Chargeback informational breakdown in the self service portal.

Figure 4–9 Chargeback Breakdown for a Self Service Application



4.3.5 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 22, "Chargeback and Metering EMCLI Verbs"](#) for more information.

Part II

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 5, "Setting Up Cloud Infrastructure"](#)
- [Chapter 6, "Setting Up and Using the Self Service Application for IaaS"](#)
- [Chapter 7, "Administering and Monitoring the Cloud for IaaS"](#)
- [Chapter 8, "Managing Cloud Policies"](#)

Setting Up Cloud Infrastructure

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](#)
- [Registering the Oracle VM Manager](#)
- [Synchronizing the OVM Manager Targets \(Optional\)](#)
- [Discovering an Oracle VM Server](#)
- [Rediscovering a Virtual Server](#)
- [Setting Up Networks](#)
- [Creating a Network Profile](#)
- [Registering Storage Servers](#)
- [Creating a Storage Repository](#)
- [Performing Administrative Operations on a 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](#)

5.1 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 5.2, "Registering the Oracle VM Manager"](#) for details.
2. You can then discover one or more Oracle VM Server targets. See [Section 5.4, "Discovering an Oracle VM Server"](#) for details.
3. Set up the storage infrastructure for the cloud. See [Section 5.8, "Registering Storage Servers"](#) for details.

4. Configure networks and VLAN Groups. See [Section 5.6, "Setting Up Networks"](#) for details.
5. Create a virtual server pool. See [Section 5.12, "Creating a Virtual Server Pool"](#) for details.
6. Create a zone by grouping one or more server pools. See [Section 5.13, "Creating a Zone"](#) for details.

Figure 5–1 shows the sequence of steps involved in setting up the virtualization infrastructure.

Figure 5–1 Setting Up the Cloud Infrastructure



5.2 Registering the Oracle VM Manager

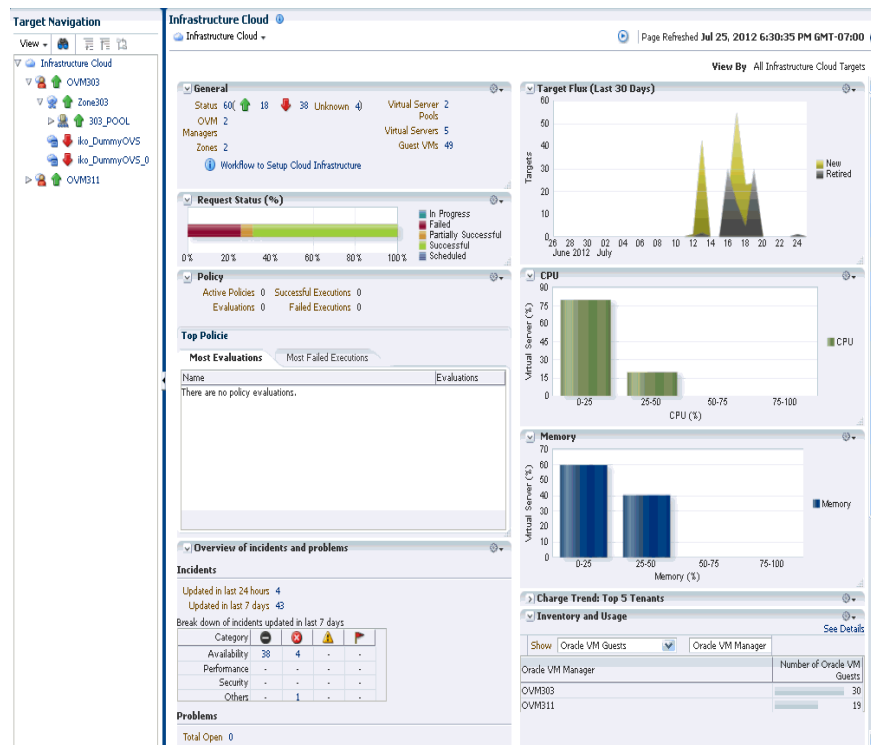
The Oracle VM Manager provides the user interface that allows you to create and manage zones, virtual server pools, virtual servers, and guest virtual machines. Before you use the Oracle VM Manager to manage virtualization targets in Enterprise Manager, you must register an existing installation of the 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 OVM Manager that you register with Enterprise Manager must be at least version 3.0.2 or higher.
 - The OVM 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 5–2 Infrastructure Cloud Home Page



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

Figure 5–3 Add Oracle VM Manager

Infrastructure Cloud

Infrastructure Cloud

Page Refreshed Sep 27, 2011 1:11:49 PM GMT-07:00

Register an existing installation of Oracle VM Manager with Enterprise Manager. Make sure that a connection can be established using the specified URL.

* Name

* EM Agent URL

Oracle VM Manager Core

* Oracle VM Manager URL

Automatic Synchronization

Monitoring Credentials for Oracle VM Manager

Specify credentials to be used for monitoring Oracle VM Manager.

* Username

* Password

Administration Credentials for Oracle VM Manager

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

Use Administration Credentials

* Username

* Password

Oracle VM Manager Console

Provide the URL of Oracle VM Manager console. This will be used to launch WVC console for Guest VM.

* Oracle VM Manager Console URL

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

1. Enter a name for the Oracle VM Manager target to be registered.
2. Select the URL for the Management Agent. For example, `https://agent.example.com:1830/emd/main/`

3. Enter the URL (`tcps://ovmmgr.example.com:54322/`) of the physical machine on which the Oracle VM Manager is installed. Refer to *Note 1456338.1* for details on configuring Oracle VM Manager for TCPS access.
4. To automatically synchronize the data from Oracle VM Manager into Enterprise Manager at fixed intervals of time, enable the **Automatic Synchronization** checkbox. The default interval is 5 minutes. This setting is optional.
5. Enter the user name and password for the Oracle VM Manager.

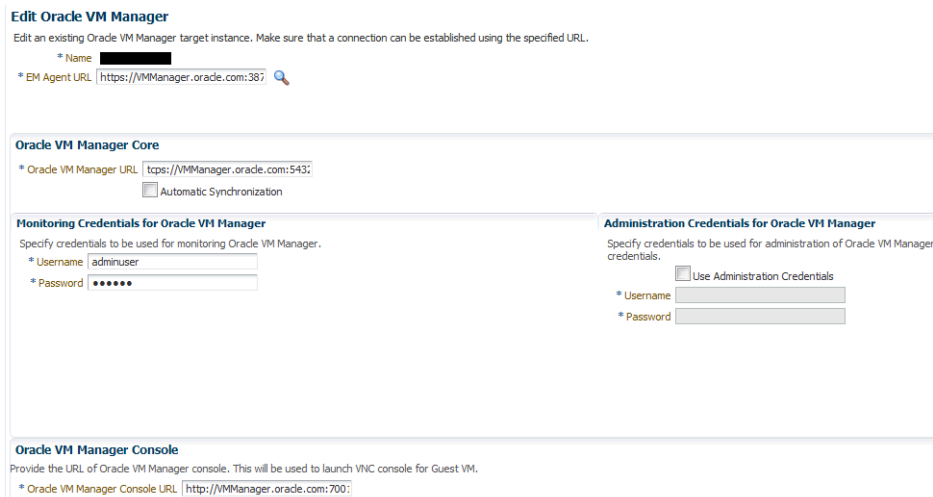
Note: These credentials are required to launch the OVM Manager. These are not operating system credentials for the host on which the OVM Manager has been installed.

6. Enter the URL for the Oracle VM Manager Console. For example, `https://ovmmgr.example.com:7002/ovm/console`
7. 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.
8. 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.

5.2.1 Registering Oracle VM Manager 3.1.x With Enterprise Manager

The older version of Oracle VM Manager 3.1.x is certified for registration with Enterprise Manager 12c Cloud Control. To register the Oracle VM Manager 3.1.x in Enterprise Manager 12c, ensure that the correct tcps URL (`tcps://hostname`) has been specified. A sample screenshot is shown below:

Figure 5–4 Registering the Oracle VM Manager



5.2.2 Upgrading from Oracle VM Manager 3.0.x

When upgrading from Oracle VM Manager 3.0.x to Oracle VM Manager 3.1.x, 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) URL in the Oracle VM Manager field.
3. You are prompted to import the OVM Manager certificate to the Agent Keystore.
 - 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>`
4. After tcps has been configured, restart the OVM Manager, then you'll see that Oracle VM Manager is now listening for tcps requests on port 54322 by default.

Note: If the OVM Manager and the Management Agent are running on the same host, edit the OVM Manager target and specify the URL as `tcp://localhost:54321`.

5.2.3 Discovering and Promoting the Oracle VM Manager (Optional)

Automatic discovery is a process that detects hosts, virtual servers 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 5.2, "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. From the **Setup** menu, select **Add Target**, then select **Configure Auto Discovery**.
2. In the Configure Auto Discovery section, click **Configure in the Discovery Setup** table for the row, Hosts and Virtual Server Discovery Using IP Scan. Enterprise Manager displays the Host Discovery (Agentless) page. Use this page to search, browse, and create host discovery specifications.
3. Enter the Name, Owner, and Description information and then click **Search** to find a host. The results appear in the Host table. Use the Scan Details tab to select an agent that performs the IP scan, and 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.
4. 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.
5. You can use the Host Discovery page to check the status of discovery and view newly discovered hosts.

6. After the target has been discovered, from the **Setup** menu, select **Add Target**, then select **Auto Discovery Results**. Enterprise Manager displays the Discovered Unmanaged Targets page where you can review discovered unmanaged targets and promote targets to be managed by Enterprise Manager for monitoring.
7. On the Non-Host targets tab, view the target discovery results. Select the OVM Manager target and click Promote. The OVM Manager Registration page is displayed. See [Section 5.2, "Registering the Oracle VM Manager"](#) for details.

5.3 Synchronizing the OVM Manager Targets (Optional)

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 use the Synchronize option to add them to Enterprise Manager. To synchronize the OVM Manager targets, follow 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**. 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.

5.4 Discovering an Oracle VM 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 5–5 Discover Virtual Servers

The screenshot shows the Oracle VM Manager interface for discovering virtual servers. The page title is 'ovm_manager' and 'Oracle VM Manager'. The main heading is 'Discover Virtual Servers'. Below this, there is a 'Virtual Servers' section with a text area for specifying host names or IP addresses. Below that is an 'Oracle VM Agent Credentials' section with input fields for 'User Name' (containing 'oracle') and '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.

5.5 Rediscovering a Virtual Server

To rediscover a virtual server, follow these steps:

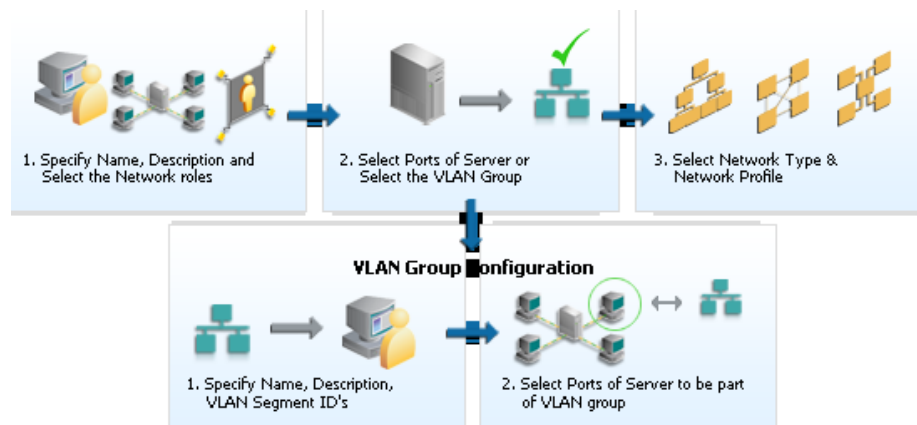
1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Right click on a Virtual Server target in the left panel and select **Rediscover Virtual Server**.
3. A confirmation message appears. Click **OK** to rediscover the virtual server and discover any new NICs and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

5.6 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 5–6 Setting Up Networks



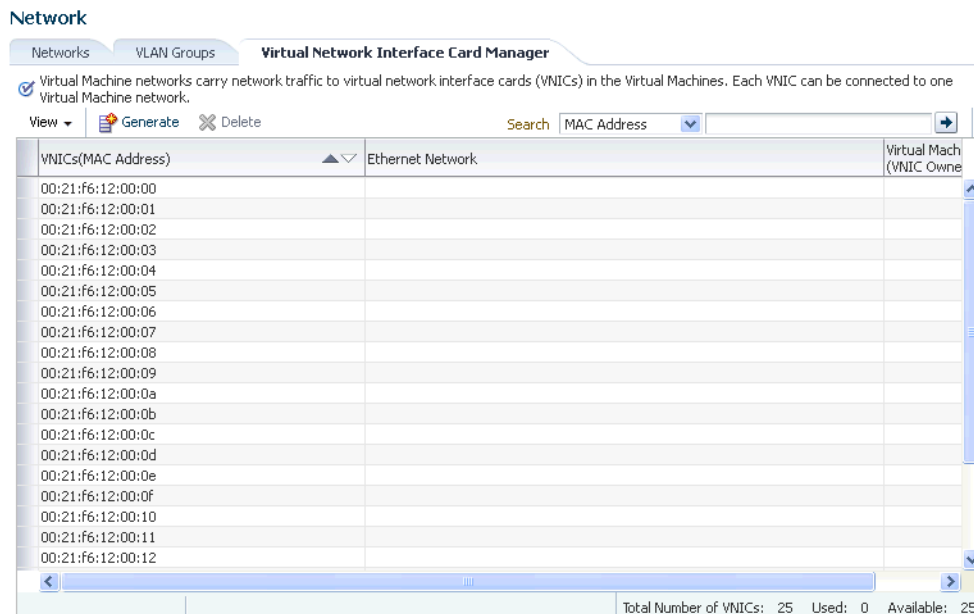
5.6.1 Generating MAC Addresses

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 5–7 Virtual Network Interface Card Manager



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.

5.6.2 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 4095.
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 two 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.

5.6.3 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 5.6.2, "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 5–8 Create Network: General Page

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 5.6.2, "Configuring a VLAN Group"](#).

Figure 5–9 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

Port	MAC Address	MTU	Address Type	IP Address	Netmask	Bonding Mode
▼ vlan_120102, 100	network.EthernetPort (5) in zowu.us.or 00:16:3e:3a:65:8a	1500	None			

Ports

View

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

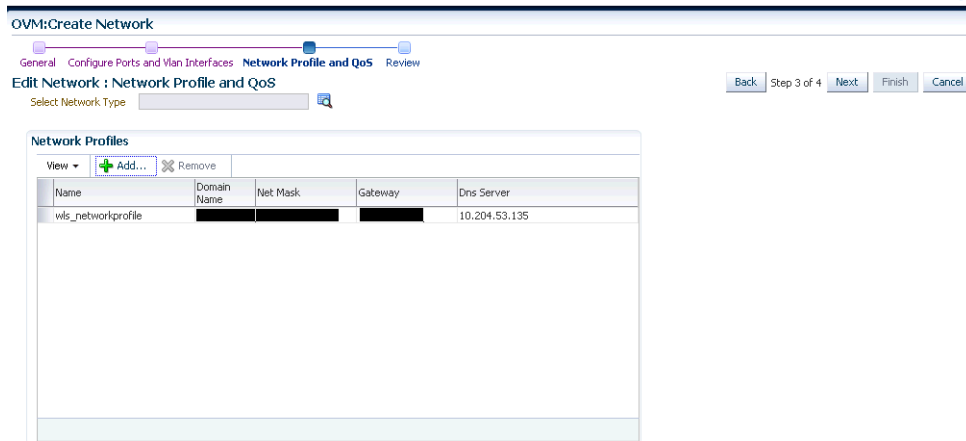
8. After adding the VLAN Segment, click **Add** the VLAN Interfaces and the ports. Select the ports to be added to network. A maximum of two ports can be added. If you select more than two ports, they will be bonded or aggregated into a single interface. Select the bonding mode which can be:
- **Active 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 5–10 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.
- **Network Profile:** A network profile defines a set of IP Addresses, their associated host names (optionally), and common networking attributes for them. See [Section 5.7, "Creating a Network Profile"](#) for details.

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

5.7 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 5–11 Network Profile



To create a network profile, follow these steps:

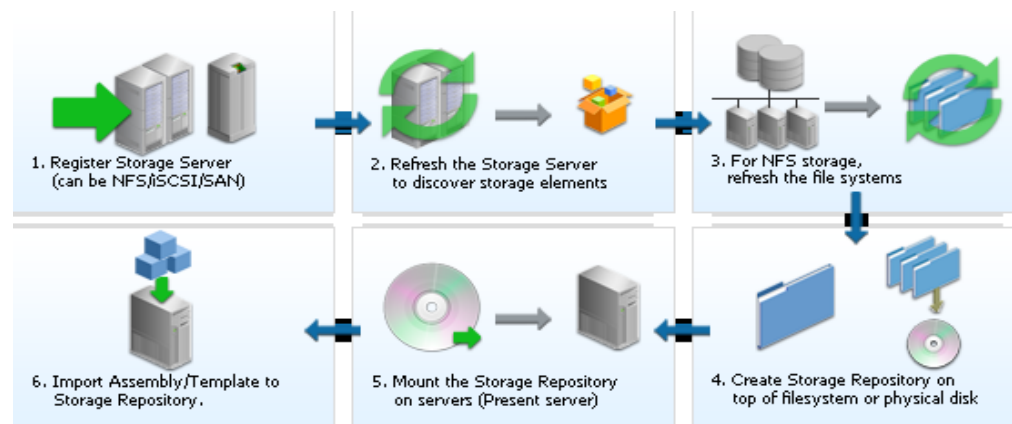
1. From the **Setup** menu, select **Provisioning and Patching**, then select **Network Profile**.

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.
 - **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.

5.8 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.

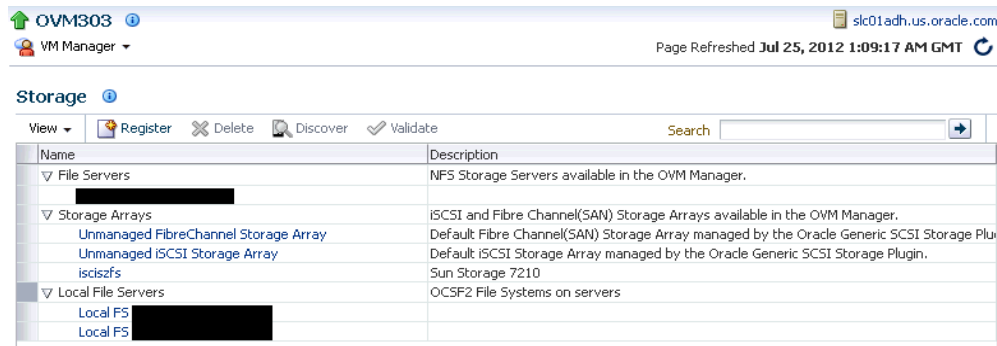
Figure 5–12 Creating Storage Servers and File Systems



To set up a storage server, follow these steps:

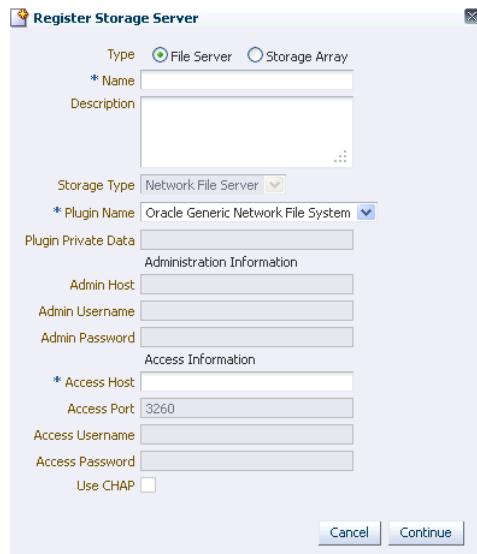
1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Right click on an OVM Manager target in the left panel and select **Manage Storage**. [Figure 5–13](#) appears.

Figure 5–13 Storage Server



3. A list of file servers, storage arrays, and local file servers that have been defined appears. Enter the following details for the Register Storage Server window:

Figure 5–14 Register Storage Server



- **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 arrays:
 - **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).
- **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).

- **Access Host:** Host name or IP address of the server offering external storage.
 - **Access Port:** Port number used to access the storage.
 - **Admin Host:** Host name of the Administration Server.
 - **Admin Username and Password:** User name and password for the Administration Server.
4. Click **Continue**. In the Admin Server window, 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.
 5. Click **Submit** to create the storage server. The new server appears in the table. At this point, the Storage Connect plug-in has established a link to the storage location. No storage has been mounted yet.
 6. 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.
 7. 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. Select an Administration Server from the list and click **Edit** to configure it.

In the Storage QOS Mapping section, map the storage type defined by the Cloud Administrator to the QOS defined by the Storage Server.
 - **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.
 8. Click a **Storage Array** link to view the Details page. This page has four tabs:
 - **General:** Provides general information on the storage array. Configuration details and Storage QoS Mapping for the array appears.
 - **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.
 9. 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.

5.9 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.

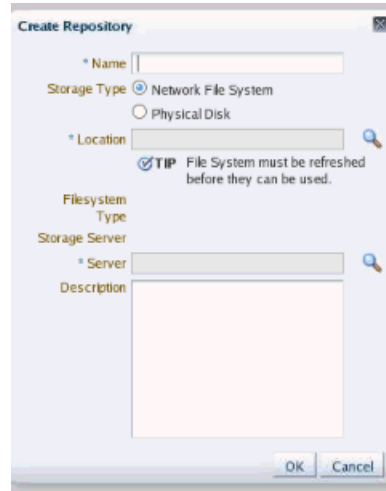
To create a storage repository, follow these steps:

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 5–15 Storage Repository

Name	Owned	File System Size	File System Size Used (%)	Disk	Share Path	Storage Server
cloud_storage_new	Another	10,896,426	85	nfs:/export/qacloudspace	/repo5/	cloud_storage_303
repo20	Another	10,896,426	85	nfs:/export/qacloudspace	/new_repo/	cloud_storage_303
repo10	Another	10,896,426	85	nfs:/export/qacloudspace	/repo10/	cloud_storage_303
nfs_repo	Another	10,896,426	85	nfs:/export/qacloudspace	/repo2/	cloud_storage_303
nfs_repo	Another	10,896,426	85	nfs:/export/qacloudspace	/repo8/	cloud_storage_303
storage_repo	Another	10,896,426	85	nfs:/export/qacloudspace	/repo7/	cloud_storage_303
Repo22	Another	10,896,426	85	nfs:/export/qacloudspace	/repo22/	cloud_storage_303
cloud_repo_ovm_303	Yes	10,896,426	85	nfs:/export/qacloudspace	/repo13/	cloud_storage_303
repo4	Another	10,896,426	85	nfs:/export/qacloudspace	/repo4/	cloud_storage_303
storageRepo-cloudsp	Another	10,896,426	85	nfs:/export/qacloudspace	/repo18/	cloud_storage_303
storageRepo-cloudsp	Another	10,896,426	85	nfs:/export/qacloudspace	/repo21/	cloud_storage_303
storage_repo_cloud	Another	10,896,426	85	nfs:/export/qacloudspace	/repo1/	cloud_storage_303
repo_311	Another	10,896,426	85	nfs:/export/qacloudspace	/repo3/	cloud_storage_303
storageRepo-cloudsp	Another	10,896,426	85	nfs:/export/qacloudspace	/repo14/	cloud_storage_303
repo19	Another	10,896,426	85	nfs:/export/qacloudspace	/repo19/	cloud_storage_303
cloud_repo_storage	Another	10,896,426	85	nfs:/export/qacloudspace	/repo11/	cloud_storage_303
cloud_repository	Another	10,896,426	85	nfs:/export/qacloudspace	/repo6/	cloud_storage_303
repo23	Another	10,896,426	85	nfs:/export/qacloudspace	/repo23/	cloud_storage_303
new_repo_cloud_auti	Another	10,896,426	85	nfs:/export/qacloudspace	/repo9/	cloud_storage_303

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 5–16 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.

5.9.1 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.
- **Delete:** Select a repository and click **Delete**. All the contents of the selected repository along with its associations will be deleted.
- **Refresh:** Select a repository and click **Refresh**. The contents of the selected repository including assemblies, templates, ISO images, templates, and virtual disks will be refreshed.
- **Discover:** Select this option to discover a repository that has been created externally through the Oracle Virtual Manager console.

5.10 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 5–17 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.

5.11 Importing Assembly, Template, and Virtual Disk Resources into a Repository

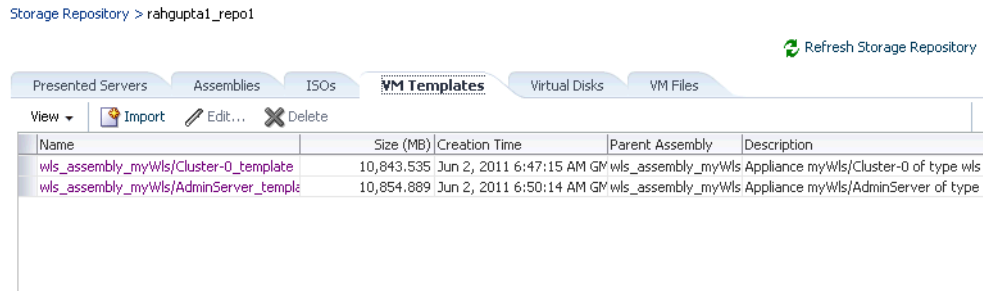
A storage repository associated with an Oracle VM server stores binary components associated with an Oracle VM servers. The following software components can be stored in the storage repository:

- 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. Assemblies can be imported as a set of `.ovf` (Open Virtualization Format) and `.img` (disk image) files, or it may all be contained in a single `.ova` file.

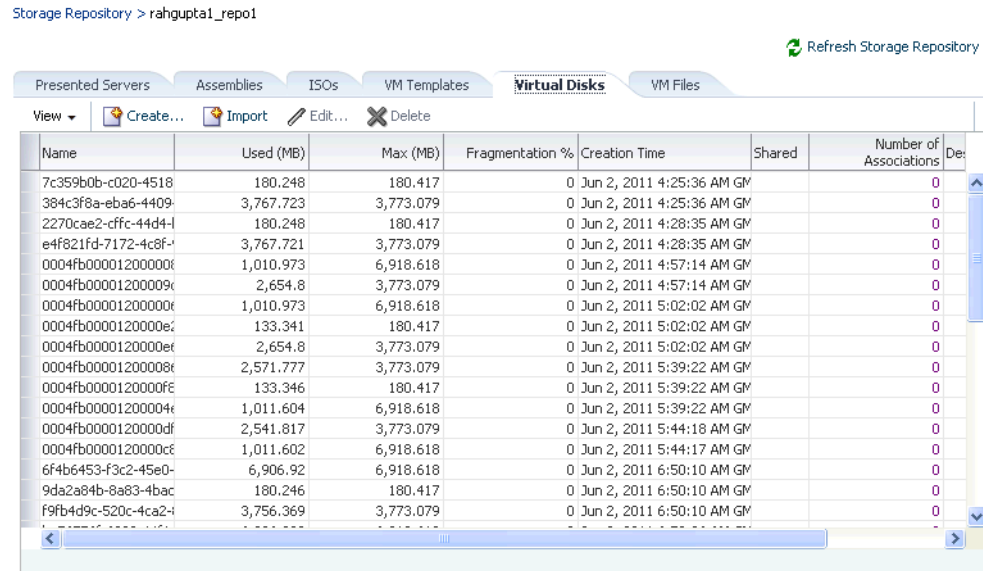
Figure 5–19 Storage Repository - VM Templates



- 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 5–20 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.
- After the disk has been created, select the template and click **Assign** to assign this disk to a server pool.
- 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.

5.11.1 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.
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.

5.12 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 Virtual Machine 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 virtual server pool can contain a maximum of 32 Oracle VM servers.
- There must be at least one virtual server in the server pool. One of the servers acts as a Master Server for the 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 and the IP address to use as the Master Server's virtual IP address are available.
- 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 5–21 Create Virtual 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**
- **Secure VM Migrate:** Select this check box to ensure that live migration of guest virtual machines will be secured.
- **Activate Cluster:** Select this check box to enable high availability on the server pool.
- **Cluster Timeout**
- **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 and cluster data. This is a dedicated file system (either a NAS export or LUN) and can be 1 GB in size.
- **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 5.4, "Discovering an Oracle VM Server"](#) for details.

5. Select the **Activate Cluster** .

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

5.12.1 Editing a Virtual Server Pool

To edit a virtual server pool, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Right click the virtual server pool that you wish to edit from the left panel and select **Edit**.
3. You can edit the following details:
 - **Description:** A description of the server pool.
 - **Master Server:** Select the master server for the virtual server pool.
 - **Keymap:** Specify the keyboard type for the guest virtual machines created in this server pool
 - **Secure VM Migrate:** Select this check box to ensure that live migration of guest virtual machines will be secured.
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 5.4, "Discovering an Oracle VM Server"](#) for details.

5. Click **Add** in the Anti-Affinity Group section. 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.
6. Click **OK** to apply the changes to the virtual server pool.

5.13 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 5–22 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 **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 5.12, "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.

5.14 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 5–23 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 **OVIManager** 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.


5.14.1 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 5–24 Network Types



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

5.14.2 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.

5.14.3 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.

5.14.4 Defining the Request Archiving 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.

5.15 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

5.15.1 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.

5.15.2 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.

5.16 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.

Setting Up and Using the Self Service Application for IaaS

Self service provisioning is an important aspect of the cloud infrastructure. It is used to delegate some of the infrastructure provisioning and management to self service administrators (SSA) so that cloud administrators can focus on more strategic initiatives. After the cloud components have been configured, you (SSA user) can log into the Self Service Application and identify the images required, reserve some resources (for example, compute power, storage, memory, and so on) to provision the images to, and specify how long the images will be required. This chapter contains the following sections:

- [Setting Up the Self Service Portal](#)
- [Using the Self Service Portal](#)
- [Requesting Servers](#)

6.1 Setting Up the Self Service Portal

The Self Service Portal allows administrators to configure the cloud and perform various operations such as defining policies for the cloud, exposing it to certain users, deciding the total amount of resources each user can reserve, and so on. Users can log in to the Self Service Portal and identify the images they need, reserve some resources (for example, compute power, storage, memory, and so on) and define policies specific for 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 2.5, "Managing Users and Roles"](#) 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 **Machines 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' interface. On the left is a navigation menu with 'Machine Sizes' selected. The main area displays a table of machine sizes with columns for Name, Description, VCPUs, Memory (MB), and Local Storage (GB). The table contains three rows: Small (2 VCPUs, 4096 MB, 250 GB), Medium (4 VCPUs, 8192 MB, 500 GB), and Large (8 VCPUs, 15360 MB, 1000 GB). Above the table are buttons for 'Create...', 'Edit...', and 'Delete'.

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 SSA user's request is being processed.

Note: If you have chosen to configure the Management Agent, you must ensure that the user has:

- Write permissions to the directory specified in the **Installation Base Directory** field.
 - Write permissions to the `/home/oracle/agent` directory specified during assembly or template deployment.
-
- **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 Feb 20, 2012 6:33:04 PM PST

Request Settings Apply Revert

Future Reservation
How far in advance a user can make a request:
 No Restriction
 Restricted Reservation Days

Request Duration
Maximum duration for which requests can be made:
 No Restriction
 Restricted Duration Days

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

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
SSA_USER_ROLE1	50	70	100	1000	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Zone	NW,NW

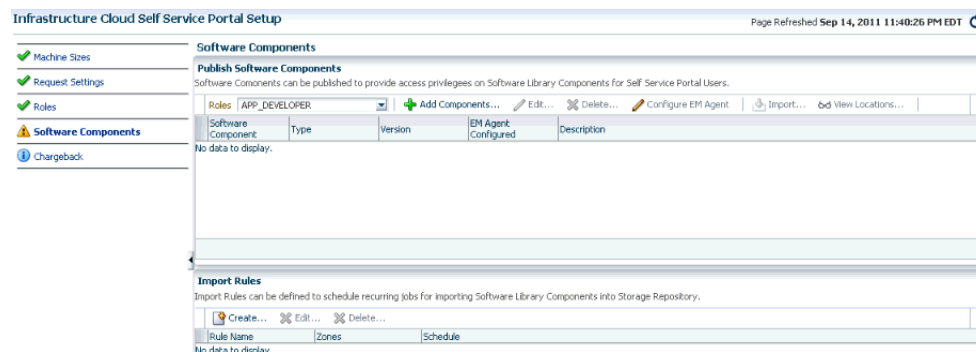
- In the Roles page, click **Assign Quota to Role** to define quota at the role level. The quota assigned to a role is available to all users with that role. For example, if the Number of Servers for the SSA_USER_ROLE1 is set to 5, all users with this role can have a maximum of 5 servers. The quota is 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.

Select a role and click **Edit**. You can modify the role mappings in the **Assign Quotas, Zones & Network Profiles to Role** page.

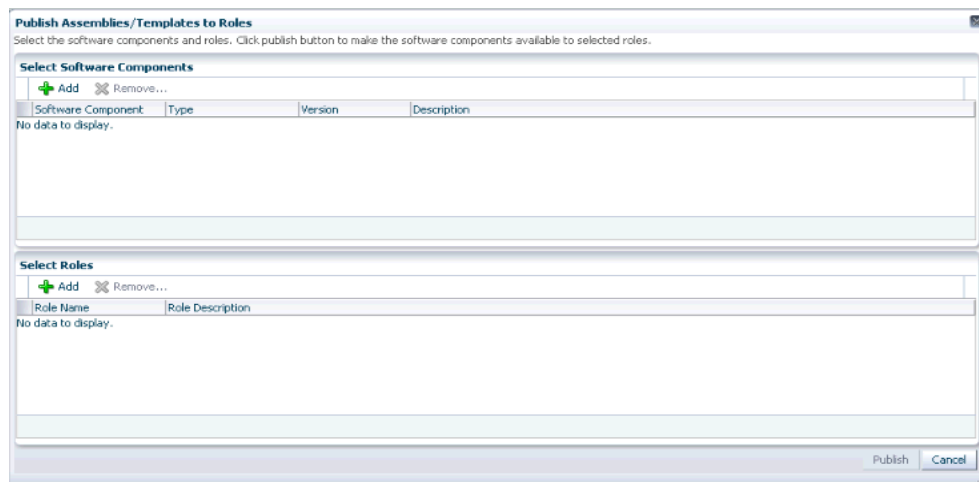
5. Click **Software Components** in the left panel to select the set of templates or assemblies that are accessible to the SSA users. You can publish the assemblies and templates to the SSA Portal and provide SSA users with access to these software components.

Figure 6–4 Self Service Portal Setup - Software Components



6. Click **Add Components**. The Publish Assemblies / Templates to Roles page appears.

Figure 6–5 Publish Assemblies / Templates to Roles



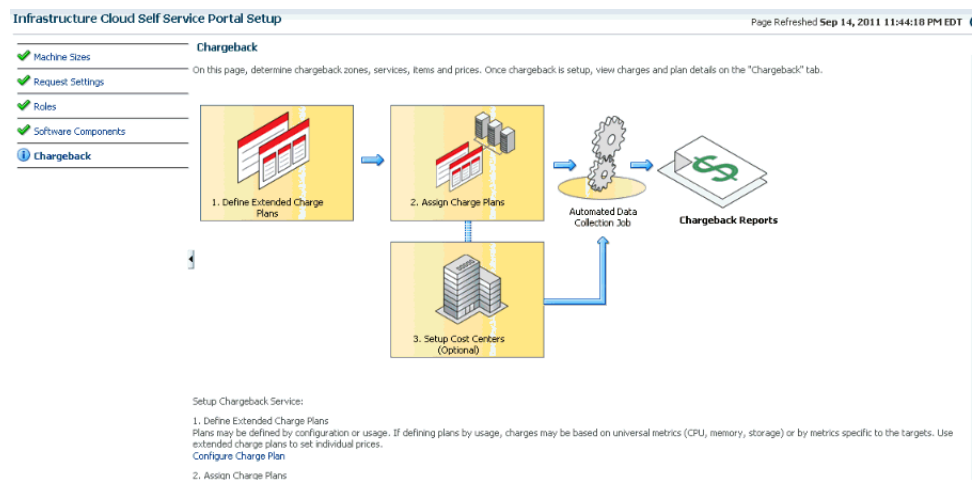
In this page, do the following:

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.

7. You can import the published assembly to a selected set of server pools and zones. Select the assembly and click **Import**. In the Zone and Server Pool selection page, select one or more zones and pools and click **Save** to import the assembly to the selected server pools and zones.
8. 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.
9. 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



10. 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 4, "Chargeback Administration"](#).

6.2 Using the Self Service Portal

The Self Service Portal 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 expiry date. Click a Server Name to view the server details. See [Section 6.2.2.1, "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 6–7 Self Service Portal

The screenshot shows the 'Infrastructure Cloud Self Service Portal' interface. The top navigation bar includes 'Home', 'My Requests', 'My Servers', 'Storage', 'Chargeback', 'My Library', and 'Policies'. The 'My Servers' tab is selected. On the left, there are 'Notifications' and 'Usage' sections. The 'Usage' section shows: Servers 1, CPUs 2, Memory (GB) 2, and Local Storage (GB) 12. The main content area is divided into two sections: '10 Last Requested Servers' and '10 Latest Requests'. The '10 Last Requested Servers' section shows a table with columns: Name, Status, Zone, CPUs, Memory (MB), Storage (GB), Charge, and Creation Date. The '10 Latest Requests' section shows a table with columns: Name, Status, Submission Date, Start Date, End Date, Type, Servers, Total CPU, Total Mem, and Total Storage (GB). The '10 Latest Requests' section currently displays 'No items found.'

3. Click **Request Servers**. The New Server Request wizard appears. See [Section 6.3, "Requesting Servers"](#) for details.

6.2.1 Viewing My Requests

The My Requests tab lists all the requests made by the SSA user.

Figure 6–8 My Requests

The screenshot shows the 'Infrastructure Cloud Self Service Portal' interface with the 'My Requests' tab selected. The page displays a table of requests with the following columns: Name, Status, Submission Date, Start Date, End Date, Type, Servers, Total CPU, Total Mem, and Total Storage. The table contains 14 rows of request data.

Name	Status	Submission Date	Start Date	End Date	Type	Servers	Total CPU	Total Mem	Total Storage
- Tue Feb 21 02:29:42 PST 2012	Execution Error	Feb 21, 2012	Feb 21, 2012 2:30		Scale Up	1	1	1024	11.17
- Tue Feb 21 02:26:29 PST 2012	Successful	Feb 21, 2012	Feb 21, 2012 2:26		Scale Down				
- Tue Feb 21 02:19:52 PST 2012	Successful	Feb 21, 2012	Feb 21, 2012 2:22		Scale Up				
- Tue Feb 21 02:17:40 PST 2012	Scheduled	Feb 21, 2012	Feb 21, 2012 2:33		Scale Up				
- Tue Feb 21 02:13:03 PST 2012	Successful	Feb 21, 2012	Feb 21, 2012 2:13		Scale Down				
- Tue Feb 21 01:50:05 PST 2012	Successful	Feb 21, 2012	Feb 21, 2012 1:51		Scale Up				
- Tue Feb 21 01:38:58 PST 2012	Execution Error	Feb 21, 2012	Feb 21, 2012 1:42		Scale Up				
- Tue Feb 21 01:34:58 PST 2012	Successful	Feb 21, 2012	Feb 21, 2012 1:35		Scale Down				
- Tue Feb 21 00:46:38 PST 2012	Successful	Feb 21, 2012	Feb 21, 2012 12:4		Assembly Deploy	3	3	3072	33.51
- Sat Feb 18 04:19:03 PST 2012	Execution Error	Feb 18, 2012	Feb 18, 2012 4:28		Assembly Deploy	6	6	7168	72.98
- Sat Feb 18 02:44:25 PST 2012	Successful	Feb 18, 2012	Feb 18, 2012 2:48		Template Deploy	2	2	1024	23.89

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.

6.2.2 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 6–9 My Servers

Name	Type	Operating System	Server Size	CPUs	Memory (MB)	Storage (GB)	Status	Expiry Date
▼ SidZone	Zone							
▼ mySite1	Assembly Instance							
▼ mySite/myWsj/Cluster-0:mySite1	Tier Instance							
us.oracle.com	Server	None	Custom	1	1024	11.94	↓	
us.oracle.com	Server	None	Custom	1	1024	11.94	↑	
us.oracle.com	Server	None	Custom	1	1024	11.94	↑	
us.oracle.com	Server	None	Custom	1	512	11.94	↓	
us.oracle.com	Server	None	Custom	1	512	11.94	↓	

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 6.3, "Requesting Servers"](#) for details.
- **Assemblies:** You can start, stop, and delete an assembly instance. 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, and launch a VNC Console.

6.2.2.1 Server Details Page

The Server Details 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.

6.2.2.2 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 7.9.4, "Provisioning Guest Virtual Machines Using Oracle Virtual Assemblies \(OVA\)"](#) for details. To scale up a tier instance, 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 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, 3 new servers can be added if the existing server is Up. If the existing server is down, 4 servers can be added to the tier. 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.

6.2.2.3 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.

6.2.3 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 delete existing shared or non-shared storage disks.

6.2.4 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.

6.2.5 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 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 **Create Software Component**. 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 **Create** to upload to software to My Library. If you see an error message while uploading the file: `Error: The file is too large. Please add a smaller file, you must modify the configuration settings in the web.xml file:`

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.

6.2.6 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, see [Section 8.1, "Managing Cloud Policies"](#).

6.2.7 Viewing My Preferences

Click **Apply** to save the preferences. The values specified here will be used when you request resources using SSA.

Click **My Preferences**. Before you request resources using SSA, you can set up your user profile by specifying the following details:

- Default Host Password: The default password for the host.
- Default Zone: The default zone for the requests.

- **Default Source Software:** The default location in which the assemblies, templates, and deployment plans are present.
- **Default Email Address for Notifications:** The default email address to be used to receive notifications.
- **Days Before the First / Second Notification:** The number of days before the first and second notifications are sent.
- **Enable Quota Related Notification:** If you select this checkbox, notifications are sent when the quota reaches a specified threshold.
- **Quota Threshold for Receiving Notifications:** If you have selected the **Enable Quota Related Notification** checkbox, specify the quota threshold in this field.
- **Database Preferences**
 - **Default Zone:** Select the default zone for the requests.
- **My Middleware Preferences**
 - **Default Zone:** Select the default zone for the requests.
- **My Servers Preferences**
 - **Default Server Root Password**
 - **Default Zone for New Requests**
 - **Default Source Software**

6.3 Requesting Servers

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 6–10 New Server Request: General

The screenshot shows the 'New Server Request: General' form. At the top, there are tabs for 'General', 'Deployment Configuration', 'Schedule', and 'Review'. Below the tabs, there are navigation buttons: 'Back', 'Step 1 of 4', 'Next', 'Save As Plan...', 'Finish', and 'Cancel'. The form is divided into several sections:

- Request Name:** 'Specify a name for your request. This will help you track it later.' The 'Name' field contains a name and a timestamp: 'Thu Jun 02 16:09:32 UTC 2011'.
- Destination:** 'Select the zone in which servers will get created.' The 'Zone' dropdown menu is set to 'West_Coast_Zone'.
- Source:** 'Select the assembly or template to be used for this request.' The 'Source' field contains a name and a search icon. The 'Description' field contains 'mySite.Jtimeout+oeb5-5.o.va'.
- Assembly Instance Name:** The 'Name' field contains 'test_assembly'.
- Deployment Plan (Optional):** 'Deployment plans allow you to record inputs and use them later while creating new requests using standard values.' The 'Name' field contains a name and a search icon. The 'Description' field is empty.

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 assemblies, see

Section 7.9.4, "Provisioning Guest Virtual Machines Using Oracle Virtual Assemblies (OVA)". For details on templates, see Section 7.9.5, "Provisioning a Guest Virtual Machine Using Oracle VM Templates".

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**. In the Deployment Configuration page, specify the resource configuration parameters such as number of servers, server size, network, and storage details for the assembly and click **Next**. See Section 5.6, "Setting Up Networks" and Section 5.9, "Creating a Storage Repository" for details.

Figure 6–11 New Server Request: Deployment Configuration

Name	Number of Servers				Auto Scalable	Server Size	Deployment Option	Server Name
	Default	Minimum	Maximum	Actual				
test_assembly								
myWls/AdminServer:1test_assembly	1	1	1	1	<input type="checkbox"/>	Custom		
myWls/Cluster-0:1test_assembly	4	1	4	4	<input type="checkbox"/>	Custom		myWls/Cluster-0_vm

Name	Default Value	Value	Category	Description
sysprop-wls-home	/mwHome/wlserver_10.3	/mwHome/wlserver_10.3	Assembly Properties	
sysprop-adminvr-jmx-protocol	iiop	iiop	Assembly Properties	
sysprop-adminvr-jmx-input	Default	Default	Assembly Properties	
sysprop-jive-workdir	/mwHome/user_projects	/mwHome/user_projects	Assembly Properties	
sysprop-admin-username	weblogic	weblogic	Assembly Properties	
sysprop-applications-dir	/mwHome/user_projects	/mwHome/user_projects	Assembly Properties	
sysprop-admin-password	{1}a613a04f736c14322f62c759actb9290		Assembly Properties	
sysprop-domain-name	base_domain	base_domain	Assembly Properties	
sysprop-domain-root	/mwHome/user_projects	/mwHome/user_projects	Assembly Properties	
sysprop-useOracleHomes	false	false	Assembly Properties	

7. Specify the schedule for the request and 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 the Cloud for IaaS

This chapter describes the monitoring and administrative tasks that you can perform on the zones, virtual server pools, virtual servers and guest virtual machines. It includes the following sections:

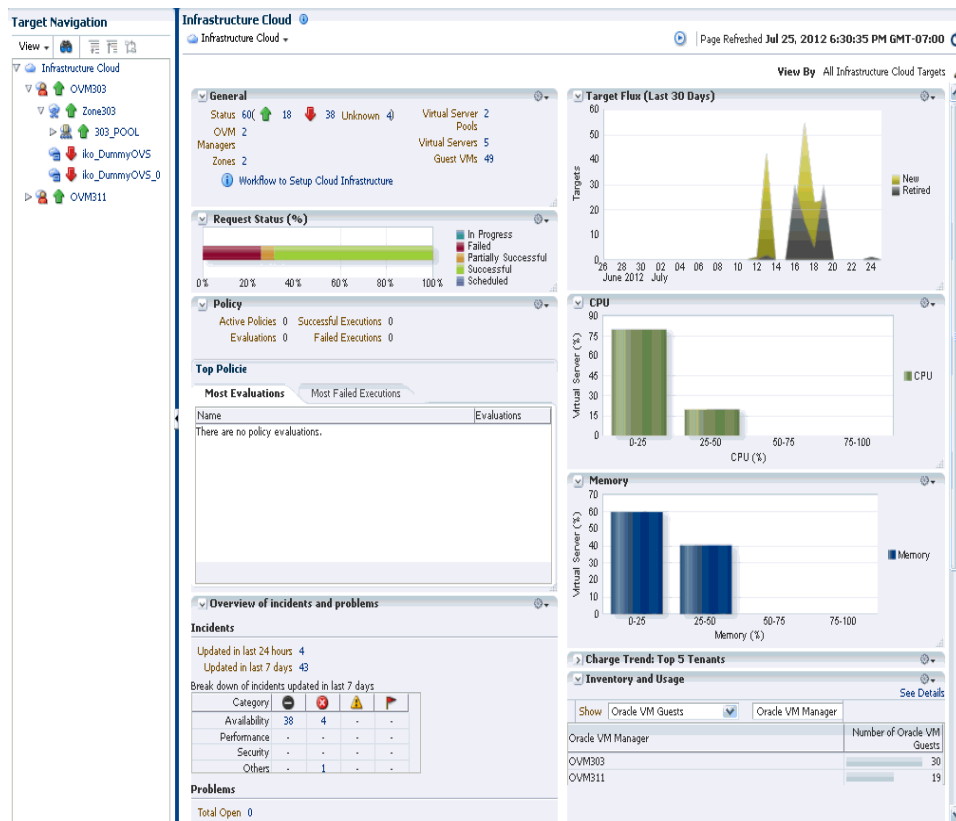
- [Viewing the Infrastructure Cloud Home Page](#)
- [Viewing the OVM Manager Home Page](#)
- [Viewing and Managing Members](#)
- [Managing Zones](#)
- [Virtual Server Pool Home Page](#)
- [Managing Virtual Servers](#)
- [Managing Guest Virtual Machines](#)
- [Deploying Guest Virtual Machines](#)

7.1 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 **Infrastructure Cloud**, then select **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 7-1 Infrastructure Cloud Home Page



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 guest virtual machines created and deleted over the last 30 days.
- **Request Status:** Shows the status of various requests and reservations. Check the **Only SSA Requests** check box to view only the Self Service Application Requests. 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.
- **Resources:** For each entity in the cloud you can view the total storage used, number of network profiles and VLAN groups. Click the Network or Storage link to drill down to the Network or Storage page of the corresponding OVM Manager target.

- **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.

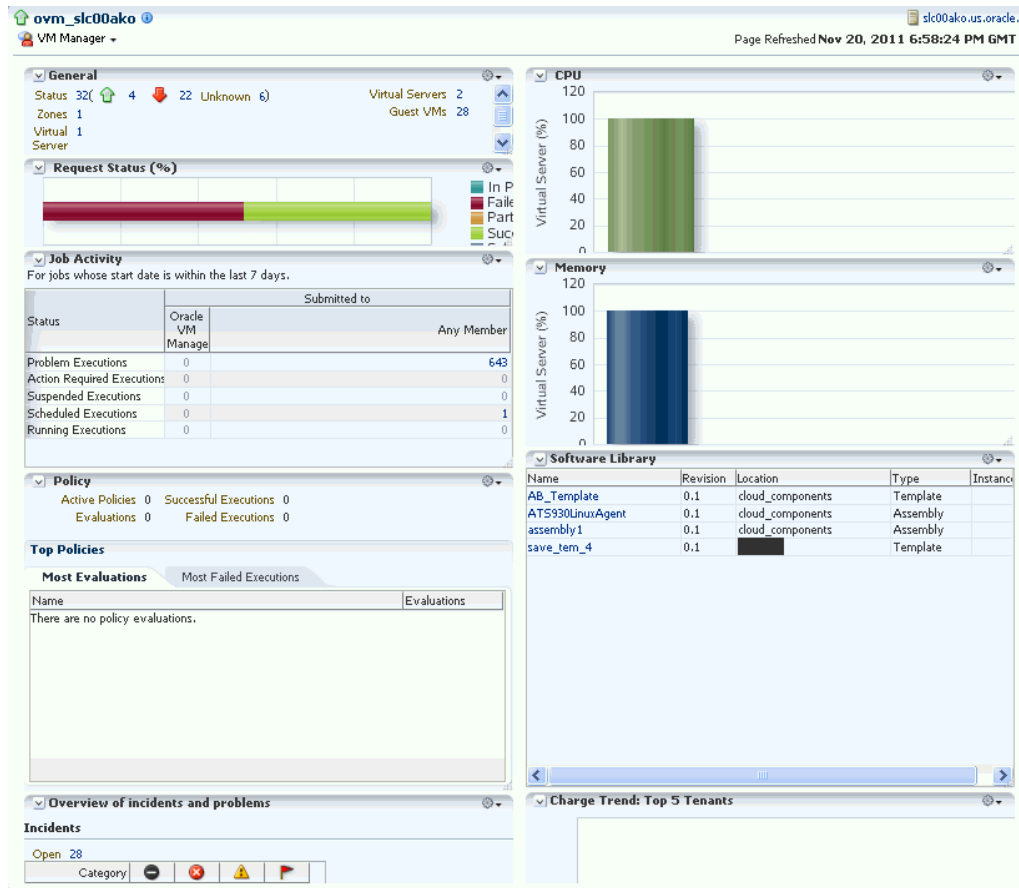
Before you can start monitoring a virtualization target, you must register the OVM Manager target. See [Section 5.2, "Registering the Oracle VM Manager"](#) for details.

7.2 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 **Infrastructure Cloud**, then select **Home**.
2. Click on an OVM Manager target from the left panel. The OVM Manager Home page shows the following details:

Figure 7-2 OVM Manager Home Page



- **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

7.3 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 Home** menu, select **Members**, then select **Show All**.

Target Name	Status	Type	Operating System	Memory Util (%)	Memory (MB)	CPU Util (%)	CPU
ovm_manager	↑	OVM Manager					
West_Coast_Zon	↑	Zone			0		0
[REDACTED]	↑	Zone			16383		2
test	↑	Zone			0		0
Pool_on_CDC	↑	Virtual Server Pool			3995		1
Unknown pool fi	↑	Virtual Server Pool			0		0
[REDACTED]	↑	Virtual Server			16383		2

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.

7.4 Managing Zones

This section describes the following:

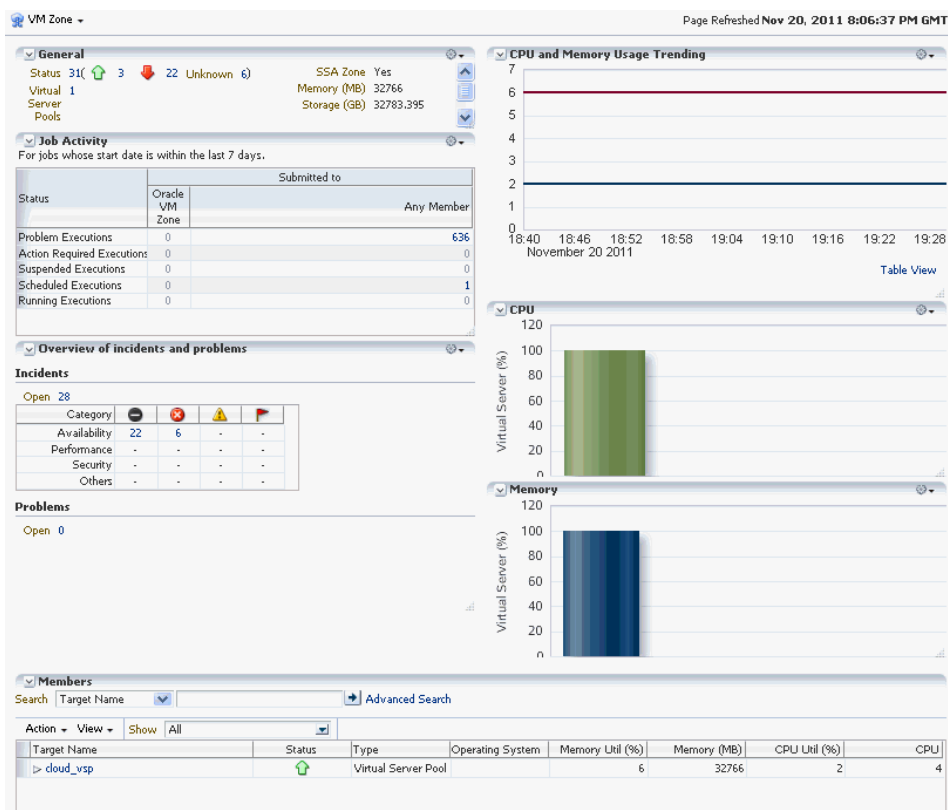
- [Viewing the Zone Home Page](#)
- [Creating or Editing a Zone](#)

7.4.1 Viewing the Zone Home Page

A zone is a collection of virtual server pools which share storage. 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 7-3 Zone Home Page



- **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.
- **Members:** This region shows the status and other details of each virtual server pool, virtual server and guest virtual machines in the zone. Click on a link to drill down to the Members page. 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.
- **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.

7.4.2 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 **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 5.12, "Creating a Virtual Server Pool"](#) for details.
6. Click **Create Zone**. 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 **Infrastructure Cloud**, then select **Home**.
 2. Select a zone target from the left panel.
 3. From the **Zone** menu, select **Target Setup**, then select **Edit VM Zone**.
-
-

7.5 Deleting a Zone

To delete a zone, follow these steps:

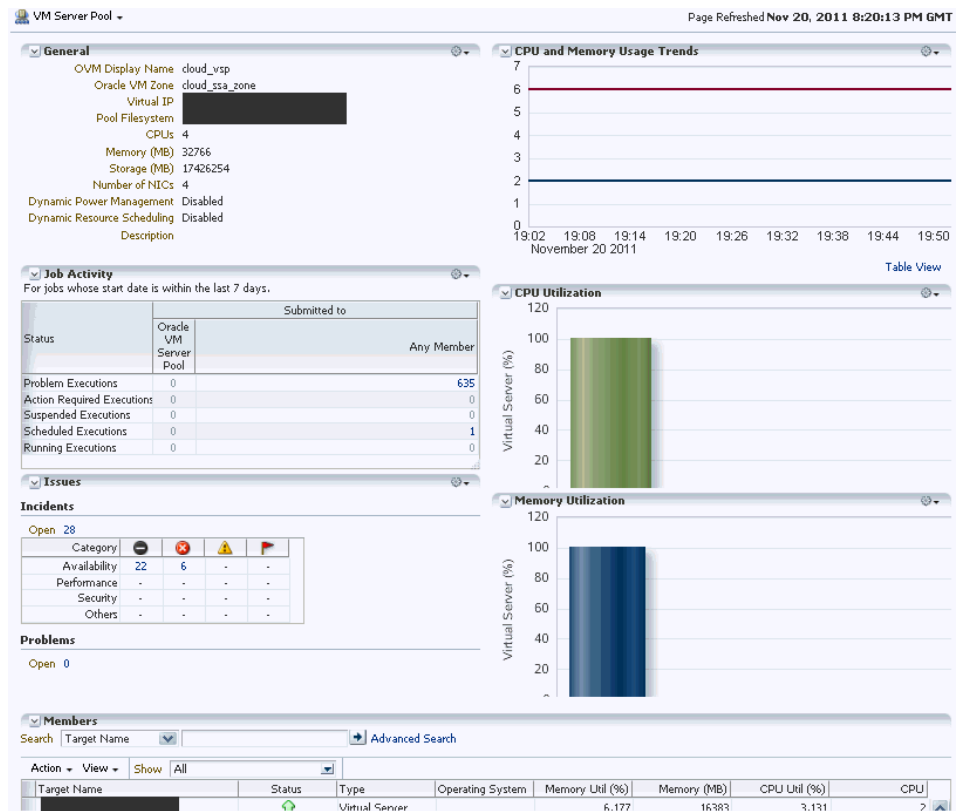
1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Home**.
2. Click on a zone target from the left panel to navigate to the Zone Home page. From the **Zone** menu, select **Target Setup**, then select **Remove Target**. A confirmation message appears.
3. Click **OK** to delete 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.

7.6 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 7–4 Virtual Server Pool Home Page



- **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.
- **Members:** This region shows the status and other details of each virtual server and guest virtual machines belonging to the virtual server pool. Click on a link to drill down to the Members page. 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.
- **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.

7.7 Managing Virtual Servers

This section describes the following:

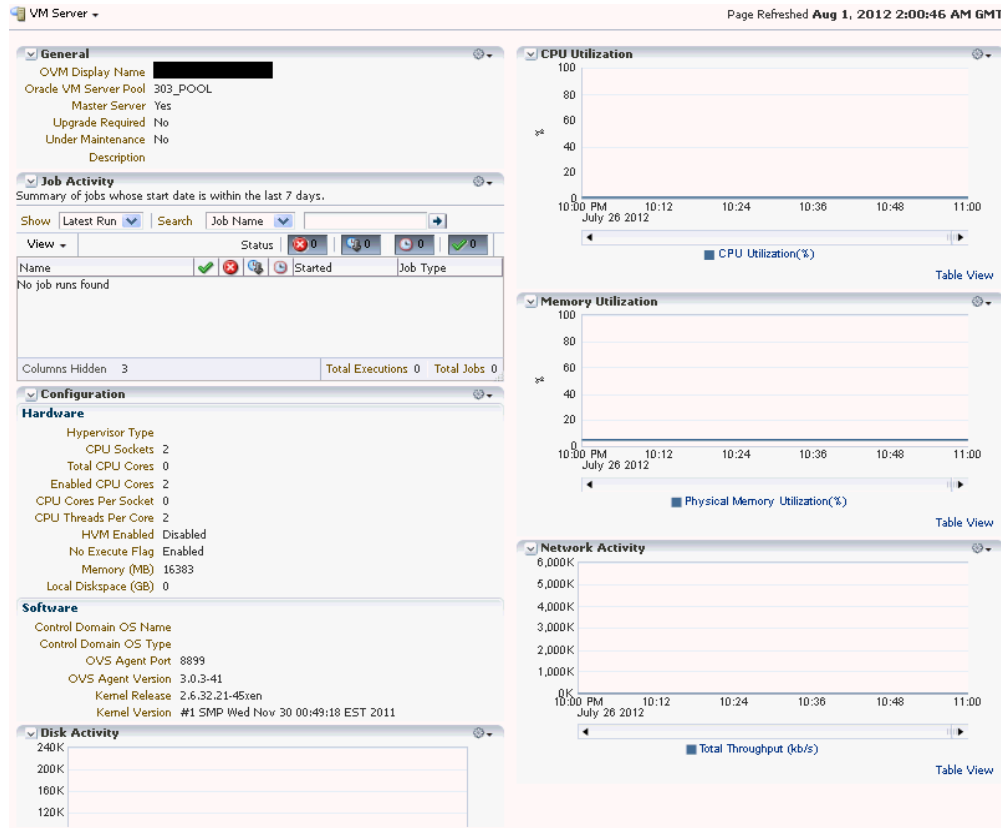
- [Virtual Server Home Page](#)
- [Editing a Virtual Server](#)
- [Upgrading a Virtual Server](#)
- [Starting and Stopping a Virtual Server](#)

- [Performing Maintenance on the Virtual Server](#)

7.7.1 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 7-5 Virtual Server Home Page



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.
- **Disk Activity:** The disk throughput of the virtual server.
- **Members:** This region shows the status and other details of all the guest virtual machines. Click on a link to drill down to the Members page. Select a target, click Guest VMs from the Actions menu. Select the action to be performed on the Guest VM.
- **CPU Utilization:** The CPU utilization of the virtual server.

- **Disk Utilization:** The disk utilization of the virtual server.
- **Network Activity:** The network throughput of the virtual server.
- **Incident Summary:** This region shows a list of incidents and problems reported on all the Guest VMs and related targets.

7.7.2 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 sever that is to be edited. The Virtual Server Home page appears.
3. From the **Virtual Server** menu, select **Target Setup**, then select **Edit Virtual Server**.
4. In the Edit Virtual Server page, you can modify the Description of the virtual server.
5. Select the **Configure 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.
6. Click **OK** to confirm the changes.

7.7.3 Upgrading a Virtual Server

After the YUM Repository has been set up (see [Section 5.15.1, "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.

7.7.4 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 7.7.2, "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. Login 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.

7.7.5 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.

7.8 Managing Guest Virtual Machines

This section describes the following:

- [Creating a Guest Virtual Machine](#)
- [Guest Virtual Machine Home Page](#)
- [Live 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](#)

7.8.1 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 7.9, "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.

7.8.2 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.
- **Job Activity:** Shows the list of jobs started within the last 7 days.
- **CPU and Memory Utilization:** The CPU and memory utilization of the guest virtual machine.
- **Disk Utilization:** The disk 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.
- **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
- Live Migrate a Guest VM
- Clone a Guest VM
- Start and Stop
- Reboot, Restart, Suspend, and Resume

7.8.3 Live Migrating a Guest Virtual Machine

To live 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.
4. Right click the selection and select **Live Migrate**. The guest virtual machines to be migrated must be in the Running state.
5. Check the **Manual** check box to select a virtual server to which the guest virtual machine is to be migrated. If you do not select this check box, the system will pick up a suitable virtual server from the server pool for the guest virtual machine live migration.
6. Click **Submit** to live migrate the guest virtual machine.

7.8.4 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](#)

7.8.4.1 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.

7.8.4.2 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**.

7.8.4.3 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**.

7.8.5 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, 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 following:
 - Destination target on which the guest virtual machine is to be deployed.
 - Source guest virtual machine that it is to be cloned.
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.
 - **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.

- **NIC Configuration:** You can add or delete virtual network interfaces. For more details, see [Section 5.6, "Setting Up Networks"](#).
 - **Disk Configuration:** 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.
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.

7.8.6 Editing a Guest Virtual Machine

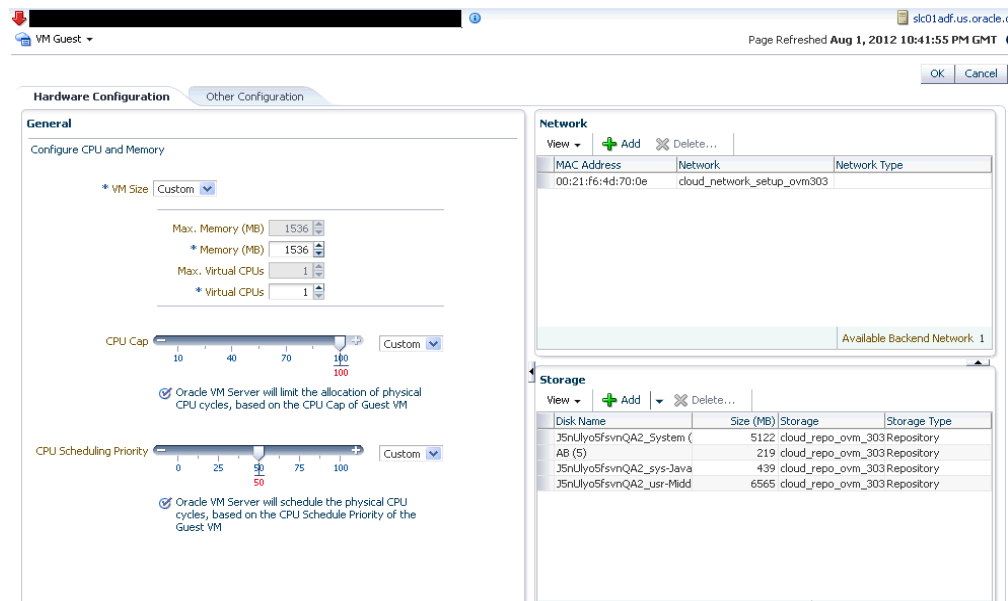
Note:

- 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 7–6 Edit Guest Virtual Machine



3. You can modify the following details:

■ Hardware Configuration (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.

■ Hardware Configuration (Network):

You can add delete virtual network interfaces. For more details, see [Section 5.6, "Setting Up Networks"](#).

■ Hardware Configuration (Storage):

- Select the desired storage configuration of your virtual machine, such as virtual disks and ISO file.

■ Other Configuration (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.

- **Other Configuration (High Availability):** If the **High Availability** checkbox is selected, the guest virtual machine will automatically restart after any outage.
- **Other Configuration (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, CDRROM, or network.
- **Other Configuration (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.

7.8.7 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.
- **NIC:**
 - You can add delete virtual network interfaces. For more details, see [Section 5.6, "Setting Up Networks"](#).
 - **Guest VM Naming:** Indicate whether the guest virtual machine names are to be automatically generated or if the names will be uploaded from a file.
- **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.

7.9 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](#)

7.9.1 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 7-1 Getting Started with Oracle VM Provisioning

Step	Description	Reference Link
Step 1	<p>Understanding the Deployment Options</p> <p>Understand the deployment plans offered by Enterprise Manager for deploying virtualization systems.</p>	To learn about the deployment options, see Section 7.9.2, "Deployment Options" .
Step 2	<p>Meeting the Prerequisites</p> <p>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.</p>	To learn about the prerequisites for provisioning guest virtual machines, see Section 7.9.3, "Prerequisites" .
Step 3	<p>Selecting the Use Case</p> <p>This section covers a few use cases for the provisioning of guest virtual machines. Select the use case that best matches your requirements.</p>	<p>To provision a guest virtual machine using:</p> <ul style="list-style-type: none"> ■ Assemblies: See Section 7.9.4, "Provisioning Guest Virtual Machines Using Oracle Virtual Assemblies (OVA)" ■ Templates: See Section 7.9.5, "Provisioning a Guest Virtual Machine Using Oracle VM Templates" ■ ISO Images: See Section 7.9.6, "Provisioning a Guest Virtual Machine Using an ISO Image" ■ PXE Boot: See Section 7.9.7, "Creating PXE Bootable Guest Virtual Machines" ■ Automated Deployment on Bare Metal Boxes: See Section 7.10, "Bare Metal Provisioning"

7.9.2 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.
- **Bare Metal Provisioning:** Allows you to provision Oracle VM server software on bare metal machines.

7.9.3 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 Management Administrator's Guide*.
- Optional Infrastructure Requirements as defined in the *Enterprise Manager Lifecycle Management 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.

7.9.3.1 Creating the Software Library Components

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](#)

7.9.3.1.1 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`

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.

7.9.3.1.2 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 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 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`

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.

7.9.3.1.3 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`

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.

7.9.3.1.4 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 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 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`

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.

7.9.4 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 Virtual Server Pool target and select **Deploy Assembly**.

Figure 7-7 Assembly Deployment: Deployment Configuration

Node Name	Start Order	Deploy
	0	<input checked="" type="checkbox"/>

3. On the Assembly Deployment: Deployment Configuration 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 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 7–8 Assembly Deployment: Network and Agent

The screenshot shows the 'Assembly Deployment: Network and Agent' configuration page. At the top, there are navigation tabs: 'Deployment Configuration', 'Network and Agent', 'Deployment Configuration', 'Schedule', and 'Review'. Below the tabs, there are buttons: 'Review', 'Save As Plan...', 'Back', 'Step 2 of 5', 'Next', and 'Cancel'. The main content area is divided into two sections.

The first section is titled 'Network'. It contains a table with the following columns: 'Assembly Network', 'Type', and 'Network'. The table has one row with the following values: 'network-1', 'Any Network', and 'System Assigned'.

The second section is titled 'Configure EM Agent Push'. It contains a checkbox labeled 'Enable EM Agent Push Configuration' which is checked. Below this, there is a note: '> 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.' Below the note is a table titled 'Assembly Node Selection for EM Agent Push Selection' with the following columns: 'Node Name', 'Installation Type', and 'Enable Agent Push'. The table has one row with the following values: a redacted node name, 'Push Agent Installation', and a checked checkbox.

Below the table, there are several input fields: 'Installation Base Directory' (with value '/home/oracle/agent'), 'Platform' (with value 'Linux x86-64'), 'Username' (with value 'oracle'), 'Password' (with value '*****'), 'Verify Password' (with value '*****'), 'Port' (with value '3872'), and 'Additional Parameters'.

5. Click on the **Browse** icon in the Network field. Specify if the network should be allocated by the system or selected manually.

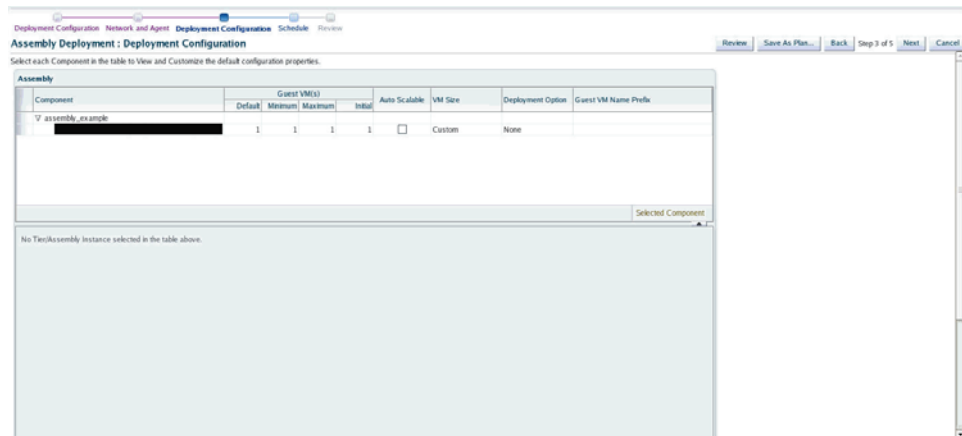
6. If you want to install the Management Agent on an assembly tier (collection of Guest VMs) or selected guest virtual machines, check the **Enable EM Agent Push Configuration** check box. You can then select the tiers and / or the guest virtual machines on which the Management Agent is to be installed using the Agent Push approach in the Assembly Node Selection for 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.
 - **Additional Parameters:** Specify additional parameters to be used during the installation of the Management Agent. For example, enter `-ignoreDirPrereq` to ignore the prerequisite checks during installation.
 8. Click **Next** to go to the Deploy Assembly: Deployment Configuration page.

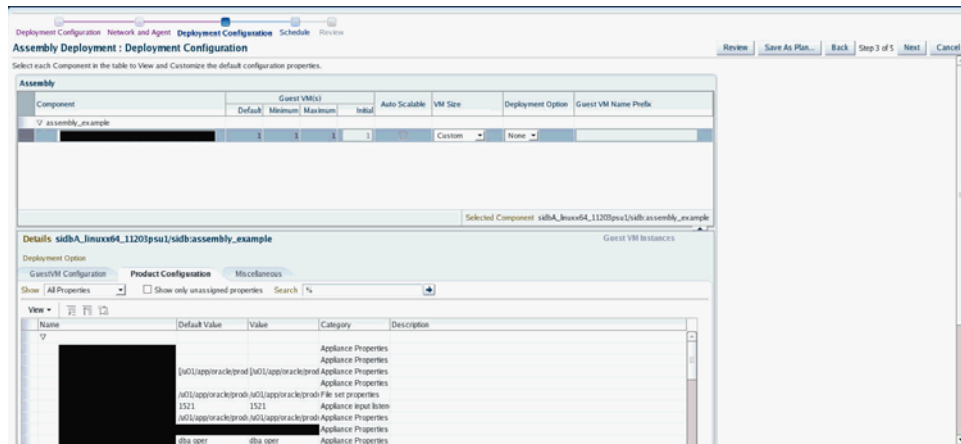
Figure 7–9 Assembly Deployment: Deployment Configuration



The following details are displayed:

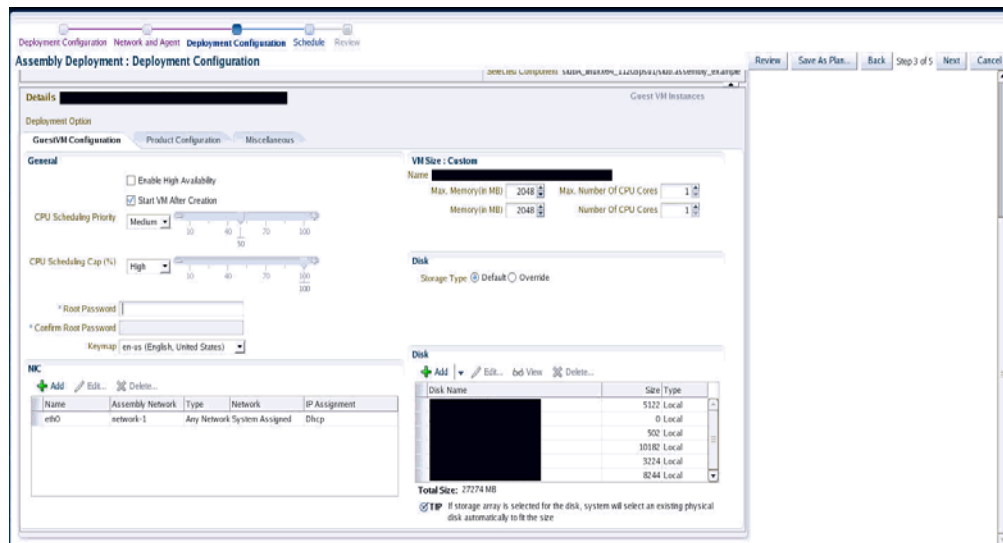
- **Component:** The assembly component that is being deployed.
 - **Guest VMs:** The default number of guest virtual machines as specified in the deployment profile appears. Specify the minimum and maximum number of guest virtual machines that are to be created.
 - **Auto Scalable:** If you select the Auto Scalable check box, it guarantees the availability of the guest virtual machines if the virtual server they are running on fails or restarts. When a virtual server is restarted or shut down, the guest virtual machines running on it are either restarted on, or migrated to, another virtual server.
Note: You can select the check box only if high availability has been enabled for the server pool during creation.
 - **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 7–10 Assembly Deployment: Deployment Configuration: Product Configuration



- **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 7–11 Tier: Deployment Configuration: GuestVM Configuration



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. 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.
- **Root Password:** Specify the operating system root password to log in to the guest virtual machine once it is up and running.
- **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.

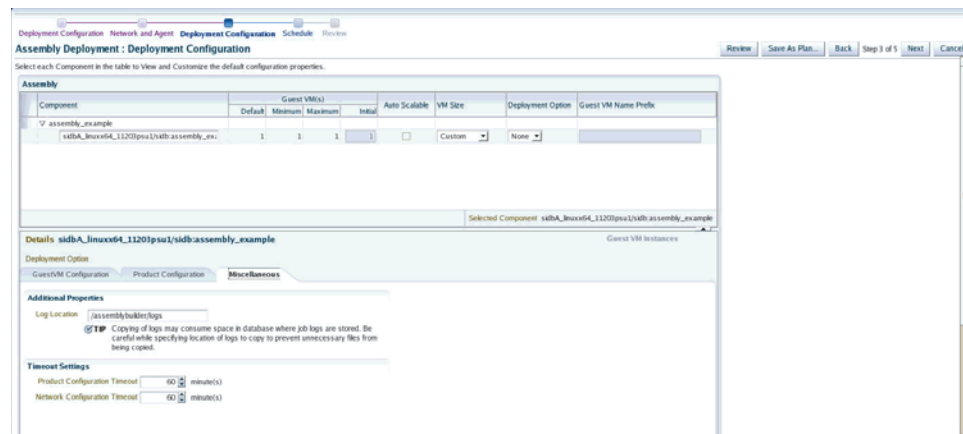
Clone Target: If you select Repository, select the storage location where the disk is to be created. If you select Storage Array, specify the type, which can be SAN or iSCSI storage server.

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.

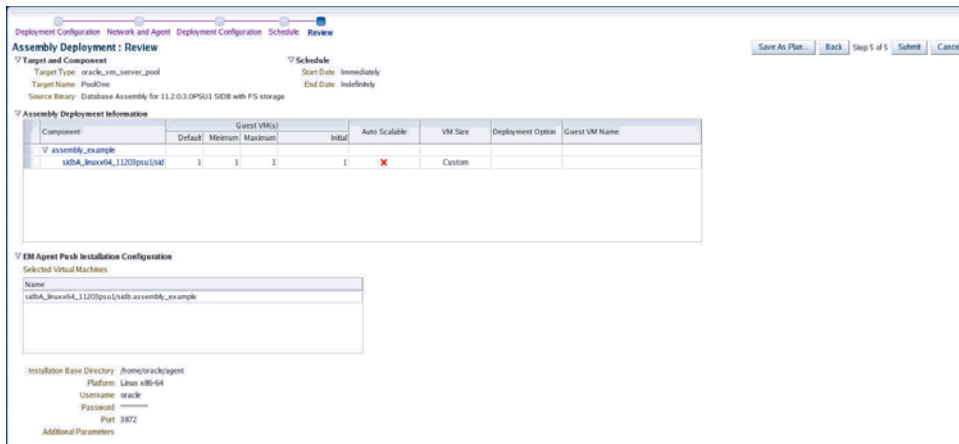
- **Disk:** Every guest virtual machine is created with at least one disk. You can specify additional disks and their sizes. These disks will be added to the guest virtual machine. After the guest virtual machine is created, log in to the guest virtual machine and reformat and mount the additional disks. 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.

Figure 7–13 Tier Deployment Configuration: Miscellaneous



- **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 7–14 Assembly Deployment: Review

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 plan and complete the process.

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.

7.9.5 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. These templates are available in the .tgz file format. Guest virtual machines are created when these templates are deployed, but 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.

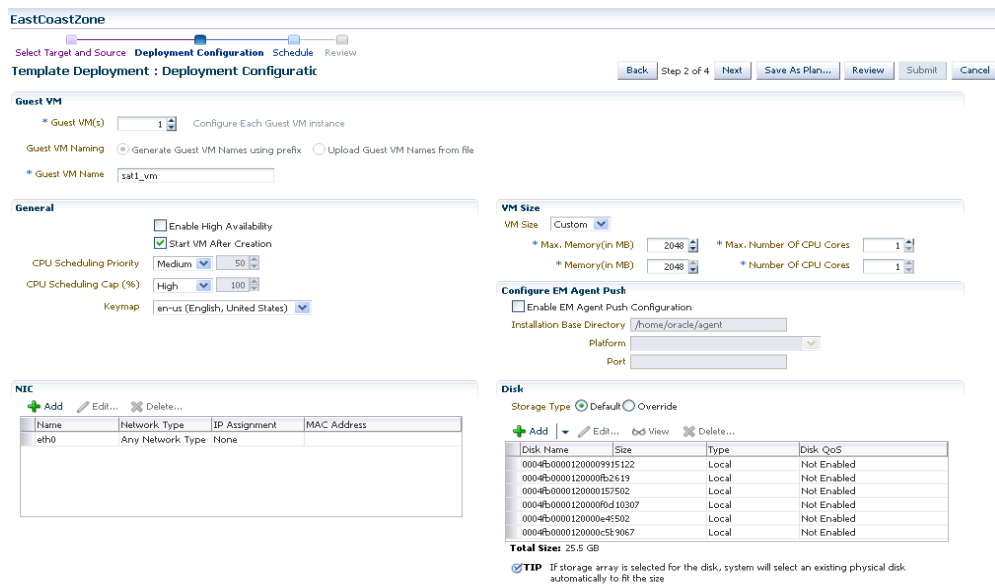
Follow this process to deploy a template:

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 7–15 Template Deployment - Select Target and Source

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 7–16 Template Deployment - Deployment Configuration



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 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 (VCPUs) 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.
 - **Configure EM Agent**: If you want to configure the Management Agent to monitor the guest virtual machine, check the Enable EM Agent Configuration check box. Specify 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.
 - **Port**: The port on which the Management Agent will communicate.
 - **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 plan and complete the process.

7.9.6 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 7–17 ISO Deployment: Select Target and Source

The screenshot displays the 'Pool1' configuration page in Oracle Enterprise Manager. The page title is 'Pool1' and the breadcrumb is 'Select Target and Source > Deployment Configuration > Schedule > Review'. The main heading is 'Iso Deployment : Select Target and Source'. The page contains three main sections:

- Destination:** A section for selecting the destination target. It shows 'Target: Pool1' and 'Target Type: Oracle VM Server Pool'.
- Source:** A section for selecting the source image. It includes a 'Use Deployment Plan' checkbox, a 'Deployment Plan' input field, an '* ISO' input field, and a 'Description' input field.
- Request Details:** A section for entering request information. It shows a '* Request Name' field with the value 'ISO_DEPLOYMENT_2011-09-25_09-48-19-689'.

 Navigation buttons at the top right include 'Back', 'Step 1 of 4', 'Next', 'Save As Plan...', 'Review', 'Submit', and 'Cancel'.

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.
- **Configure EM Agent:** If you want to configure the Management Agent to monitor the guest virtual machine, check the Enable EM Agent Configuration check box. Specify 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.
 - **Port:** The port on which the Management Agent will communicate.
 - **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.

7.9.7 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.

7.10 Bare Metal Provisioning

Bare metal provisioning involves automated discovery and deployment of the Oracle VM Server on bare metal boxes. It also allows the virtual management servers to be automatically discovered and registered to the virtual management server pool. You can create customized Oracle virtual server images and automatically deploy the Oracle VM Server software on multiple servers based on their MAC or subnet addresses. You can automatically register the provisioned Oracle VM Servers to the specified server pool.

Oracle VM Server image is modeled as a software library component that captures properties like RPM repository, disk layout, packages to be installed, operating system password, Oracle VM Agent password, NFS information, NIS information, NTP information, and so on.

During deployment, you can use a pre-created Oracle VM Server software image or specify all the required properties in the deployment flow itself. You can provision Oracle VM Server software on the following types of targets:

- **MAC Address:** You can specify a list of MAC addresses of bare metal boxes to provisioned.
- **Subnet:** You can specify a subnet for which the Oracle VM server software has to be staged.
- **Enterprise Manager Host Target:** You can specify existing host targets and reprovision them with Oracle VM Server software images.

Follow this procedure to provision virtual management servers and automatically add them to the server pool.

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select Bare Metal Provisioning.
2. The Bare Metal Provisioning Dashboard appears. Click the **Deployments** tab.
3. From the Provisioning menu, select **Oracle VM Server**. The Provision Oracle VM Server: Target Selection page appears.
4. Enter a name and description for the deployment procedure.
5. If you select MAC Addresses, click **Add** and add the MAC addresses of the servers that need to be provisioned. Click **Add from File** to add MAC addresses stored in a file.
6. If you select Subnet, specify the Subnet IP, Netmask, Number of Network Interfaces, and Bootable Network Interface.
7. Select the OVM Manager target that will be used to manage the Oracle VM Servers.

8. In the Provision Oracle VM Server: Deployment page, select the following:
 - **Stage Server:** During provisioning of an image on hardware servers, the required binaries and files are first transferred to a stage server. This is known as Staging phase and is responsible for preparing images to be installed over the network, and exposing installable or executable software elements over the network to the target hardware server being provisioned.

The Provisioning application requires at least one stage server on which all the activities related to staging can be performed. From the networking perspective, you are advised to keep the stage server as close to the target machines as possible. The Stage Server must be an Enterprise Manager target. If you select the **Run Stage Server Pre-requisites Checks** check box, you can ensure that the Stage Server meets requirements such as large storage, high memory, bandwidth and so on.
 - **Boot Server:** One of the key requirements of the application is the ability of the hardware server to boot up over the network (rather than from a local boot device). A boot server must be set up so that it is able to service the requests from the designated hardware servers in order for them to boot over the network. The Boot Server must be an Enterprise Manager target and should be able to receive the BOOTP and Trivial File Transfer Protocol (TFTP) requests over the network from the hardware server. Select the Boot Server and check the **Run Boot Server Pre-requisites Checks** to ensure that the Boot Server requirements are met.
 - **Local RPM Repository:** The Provisioning application picks up the RPMs for the operating system from the RPM repository. At least one repository needs to be setup for use by the Provisioning application.
 - **Fetch Configuration Properties From Pre-Created Components:** In this section, you can select preconfigured components from the Software Library.
 - **OS Component:** You can select a preconfigured OS Component which provides operating system details for the machines you are provisioning.
 - **Disk Layout Component:** You can select a preconfigured disk layout component to provide profiles for hard disk, physical disk partition, RAID configuration, logical volumes, and logical volume groups for the machines you are provisioning.
 - **Provisioning Directive:** Provisioning Directives are entities in the Software Library that represent a set of instructions to be performed. These are constructs used to associate scripts with software components and images. These scripts contain directions on how to interpret and process the contents of a particular component or an image.
9. Click **Next** to continue to Provisioning VM Server: Basic OS Details page. Set the Time Zone, OS Root Password, and OVM Agent Password. In the Add Operating System Users list section, click Add. Specify the User Name, Password, Primary Group, and Additional Groups to add the operating system users. Enable or Disable sudo access. Click **Next**.
10. The Provision Oracle VM Server: Additional OS Details page contains the following sections:
 - **Dom0 Configuration:**
 - **Dom0 Memory:** Specify the memory size for the virtual server target.
 - **Additional Configuration and Power Interface:** Specify the advanced power configuration for the virtual server target.

- **Additional OS Details**
 - **Mount Point Settings:** Specify entries for the `/etc/fstab` file. You can specify mount points on the newly provisioned Linux machine. By default, mount point settings from the reference Linux machine are inherited.
 - **NIS Settings:** Specify entries for the `/etc/yp.conf` file. You can specify NIS settings for the newly provisioned Linux machine. By default, NIS settings from the reference Linux machine are inherited.
 - **NTP Settings:** Specify entries for the `/etc/ntp.conf` file. You can specify NTP settings for the newly provisioned Linux machine. By default, NTP settings from the reference Linux machine are inherited.
 - **Kernel Parameter Settings:** You can add a script in the Kernel Parameters field.
 - **Inittab Settings:** You can add a script to describe how the `init` process should set up the system at run-time.
 - **Firewall Settings:** Specify firewall settings for the Linux target. Firewall settings are disabled by default and can be configured. Ensure that the port used by the Management Agent is open for its communication with the Management Service. By default, the Management Agent uses port 3872 or a port number in the range 1830-1849, unless configured to use some other port.
 - **Boot Configurations**
 - **Post Install:** Specify any set of commands that need to be executed on the newly provisioned machine. These commands will be appended to the post section of the kickstart file.
 - **First Boot:** Enter the commands to be run on the system when it boots up for the first time after installation.
11. Click **Next** to continue to the Provision Oracle VM Server: Disk Layout page. In the Disk Layout Configuration page, specify hard disk profile, partition configuration, RAID configuration, Logical Volume Group configuration, and Logical Volume configuration.
- **Hard Disk Profile:** Click **Add** and specify the **Device Name** and **Capacity**.
 - **Logical Volume Group Configuration:** Click **Add** and specify the **Group Name**, **Partitions**, and **RAIDs**.
 - **Partition Configuration:** Click **Add** and specify the **Mount Point**, **Device Name**, **File System Type**, and **Size (MB)**.
 - **Logical Volume Configuration:** Click **Add** and specify the **Mount Point**, **Logical Volume Name**, **Logical Group Name**, **File System Type**, and **Size (MB)**.
- Click **Next**.
12. In the Network Configuration page, the list of MAC addresses to be provisioned and the network interface components are displayed. You can edit or retain these values. Click **Add** to add a network interface.
13. Click **Next**. In the Schedule page, the Job Name for the Oracle VM Server Bare Metal Provisioning deployment appears. Provide a schedule for the job, either immediately or at a later date. Click **Next**.

14. In the Review page, verify that the details you have selected are correctly displayed and submit the job for the deployment. Click **Submit Job**.
15. In the Jobs page, view the job steps and verify that Status is Success. If the status is Failed, view the steps that have failed, and fix them and resubmit the job.
16. Start Provisioning and plug in the bare metal machines when required to add Oracle VM Servers to the Server Pool.

Managing Cloud Policies

This chapter describes the process to define policies for 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](#)
- [Defining Policy Groups](#)
- [Managing DRS and DPM Policies](#)

8.1 Managing Cloud Policies

Cloud management policies are rules that a `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.

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 Based Policy
- Activate or Deactivate a Policy
- Define Policy Groups

8.2 Viewing Cloud Policies

You can view all the cloud policies 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.

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Policies**.
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 8–1 Cloud Policies Page

The screenshot shows the 'Cloud Policies' page in Oracle Enterprise Manager. At the top, there is a navigation bar with 'Cloud Policies' and a refresh button. Below the navigation bar, there is a toolbar with options like 'Create', 'View', 'Create Like...', 'Edit...', 'Delete', 'Activate', and 'Deactivate'. A search box is also present. The main area contains a table of policies with columns: Name, Type, Associations, Target Type, Status, Owner, Date Created, and Last Modified. The table lists several policies, including 'Test Performance Policy 2', 'test', 'testew', 'yerte', 'esyte', 'hjh', 'twetwe', and 'guest_vm_policy'. Below the table, there is a detailed view for 'Test Performance Policy 2'. This view is divided into two sections: 'General' and 'Associations'. The 'General' section shows details such as Type (Performance), Associations (0), Description (Test performance policy description), Status (Active), Owner (EMCLOUD_ADMIN), Date Created (7/23/2012), Last Modified (EMCLOUD_ADMIN), and Target Type (Host). The 'Associations' section shows a table with columns 'Target', 'Created By', and 'Last Updated On', and a message stating 'There are no associations for the selected policy.'

3. Select a policy from the table. The policy details like the type, associations, and status are listed.
4. You can view the associations for a selected policy. A policy can be associated with one or more targets. Click **Add Association** to associate a target with the policy. You can:
 - Click a target in the Associations region to drill down to the policies page for the target.
 - Click **View History** to view the history of policy associations for this target.

8.3 Viewing Policies for a Target

You can view the policies associated with a target on this page. You can associate or disassociate policies with a target.

- Click **Associate** to associate it with the target.
- Select a policy from the list and click **Disassociate** to delete the policy association.
- The History section shows the history of the policy associations performed on this target.

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 History section shows the history of associations for this target.

8.4 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.

Prerequisites

- User has administrator privileges.

To define a performance policy, follow these steps:

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Policies**.
2. On the Policy Home page, you will see a list of a policies that you own and for which you have administration privileges.
3. Click **New Policy** and select the **Performance** Policy type.
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 8–2 Create Performance Policy - General

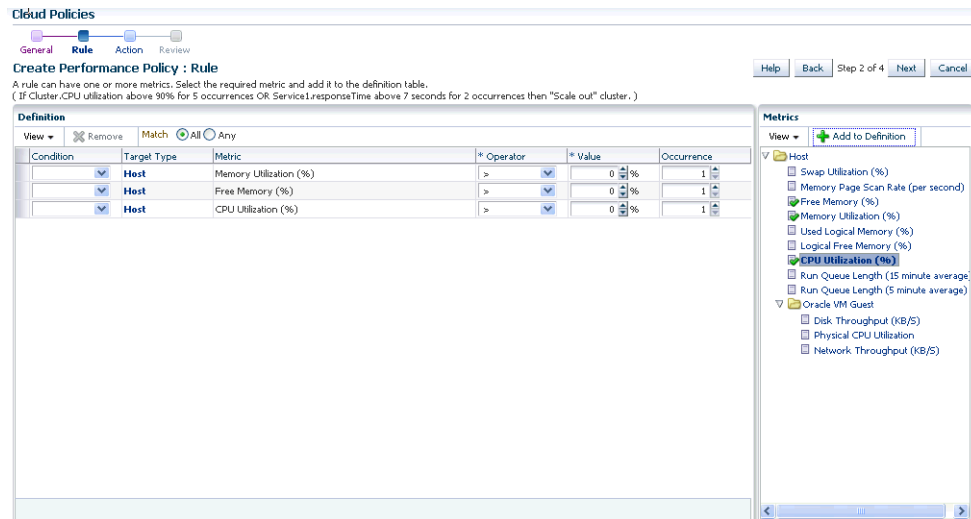
The screenshot shows the 'Create Performance Policy : General' form. At the top, there are tabs for 'General', 'Rule', 'Action', and 'Review'. The 'General' tab is active. The form contains the following fields:

- Name:** performance_policy
- Description:** (empty text area)
- Type:** Performance
- Target Type:** Host (dropdown menu)

At the top right of the form, there are navigation buttons: Help, Back, Step 1 of 4, Next, and Cancel.


5. Click **Next**. In the Rule page, select a metric from the left panel and click **Add to Condition**. Specify the condition, value, and number of occurrences. When these conditions are met, the policy is executed.

Figure 8–3 Create Performance Policy - Rule



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.

Figure 8–4 Create Performance Policy - Action



7. 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.

8.5 Defining a Schedule Based Policy

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

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Policies**.
2. On the Policy Home page, you will see a list of a policies that you own and for which you have administration privileges.
3. Click **New Policy** and select the **Schedule Policy** type.

Figure 8–5 Create Schedule Policy

The screenshot shows the 'Create Schedule Policy' dialog box. It includes the following fields and options:

- Name:** A text input field.
- Description:** A larger text area.
- Type:** A dropdown menu set to 'Schedule'.
- Target Type:** A dropdown menu set to 'Oracle VM Guest'.
- Policy Action:** A dropdown menu set to 'Start'.
- Start:** A date and time picker set to '6/2/2011 4:26:55 PM' with a UTC+00:00 indicator.
- Repeat:** A dropdown menu set to 'Every N minutes'.
- Frequency:** A dropdown menu set to 'Every 1 minute(s)'.
- Grace Period:** A checkbox labeled 'Do not run if it cannot start within' followed by a spinner set to '1' and a dropdown set to 'hours'.
- Duration:** Radio buttons for 'Indefinitely' (selected) and 'For 1 hours'.

4. 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.
5. Specify the action that is to be performed when the policy conditions are met. The list of available actions can vary depending on the target type with which the policy is associated.
6. Specify the frequency and the schedule for the policy.
7. Click **Save** to define this policy.

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

8.5.1 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 8.5.1.1, "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 8.5.1.2, "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 8.5.1.3, "Creating a Policy Group \(Start Up and Shut Down\)"](#).

8.5.1.1 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. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Policies**.
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 8–6 Create Schedule Policy (Example 1)

The screenshot shows the 'Create Schedule Policy' dialog box. The 'Name' field contains 'guestvm_policy1'. The 'Description' field is empty. The 'Type' is set to 'Schedule'. The 'Target Type' is 'Oracle VM Guest'. The 'Policy Action' is 'Stop'. The 'Description' for the action is 'Stop a Guest Virtual Machine'. The 'Start' time is '8/10/2012 6:00:00 PM' in GMT. The 'Repeat' is set to 'Specific days of the week'. Under 'Days of Week', the 'Friday' checkbox is checked. The 'Grace Period' is set to '1' hour.

3. Click **OK** to save this policy and return to the Cloud Policies page.
4. Click on the policy that you have created to view the details of the policy. In the **Associations** region, click **Add**.
5. In the Search window, select the Oracle VM Guest checkbox and click **Search** to view a list of guest virtual machines.
6. Select the virtual guest machine with which the policy should be associated and click **Select**.
7. The policy is now associated with the selected guest virtual machine and the policy conditions will be applied.

8.5.1.2 Policy for Starting Up Guest Virtual Machine

To restart 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.
4. Click on the policy that you have created to view the details of the policy. In the **Associations** region, click **Add**.

5. In the Search window, select the Oracle VM Guest checkbox and click **Search** to view a list of guest virtual machines.
6. Select the virtual guest machine with which the policy should be associated and click **Select**. The Guest VMs you select must be the same as the ones selected in the earlier example.
7. The policy is now associated with the selected guest virtual machine and the policy conditions will be applied. You can click on the associated target to view details about the policy.

8.5.1.3 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. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Policies**.
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.

8.6 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. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Policies**.
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.

8.7 Defining Policy Groups

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

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Policies**.
2. Click on the **Policy Groups** tab. On the Policy Groups page, you will see a list of a policies that have been defined.
3. Click **Create**. In the New Policy Group page, click **Select and Add Policy** to add one or more policies to the group.
4. A list of policies are displayed. Select one or more policies from the list and click **Add**.
5. The selected policy will be added to the policy group. Based on the schedule details defined for each policy, the policy action will be scheduled.
6. Click **Save** to save this policy group.

You can perform the following operations:

- **Edit Group:** In the **Policy Group** page, select a policy group from the list and click **Edit Group**. The policies included in the group and the associations are displayed. You can add or delete policies and associations to the group.
- **Delete Group:** In the **Policy Group** page, select a policy group from the list and click **Delete Group**. The policy group and its associations will be deleted.

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

8.7.1 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 based on certain metric values. You can create the following policies:

- First, create a performance policy to scale down the tier instance to 1. See [Section 8.7.1.1, "Policy to Scale Down a Tier Instance"](#).
- Scale up the tier instance to 2 every Monday. See [Section 8.7.1.2, "Policy to Scale Up a Tier Instance"](#).
- Combine these policies and create a Policy Group. See [Section 8.7.1.3, "Creating a Policy Group \(Scale Up and Scale Down\)"](#).

8.7.1.1 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. From the Enterprise menu, select Cloud, then select Infrastructure Policies.
2. In the Cloud Policies page, select Performance Policy from the Create menu.
3. In the Create Performance Policy: General page, enter a name and description for the policy and select the Target Type as Oracle Tier Instance.

Figure 8–7 Performance Policy: General Page (Example 2)

Cloud Policies

General Rule Action Review

Create Performance Policy : General

Back Step 1 of 4 Next Cancel

* Name perf_policy1

Description Policy to scale down a tier instance

Type Performance

* Target Type Oracle Tier Instance

4. Click **Next**. In the Create Performance Policy: Rule page, add the following metrics to the Definition table.
 - **Physical CPU Utilization**: Specify the **Condition** as **Average of, Value** greater than **90, Occurrences** as **5**.
 - **Total Throughput**: Specify the **Condition** as **Average of, Value** greater than **75, Occurrences** as **3**.

Select **Match Any** option to indicate that the policy is applicable if any of the conditions are met.

Figure 8–8 Performance Policy: Rule Page (Example 2)

Cloud Policies

General Rule Action Review

Create Performance Policy : Rule

Back Step 2 of 4 Next Cancel

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 Service.responseTime above 7 seconds for 2 occurrences then "Scale out" cluster.)

Definition

View Remove Match All Any

Condition	Target Type	Metric	Value
Average of	Oracle VM Guest	OVM Guest Load.Physical CPU Utilization	> 90
Average of	Oracle VM Guest	OVM Guest Network Activity Summary.Total Throughput	> 75

Metrics

View Add to Definition

- Oracle Tier Instance
 - Oracle VM Guest
 - OVM Guest Load.Physical CPU Utilization
 - OVM Guest DiskActivity Summary.Total Throughput (kb/s)
 - OVM Guest Network Activity Summary.Total Throughput**

5. Click **Next**. In the Performance Policy: Action page, enter the following details:
 - In the Applies to Target field, select **Oracle Tier Instance**.
 - In the Action field, select **Scale Down**.
 - In the Action Parameters field, enter the Number of Nodes to be scaled down as **1**.

Figure 8–9 Performance Policy: Action Page (Example 2)

Cloud Policies

General Rule Action Review

Create Performance Policy : Action

Back Step 3 of 4 Next Cancel

Rule Average of Oracle Tier Instance/OVM Guest/OVM Guest Load.Physical CPU Utilization is greater than 90 for 1 consecutive occurrence(s) or Average of Oracle Tier Instance/OVM Guest/OVM Guest Network Activity Summary.Total Throughput (kb/s) is greater than 75 for 1 consecutive occurrence(s)

* Applies to Target Oracle Tier Instance

Action Scale Down

Description

Action Parameters

Parameter	Value
Number of Nodes	1

6. Click **Next**. Review the details of the policy and click **Save**.

8.7.1.2 Policy to Scale Up a Tier Instance

To scale up a tier instance, navigate to the Cloud Policy page and select the Performance Policy that you just created. Follow these steps:

1. Click **Create Like** to make a copy of the policy.
2. Modify the name and description of the policy. Click **Next**.
3. In the Performance Policy: Rule page, do the following:
 - **Physical CPU Utilization:** Specify the **Condition** as **Average of, Value** greater than **75**, **Occurrences** as **5**.
 - **Total Throughput:** Specify the **Condition** as **Average of, Value** greater than **50**, **Occurrences** as **3**.Select **Match Any** option to indicate that the policy is applicable if any of the conditions are met.
4. In the Performance Policy: Action page, modify the Action to **Scaleup** and in the Action Parameters field, specify the **Number of Nodes** as **2**.
5. Click **Next**. Review the policy and click **Save**.

8.7.1.3 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. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Policies**.
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 Oracle Tier Instance targets associated with this policy. Click **Add** to associate more Oracle Tier Instance targets with this policy group.
10. When this policy group is active, the Oracle Tier Instance targets are scaled down and scaled up based on the policy definition.

8.8 Managing DRS and DPM Policies

Apart from managing and defining policies for a zone, you can also manage Dynamic Resource Scheduling (DRS) and Distributed Power Management (DPM) policies for a virtual server pool. These automated policies 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. 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 8–10 Configure DRS/DPM: General Page

The screenshot shows the 'Configure DRS/DPM: General' page. At the top, there are tabs for 'General' and 'Network', with 'General' selected. Below the tabs, the title is 'Configure DRS/DPM: General' and 'Policy Control'. There are navigation buttons: 'Back', 'Step 1 of 2', 'Next', and 'Cancel'. The 'Policy Type' is set to 'Distributed Resource Scheduler' and the 'Time Period (Minutes)' is set to 10. Below this is a section for 'Virtual Servers' with a table for adding or removing targets. At the bottom is a section for 'Constraints' for 'Virtual Server CPU', with 'Enable' checked and 'Threshold (%)' 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 8–11 Configure DRS/DPM Policies: Network Page



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 8.5, "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.

8.8.1 Creating Automated DRS and DPM Policies

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

1. From the **Enterprise** menu, select **Cloud**, then select **Infrastructure Policies**.
2. Click **New Policy** and select the **Schedule Policy** type.
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.

Part III

Platform as a Service

This section describes the steps required to enable platform as a service.

It contains the following chapters:

- [Section 9, "About Platform as a Service"](#)
- [Section 10, "Using the Self Service Framework"](#)

About Platform as a Service

This chapter provides an introduction to platform as a service (PaaS). It contains the following sections:

- [About Platform as a Service \(PaaS\)](#)
- [Platform as a Service with Enterprise Manager](#)

9.1 About Platform as a Service (PaaS)

Enterprises today must deliver new services and constantly innovate to stay ahead of their competitors. However, enterprises face several challenges in delivering quality service on schedule. Provisioning infrastructure for new projects takes days and weeks, instead of hours. This severely limits enterprise agility. Users, such as Developers or QA Engineers, need rapid access to development platforms, without intervention by IT administrators.

Oracle Cloud Management Pack for Middleware offers all capabilities necessary for deploying and managing middleware-centric Platform as a Service (PaaS) clouds within the enterprise. PaaS enables enterprise users to acquire resources and deploy J2EE applications quickly, while focusing on application development, instead of infrastructure administration. Users can directly access applications without IT intervention and achieve massive enterprise agility.

9.2 Platform as a Service with Enterprise Manager

Enterprise Manager provides the most comprehensive solution for rolling out a middleware-centric PaaS Cloud for users in an enterprise. It allows users to build, deploy and manage the entire lifecycle of the cloud from a single console. The pre-integrated solution ensures that IT administrators can quickly deliver value to users, while retaining control of the environment. Oracle platforms ensure developers can leverage the features of the PaaS Cloud without having to learn new non-standard technologies, ensuring continuity of skills for the enterprise.

Enterprise Manager provides the following features:

- **Self Service Provisioning of J2EE Applications:** Enterprise Manager enables IT to deliver WebLogic-based application deployment environments as Cloud services. Administrators can define different types of service (based on sizing or other configuration characteristics of the WebLogic runtime) to meet their enterprise standards. The Self Service Portal can be used to request these services on demand, deploy applications to them, and manage the lifecycle of each of their service instances. Users can monitor their services instance and the J2EE applications deployed to them for performance and availability.

- **Self Service Provisioning of Multi Tier Applications:** Using the Oracle Virtual Assembly Builder, Enterprise Manager, allows IT to package complex, enterprise platforms and applications into Assemblies and deliver them as Cloud services. Oracle Virtual Assembly Builder allows administrators to quickly configure and provision entire multi-tier enterprise applications Cloud environments.

Administrators can create different assemblies, each capturing enterprise best practices for deployment, and enable rapid self-service provisioning of these platforms. Developers can use the Self Service Portal to deploy these assemblies on demand, manage the lifecycle of these instances, as well as monitor the various tiers.
- **Elastic Resource Provisioning:** Enterprise Manager enables users to scale up or scale down resources attached to an application in an elastic fashion. For Java services, developers can dynamically scale up or down the number managed servers available within a WebLogic service by way of the Self Service Portal.
- **Metering and Chargeback:** Enterprise Manager provides tools to define detailed charge plans spanning different metrics collected for physical and virtual resources. Enterprise Manager supports metering for all of the resources necessary for an Oracle-based PaaS cloud, including hosts, virtual machines, database, and WebLogic Server instances. A hierarchy of cost centers can be created or imported from the LDAP for cost allocation and reporting aggregation.
- **Consolidation Planner:** IT administrators can use Consolidation Planner to run different scenarios for redistributing workloads onto existing systems or new environments, and determine if this will result in SLA violations. Based on these scenarios, an administrator can select the best approach for establishing the infrastructure to support the Cloud.
- **Cloud Administration and Monitoring:** Enterprise Manager provides end-to-end monitoring and diagnostic capabilities for identifying bottlenecks in the cloud infrastructure and for taking remedial actions against such bottlenecks. End User monitoring¹ (real and synthetic), Business Transaction Management, Java and Database monitoring and Diagnostics², and system monitoring agents are deployed into the cloud infrastructure and provide complete performance and availability monitoring of the business application, individual infrastructure, and the software components on which the application is deployed.

Using the Self Service Framework

This chapter covers the common self service operations that need to be performed before setting up DBaaS or MWaaS. It also describes the monitoring operations that can be performed on DBaaS or MWaaS. It contains the following sections:

- [Creating a PaaS Infrastructure Zone](#)
- [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](#)
- [PaaS Infrastructure Zone Home Page](#)

10.1 Creating a PaaS Infrastructure Zone

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 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.

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: 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 example, if you have a Cloud Plug-in installed, and are using the Cloud features offered in Enterprise Manager, then ensure that the `CLOUD_ENGINE_USER` is also granted view privileges on credentials associated with Software Library. Since the `CLOUD_ENGINE_USER` is a hidden user account, the owner of the named credential will not be able to grant him View privileges from the Enterprise Manager UI. To handle this situation, (especially on a Windows host where OMS Agent Filesystem is the recommended approach for setting up Software Library) you can run the following EMCLI commands:

```
emcli login -username=<username -password =<>

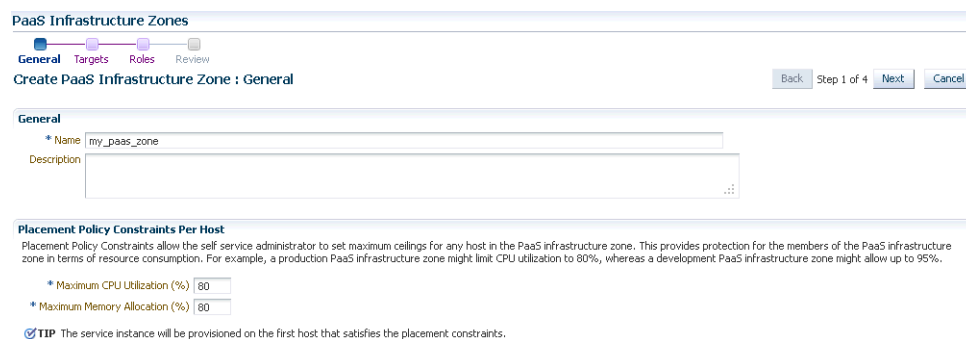
emcli grant_privs -name=CLOUD_ENGINE_USER -privileges="GET_CREDENTIAL;CRED_ANME=<>:CRDED_OWNER=<>"
```

- 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 **Enterprise** menu, select **Cloud**, then select **Middleware and Database Home**.
2. From the **Middleware and Database Cloud** menu, select **Create PaaS Infrastructure Zone**. The Create PaaS Infrastructure Zone: General page appears.

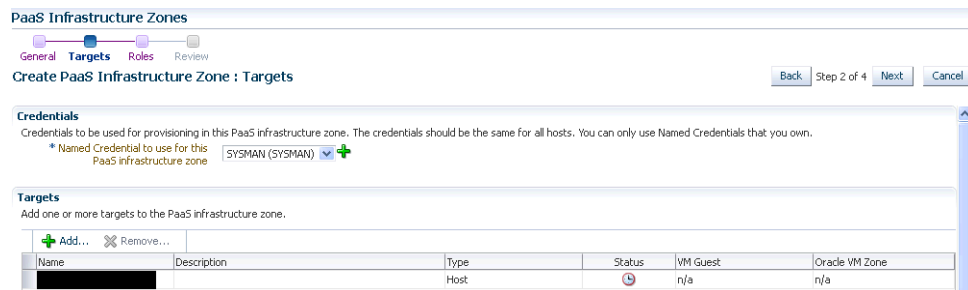
Figure 10–1 Create PaaS Infrastructure Zone: General



3. Enter a name and description for the zone, and specify appropriate values for placement policy constraints based on the resources available to the members. The `EM_SSA_ADMINISTRATOR` can use these 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 over a specified period.
 - Maximum Memory Allocation: The average memory allocation over a specified period.

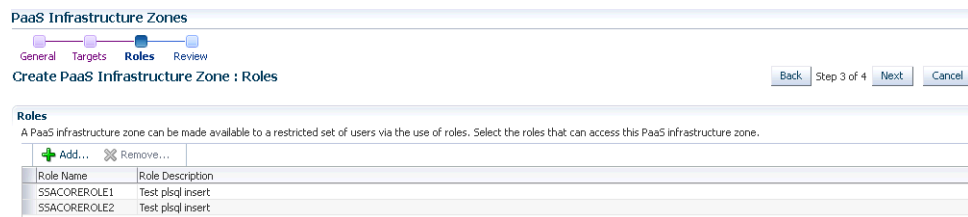
During provisioning, hosts with these constraints will not be considered for provisioning operations. Click **Next**. The Create PaaS Infrastructure Zone page appears.

Figure 10–2 Create PaaS Infrastructure Zone: Targets Page



4. 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 the same for all the hosts in this zone.
5. In the Targets region, click **Add** to add one or more targets to the zone. A target can be a host or an OVM Zone. See [Section 5.13, "Creating a Zone"](#) for details.
6. Click **Next**. The Create PaaS Infrastructure Zone: Roles page appears.

Figure 10–3 Create PaaS Infrastructure Zone: Roles Page



7. A PaaS Infrastructure Zone can be made available to a restricted set of users through the use of roles. Click **Add** to add and select the roles that can access this PaaS Infrastructure Zone.
8. 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.

10.1.1 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**.
4. Follow the steps in the wizard to edit the necessary parameters and click **Submit** to accept the changes.

10.1.2 Deleting a PaaS Infrastructure Zone

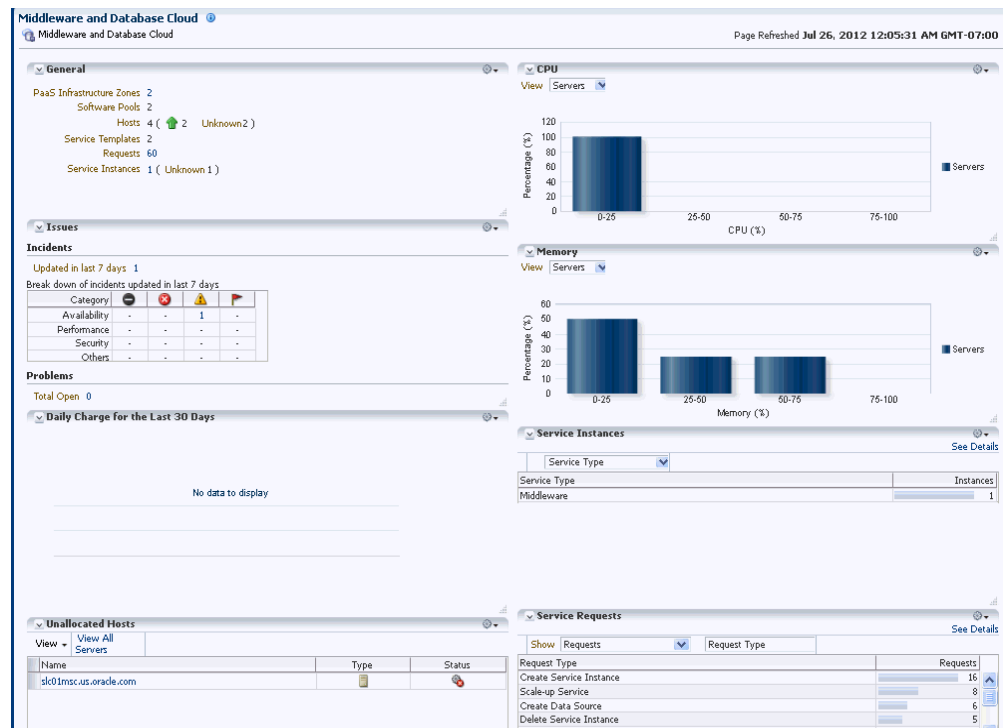
To delete a PaaS Infrastructure Zone, follow these steps:

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.

10.2 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 10–4 Middleware and Database Cloud Home Page



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.
 - **Oracle VM Zones:** The Oracle VM Zones associated with the PaaS Infrastructure zones.
 - **Hosts:** The physical hosts associated with the software pools.
 - **Service Templates:** The 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 utilizations 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.
- **Service Instances:** The total number of instances by service type, zone, service template, status or owner. Select an 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.

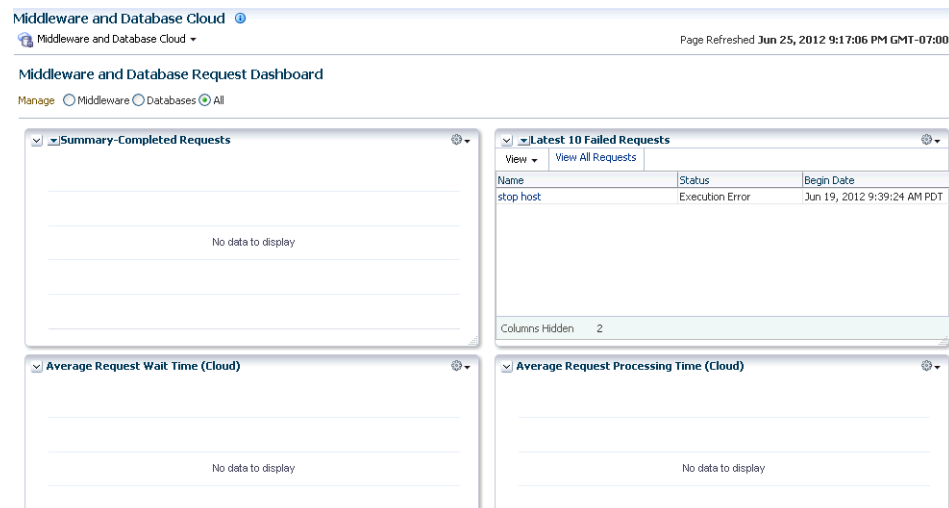
From the Middleware and Database Cloud menu, you can do the following:

- **Create PaaS Infrastructure Zone:** You can launch the wizard to create the zone. See [Section 10.1, "Creating a PaaS Infrastructure Zone"](#) for details.
- **Service Instances:** You can view detailed information of all the members. See [Section 10.5, "Viewing the Service Instances Page"](#) for details.
- **Request Dashboard:** You can view all the database and middleware requests. See [Section 10.3, "Viewing the Middleware and Database Request Dashboard Page"](#) for details.

10.3 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 10–5 Middleware and Database Request Dashboard Page

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.

10.4 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 10–6 Request Dashboard Details Page

Request Dashboard > All Requests > Create MW Domain1

Create MW Domain1

General Deployment

Request Name Create MW Domain1
 Type Create Service Instance
 Status Success
 Description Create mwdomain target
 Service Type Middleware
 Deployed On test paas zone
 Submission Date Jun 12, 2012 4:39:24 PM PDT
 Created By SSACOREUSER1
 Begin Date Jun 12, 2012 9:39:24 AM PDT
 End Date Jun 20, 2012 9:39:24 AM PDT
 Last Modified By SSACOREUSER1
 Last Modified On Jun 12, 2012 9:39:24 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.
- **Deployment:** This tab shows a list of the targets that have been deployed. Click on the link to view the details of the deployment procedure.

10.5 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 10–7 Service Instances Page

Middleware and Database Cloud

Middleware and Database Cloud Page Refreshed Jun 25, 2012 10:32:33 PM GMT-07:00

Service Instances

Search Saved Search Default Search

Match All Any

Service Instance Name Starts with Service Type Equals

View

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.

10.6 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.

10.7 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.

Figure 10–8 PaaS Infrastructure Zones Page

The screenshot shows the 'PaaS Infrastructure Zones' page. At the top, it says 'Middleware and Database Cloud' and 'Page Refreshed Aug 18, 2012 12:43:13 AM GMT+00:00'. Below the search bar, there are buttons for 'Create', 'Edit', and 'Delete'. The table below has columns for 'PaaS Infrastructure Zone', 'Service Instances', 'Members', 'Servers', 'Server Status Summary', and 'Server Incidents'. The 'Server Status Summary' column contains icons for various server states and counts. The 'Server Incidents' column contains icons for different incident types and counts.

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

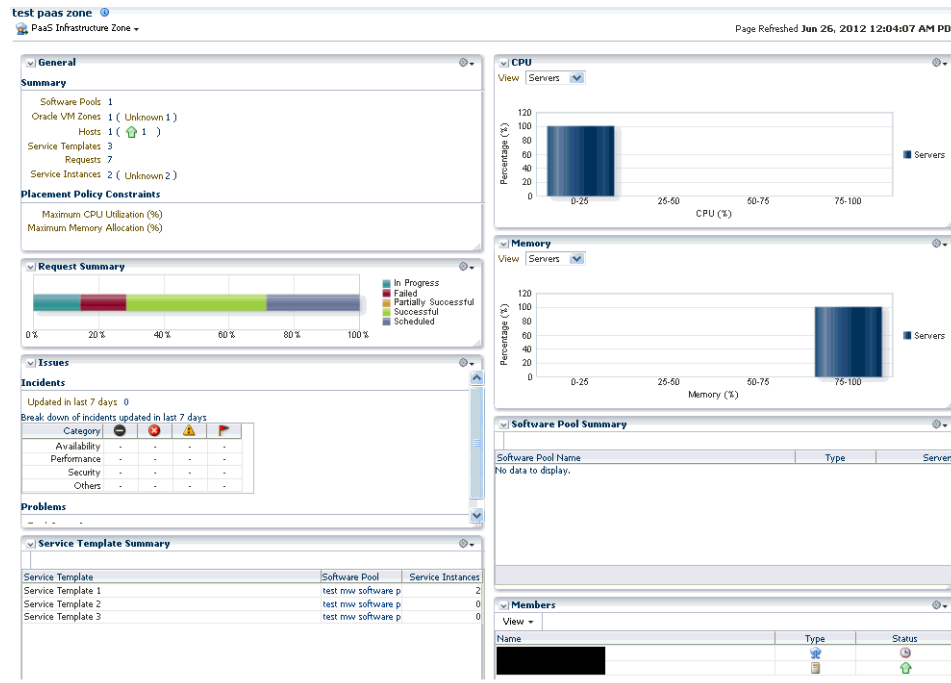
You can do the following:

- **Create:** Click **Create** to launch the Create PaaS Infrastructure Zone wizard. See [Section 10.1, "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.

10.8 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 10–9 PaaS Infrastructure Zone Home Page



The PaaS Infrastructure Zone Home page contains the following regions:

- **General:** This region lists the total number of Software Pools, Oracle VM Zones, Hosts, Servers, 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.
- **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. 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.

Part IV

Database as a Service

This section provides an overview of database as a service (DBaaS), describes the configuration steps and database self service portal.

It contains the following chapters:

- [Chapter 11, "About Database as a Service \(DBaaS\)"](#)
- [Chapter 12, "Setting Up Database as a Service"](#)
- [Chapter 13, "Setting Up and Using the Self Service Application for DBaaS"](#)

About Database as a Service (DBaaS)

This chapter provides an introduction to database as a service. It contains the following sections:

- [About Database as a Service \(DBaaS\)](#)
- [Database as a Service in Enterprise Manager](#)

11.1 About Database as a Service (DBaaS)

Enterprises need to support hundreds or even thousands of applications to meet growing business demands and this growth has driven up the cost of running and managing the databases under those applications. This puts a lot of stress on the IT budget and makes it harder to provide databases to support new requirements such as Web 2.0 applications or other emerging collaboration solutions or even to support other uses such as increased application testing.

Private Clouds enable customers to consolidate servers, storage, and database workloads onto a shared hardware and software infrastructure. By providing on-demand access to database services in a self-service, elastically scalable and metered manner, databases deployed as DBaaS on Private Cloud offer compelling advantages in cost, quality of service, and agility.

DBaaS extends the Oracle Private Cloud Management solution by automating the lifecycle of a database and allowing users to request database services through self-service portal. With this solution, IT Managers no longer have to perform mundane administrative tasks for provisioning databases. Database users can get instantaneous access to new database services through the Self Service Portal.

DBaaS in Enterprise Manager extends the capabilities of the Private Cloud by:

- **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 Private 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 Private 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.

11.2 Database as a Service in Enterprise Manager

In Enterprise Manager, Database as a Service (DBaaS) is implemented through the following options:

- **Virtual Machine Based:** The database is deployed as a part of Virtual Assembly or Template and several virtual machines share the same physical server. This offers the maximum level of isolation (at OS level) but creates manageability problems owing to VM sprawl. This option is available through the Infrastructure as a Service (IaaS) Self Service Portal.
- **Shared Cluster:** The database is deployed on existing clusterware. Typically the Grid Infrastructure (Clusterware, ASM) and Database software are preinstalled. The cloud service entails the deployment of databases on top of that infrastructure. This option is available in the DBaaS Self Service Portal.
- **Shared Installation:** The database is deployed as a single instance database on an existing installation. This option is available in the DBaaS Self Service Portal.

Before you implement DBaaS, it is important to select a delivery model that best satisfies the requirements of your application or intended use of database services. It is also possible to adopt multiple delivery models. For example, some simple applications may be hosted on a shared database model, if their requirements are limited to just a schema. Other more complex applications like Oracle Applications may require significant tweaking of the database parameters and would need the shared cluster or shared installation model. The Enterprise Manager DBaaS offering will allow you to implement and manage multiple such delivery models.

Setting Up Database as a Service

This chapter describes the steps required to setup the Database Cloud Infrastructure. In particular, it contains the following sections:

- [Getting Started](#)
- [Defining Roles and Assigning Users](#)
- [Adding Hosts](#)
- [Configuring Privilege Delegation Settings](#)
- [Configuring the Software Library](#)
- [Setting Up Credentials for Provisioning](#)
- [Provisioning Database Software](#)
- [Creating Oracle Home and Configuring the Oracle Listener](#)

12.1 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 2.1, "Installing Virtualization and the Cloud Plug-ins"](#).

Table 12–1 *Getting Started with DBaaS*


Step	Task	Role
1	Define roles for administrators and self service users. See Section 12.2, "Defining Roles and Assigning Users" .	Super Administrator
2	Install the Management Agent on unmanaged hosts so that they can be monitored by Enterprise Manager. Section 12.3, "Adding Hosts" .	Cloud Administrator
3	Configure Privilege Delegation Settings on your managed hosts. Section 12.4, "Configuring Privilege Delegation Settings" .	Super Administrator
4	Configure the storage location for the Software Library. See Section 12.5, "Configuring the Software Library"	Cloud Administrator
5	Set up provisioning credentials. See Section 12.6, "Setting Up Credentials for Provisioning" .	Self Service Administrator
6	Configure the Listener. See Section 12.8, "Creating Oracle Home and Configuring the Oracle Listener" .	Self Service Administrator

Table 12–1 (Cont.) Getting Started with DBaaS

Step	Task	Role
7	Provision the database software or create Oracle Homes. See Section 12.7, "Provisioning Database Software" .	Self Service Administrator

Figure 12–1 shows the sequence of steps required to enable database as a service.

Figure 12–1 Enabling Database as a Service



12.2 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 built-in roles for the Database Cloud 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.
- **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

12.2.1 Creating a Custom Role for Self Service Application Users

The Cloud Management Self Service Portal allows self service users to provision and manage their own cloud services. Self service users need access only to the Self Service Portal and the resources allocated to them. These capabilities are inherent in the predefined **EM_SSA_USER** role.

Since the functions performed by the **EM_CLOUD_ADMINISTRATOR** and **EM_SSA_ADMINISTRATOR** roles are consistent across Enterprise Manager, these ready-to-use roles can be used as they are. On the other hand, the **EM_SSA_USER** role is used for quota assignment, and to limit access to Database Pools and Service Templates. Hence the predefined role needs to be modified to create custom SSA user roles that are based on the standard **EM_SSA_ROLE** role.

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 Siebel DBA team, BRM DBA team, 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.

12.2.2 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.

Note: Repeat these steps to create other users. For the Cloud Administrator and SSA Administrator users, the `EM_USER` and `PUBLIC` roles must not be removed.

12.3 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. To install the Management Agent on an unmanaged server, follow these steps:

1. Log in to Enterprise Manager as an `EM_CLOUD_ADMINISTRATOR` user.
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. For a detailed description of this procedure, see [Section 12.3, "Adding Hosts"](#).

Note: You can deploy the Management Agent on a maximum of 16 servers at a time. The total time required will thus increase if the Management Agent is deployed to a large number of servers.

12.4 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.

12.5 Configuring the Software Library

Oracle Software Library (Software Library) is a repository that stores software patches, virtual appliance images, reference gold images, application software and their associated directive scripts. It allows maintaining versions, maturity levels, and states of entities.

The software entities can be automatically mass-deployed to provision software, software updates, and servers using Enterprise Manager Cloud Control in a reliable and repeatable manner. These provisioning operations, which are unattended and can be scheduled, lead to substantial cost savings.

Besides acting as a repository for certified software entities, the Software Library is a logical interface between the deployment models and the automation framework required to perform a large number of patching and provisioning tasks. To configure the storage location for the Software Library, follow these steps:

1. Log in to Enterprise Manager as an `EM_CLOUD_ADMINISTRATOR` user.
2. From the **Setup** menu, select **Provisioning and Patching**, then select **Software Library**.
3. Select **OMS Shared File System** in the Storage Type list and click **Add...**

If you are using the OMS Agent Filesystem location, see the note in [Section 10.1, "Creating a PaaS Infrastructure Zone"](#) for setting up credentials for administrators.

4. 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.

5. A job is executed to upload all the ready-to-use content. This may take 15 to 30 minutes depending on your disk speed.

12.6 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 two types of credentials:

- Normal credentials are the host operating system credentials used to install the database software. For example, `oracle/oracle`.
- Privileged credentials are the host operating system credentials used to perform privileged actions like executing root scripts. For example, `oracle/oracle` with `sudo` to root.

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 this credential against a host target and click **OK**.

12.7 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**
 - Create a snapshot 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.

Note: Do not create a new database as part of this deployment procedure.
- **Using the Database Installer**
 - Run the Clusterware/ASM or Grid Infrastructure installer to set up the cluster and ASM (for RAC databases).
 - Run the Database Installer with the **software only** option on all hosts.

- Add the database Oracle Homes as Enterprise Manager targets. See [Section 12.8, "Creating Oracle Home and Configuring the Oracle Listener"](#) for details.

For more details on provisioning the database software, see the *Enterprise Manager Lifecycle Management Administrator's Guide*.

12.8 Creating Oracle Home and 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 12.3, "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.

Setting Up and Using the Self Service Application for DBaaS

This chapter describes the process to set up and use the database self service application. In particular, it contains the following sections:

- [Customizing and Saving the Database Deployment Procedure](#)
- [Setting Up the Self Service Portal](#)
- [Using the Database Self Service Portal](#)
- [Requesting Databases](#)
- [Troubleshooting Tips for Database Request Failures](#)

13.1 Customizing and Saving the Database Deployment Procedure

A Database Deployment Procedure is a series of steps that is used to create a database. With the Database Self Service Portal, users can select a service template that closely resembles their requirements which identifies the size and configuration of the database to be created.

To customize the database deployment procedure, follow these steps:

1. Log in to Enterprise Manager as an administrator with the `EM_SSA_ADMINISTRATOR` role.
2. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Database Provisioning**.
3. Select the **Create Database Procedure** option and click **Launch**.
4. Select or enter appropriate values for procedure inputs, and click the **lock** icon once the values are finalized.

Note: You must lock the variables as specified in the table below. If any of these variables are not locked, the database deployment procedure cannot be used.

Field	Description	Locked Y/N?
Version	Select from a list of supported database versions	Y
Database Type	Select among different database configurations	Y

Field	Description	Locked Y/N?
Hosts (Oracle Home + Credential)	Select a single sample host, Oracle Home and credentials. While creating Real Application Cluster databases, select a cluster, the desired Oracle Home, credentials for the cluster, and a sample node. Note: These inputs will be ignored when the procedure is published as a service template.	N
Database Template	Select the database template location. The location can be Software Library or Oracle Home. The template selected must be compatible with the selected Oracle Home version. If you have selected Oracle Home, select the template from the Oracle Home. The default location is ORACLE_HOME/assistants/dbca/templates. Note: Due to a known issue, when this field is locked, it may be difficult to identify the template selected from the Oracle Home.	Y
Cluster Database Configuration Type (RAC Only)	Select the configuration type as Administrator Managed and suitable nodes are selected by default.	NA
Global Database Name	Enter a dummy value. This value will be auto-generated when this procedure is published as a service template.	N
SID	Enter a dummy value. This value will be auto-generated when this procedure is published as a service template.	N
Passwords (SYS, SYSTEM, DBSNMP)	Specify all three passwords here. The SYS password is used and the SYSTEM and DBSNMP are displayed.	Y
Storage Type	Select the storage type, whether File System or Automatic Storage Management (ASM).	Y
Database File Location	Specify the location or disk group where data files, temporary files, redo logs, and control files will be stored.	N
Recovery File Location	To enable backup and restore for Self Service users, check the Use Fast Recovery Area and Enable Archiving check boxes. When you check the Use Fast Recovery Area check box, you are prompted for the location or disk group and size required for storing recovery related files. If these check boxes are selected, the Self Service user can schedule Oracle recommended RMAN backups for the database.	Y
Memory Parameters	Specify this value according to your database size definition. For example, small = 1 GB, medium= 2 GB, large = 4 GB, and so on.	Y
Database Sizing	Specify the Block Size and number of Processes. If you have selected a database template with data files in the Database Template page, you cannot edit the Block Size.	Y
Host CPU Count	Specify this value according to your database size definition. For example, small = 1 CPU, medium=2 CPU, large= 4 CPU, and so on.	Y

Field	Description	Locked Y/N?
Character Set	Select the required character set. The default character set is based on the locale and operating system.	Y
Database Connection Mode	Select the server mode. For shared server mode, specify the number of shared servers.	Y
Listener Configuration	As part of the prerequisites, the listener has to be configured for each host and Oracle Home. After it has been configured, select the correct listener and lock the value.	Y
Custom Script	Optional custom script that will be executed after the database has been created.	Y
Schedule	This field will be ignored. The schedule is specified by the Self Service user as part of the database request.	NA
Review	Review all the inputs provided in the deployment procedure.	NA

- On the Review page, click **Save** and enter a meaningful name for the deployment procedure.

Note: Do not click **Submit** on the Review page.

- After the deployment procedure has been saved, click **Cancel** to exit the wizard.

13.2 Setting Up the Self Service Portal

The 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.

Prerequisites

- DBaaS must be enabled by following the steps described in [Chapter 12, "Setting Up Database as a Service"](#).
- PaaS Infrastructure Zone must have been created. See [Section 10.1, "Creating a PaaS Infrastructure Zone"](#).

Follow this process to set up the database cloud self service portal:

- Log in to Enterprise Manager Grid Control as an administrator with the `EM_SSA_ADMINISTRATOR` role.
- From the **Setup** menu, select **Cloud**, then select **Database**.
- The Database Cloud Self Service Portal Setup page appears. Click **Database Pools**. A database pool is a collection of servers or clusters with database software installed.

Figure 13–1 Database Self Service Portal Setup - Database Pools

Database Cloud Self Service Portal Setup Page Refreshed Jul 26, 2012 5:56:52 AM UTC

Database Pools
A database pool is a collection of servers or clusters with database software (Oracle Home) pre-installed.

Request Settings | Quotas | Service Templates | Chargeback

Pool Name	Owner	PaaS Infrastructure Zone	Member Target Type	Targets	Description
pool1		comp zone	Oracle Home	2	
pool1		zone1	Oracle Home	1	

4. Click **Create**. The Create Software Pool: General page appears. Enter a name and description for the database pool.
5. Enter the placement policy constraints. 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.

Figure 13–2 Database Self Service Portal Setup - Create Pool: Name

Software Pools
Create Software Pool : General Back Step 1 of 2 Next Cancel

General

* Name
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 Database Instances

TIP The service instance will be provisioned on the first member that satisfies the placement constraints.

6. Click **Next**. The Create Software Pool: Targets page appears. In this page, you can add one or more Oracle homes to the database 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 10.1, "Creating a PaaS Infrastructure Zone"](#) for details.
 2. Specify the Version of the Oracle home being added to the database pool.
 3. Click **Add** and select one or more Oracle homes to be added to the database pool. Click **Submit**.

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.

7. Click the **Request Settings** tab. Specify the following:
 - Future Reservation: How far in advance can a request be made.
 - Request Archive Retention: Period for which archived requests are stored before they are automatically deleted by the system.
 - Request Duration: Maximum duration for which requests can be made.

Click **Apply**.

8. Click the **Quotas** tab. 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 database provisioned through the Self Service Application.

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.

9. 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:
 - Total allocated memory that can be assigned for all the databases owned by a specific user.
 - Total allocated storage that can be assigned for all the databases owned by a user.
 - The number of databases 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.

10. Click **Create**. In the Create New Quota window, select the Role Name for which the quota is to be granted.

Note: Only roles that have the `EM_SSA_USER` privileges are displayed in the Select Roles dialog box.

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:

Role Name	Attribute	Value
Dev Role	Memory	2 GB
	Storage	15 GB
	DB Count	10
Test Role	Memory	5 GB
	Storage	10 GB
	DB Count	12
PM Role	Memory	Quota not set
	Storage	Quota not set
	DB Count	Quota not set

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:

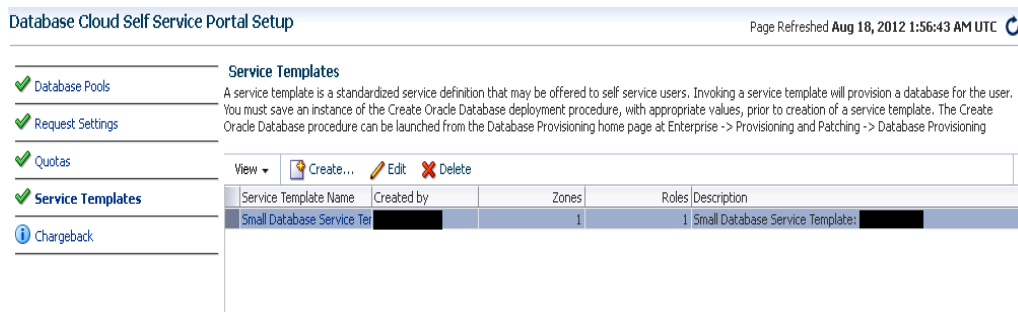
Value for Memory = MAX (Dev, Test, PM) = MAX (2, 5, NA) = 5 GB
 Value for Storage = MAX (Dev, Test, PM) = MAX (15, 10, NA) = 15 GB
 Value for DB Count = MAX (Dev, Test, PM) = MAX (10, 12, NA) = 12 Databases

A sample quota validation log generated by the user request follows:

```
***** UserName - SSA_USER
***** Quota Limits:
*** Memory: 10
*** Storage: 10
*** Instances: 10
*** Memory Used: 6.732421875
***** Current Usage:
*** Storage Used: 5.28
*** Instances Used: 5
```

11. Click the **Service Templates** tab. A service template is a standardized service definition that can be selected by self service users when requesting a new database. You can create service templates that can be used to provision databases in a set of selected zones.

Figure 13–3 Database Cloud Self Service Portal Setup - Templates



12. Click **Create** in the Service Templates page to launch the Create New Service Template Wizard.

Note: Before you create a service template, you must have saved a copy of a Create Database Deployment Procedure with appropriate values.

13. Enter a name and description for the service. Select a saved deployment procedure that is to be used to create the service template. Click **Next**. The service template description appears, and you can select the right template. Sample service template descriptions are shown in [Figure 13–3](#).
14. Click **Add** to add one or more roles. Users belonging to the roles selected here can access the service template. Click **Submit**. The newly created service template can now be accessed through the self service portal while requesting databases.

To edit an existing 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 database 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 access this template.
15. Click the **Chargeback** tab. 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 [Chapter 4, "Chargeback Administration"](#).

13.3 Using the Database Self Service Portal

The Database Self Service Portal allows self service users to view, monitor, and deploy databases on selected zones. Users can request databases using any of the predefined service templates or database specifications that closely resemble their requirements. Users can provision databases easily without being aware of the hardware and software infrastructure details. The entire database lifecycle of creation, maintenance and retirement can be done by the self service users with zero or minimal intervention from an official Database Administrator.

Figure 13–4 Database Cloud Self Service Portal

The screenshot shows the Database Cloud Self Service Portal interface. At the top, it says 'Database Cloud Self Service Portal' and 'Page Refreshed Aug 6, 2012 12:37:00 AM UTC'. Below that are navigation tabs for 'Manage', 'My Servers', 'My Databases', and 'My Middleware'. The 'Home' tab is selected, and there is a 'Chargeback' tab. The main content area is divided into four sections:

- Notifications:** Shows 'Databases Due to Expire in Next 7 Days: 0' and a link to 'Service Templates Displayed in Last 7 Days: 7'.
- Usage:** Shows 'Databases: 1', 'Memory (GB): 1', and 'Storage (GB): 1.23'.
- Databases:** A table with columns: Name, Type, Status, Start, End, Zone. One row is visible: Name (redacted), Type: Cluster Database, Status: (green up arrow), Start: Aug 5, 2012 11:27:21, End: Indefinitely, Zone: 11203 RAC ZONE. Above the table are buttons for 'View', 'Request Database...', and 'Delete'.
- Requests:** A table with columns: Name, Status, Request Type, Submission Date, Begin Date. It shows a list of recent requests with various statuses like 'Success', 'Insufficient Resource', and 'Scheduled'.

The Database Cloud Self Service Portal Home page displays the following details:

- **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.
 - **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 cancel 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: 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 menu. Click **Export** to export the details of the charge plan to a .csv file.
 - 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 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 appears.
- Click **My Preferences**. Before you request resources using SSA, you can set up your user profile by specifying the following details:
 - Default Host Password: The default password for the host.
 - Default Zone: The default zone for the requests.
 - Default Source Software: The default location in which the assemblies, templates, and deployment plans are present.
 - Default Email Address for Notifications: The default email address to be used to receive notifications.
 - Days Before the First / Second Notification: The number of days before the first and second notifications are sent.
 - Enable Quota Related Notification: If you select this checkbox, notifications are sent when the quota reaches a specified threshold.
 - Quota Threshold for Receiving Notifications: If you have selected the **Enable Quota Related Notification** checkbox, specify the quota threshold in this field.
 - **Database Preferences**
 - Default Zone: Select the default zone for the requests.
 - **My Middleware Preferences**
 - Default Zone: Select the default middleware zone for the requests.
 - **My Servers Preferences**
 - Default Server Root Password
 - Default Zone for New Requests
 - Default Source Software

Click **Apply** to save the preferences. The values specified here will be used when you request resources using SSA.

13.4 Requesting Databases

You can deploy databases to a specific zone with a selected service template. To request a database, 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. On the Self Service Portal Homepage, click **Databases** on the upper right corner of the screen.
3. The Home page of the Database Cloud Self Service Portal appears. Click **Request Database** in the Home page.

Note: If the Infrastructure Cloud Self Service Portal appears, click the **Databases** option at the top left-hand corner of the page to toggle to the Database Cloud Self Portal page.

4. Choose a service template from the list and click **Select**. If you have only one template, it will be automatically selected. Typically, the service template description should provide the configuration details for the database.

Figure 13–5 New Database Request: General

5. In the New Database Request: General page, enter a meaningful name for the request and select a zone on which the database is to be deployed. The request is tied to the provisioned database until the database is retired. This is indicated by the status of the request.
6. Click **Next**. In the New Database Request: Deployment Inputs page, specify the custom variables. Click **Next**.

Currently, the only inputs are the set of credentials for the database. For security reasons, users do not get access to SYS and SYSTEM accounts, instead a custom master account is created for the user, and this is the primary administrative account for the database.

The master account is granted the DBA role, and following privileges are revoked from it:

- Alter database
- Alter system
- Create any directory
- Drop any directory

- Grant any object privilege
- Grant any privilege
- Grant any cycle

Note: You cannot use any of the default Oracle user names for the master account. The user name is case insensitive. The following user names cannot be used as the master account user name:

- USER
- SYS
- SYSTEM
- DBSNMP
- OUTLN
- MGMT_VIEW
- FLOWS_FILES
- MDSYS
- ORDSYS
- EXFSYS
- WMSYS
- APPQOSSYS
- APEX_030200
- OWBSYS_AUDIT
- ORDDATA
- CTXSYS
- ANONYMOUS
- SYSMAN
- XDB
- ORDPLUGINS
- OWBSYS
- SI_INFORMTN_SCHEMA
- OLAPSYS
- SCOTT
- ORACLE_OCM
- XS\$NULL
- MDDATA
- DIP
- APEX_PUBLIC_USER
- SPATIAL_CSW_ADMIN_USR
- SPATIAL_WFS_ADMIN_USR

If a default Oracle user name is used, an error message appears.

Figure 13–6 New Database Request: Deployment Inputs

7. Specify the schedule for the request. The Start Date (immediate or deferred) determines when the database will be provisioned, while the End Date determines when the database will be retired.
8. Review the details of the request.
9. Click **Submit** to deploy the database to the selected zone. You are directed back to the Database Self Service Portal where you can monitor the progress of your request. The process can take 5 to 10 minutes, depending on the performance and load profile of the Cloud Infrastructure. The resource usage (against your quota limit) charts will also be updated once the database has been created.

Once the request has been completed, the new database is visible in the My Databases table. You can click the database name to view its Home page.

13.5 Troubleshooting Tips for Database Request Failures

This troubleshooting section can be used to determine the possible cause of a database request failure and to take suitable remedial actions.

Submitting a database request involves the following steps:

1. Pre-processing
 1. Initialization
 2. User Quota Validation
 3. Placement logic execution to find a suitable host
2. Database Creation
3. Post Processing

Each of these steps is executed sequentially and each step must be successfully completed before the consecutive step can be executed. Failure in any one of the steps causes the database request to fail.

13.5.1 Viewing the Output Log of a Request

To view the output log, follow these steps:

1. Follow the steps listed in the [Section 10.3, "Viewing the Middleware and Database Request Dashboard Page"](#) section and select the Deployment Procedure associated with the request to view the execution details.

Figure 13–7 Procedure Execution Details

The screenshot shows the 'Provisioning' page for a procedure named 'SYSMAN - Mon Dec 05 03:02:05 PST 2011'. It includes a 'General Information' section with details like 'Run SYSMAN - Mon Dec 05 03:02:05 PST 2011', 'Procedure Database as a Service Procedure for SI', and 'Status Running'. Below this is a 'Status Detail' section with a table of steps:

Name	Status	Type	Description	Error Handling Mode
Database as a Service Procedure for SI	Running		This Database as a Service Procedure creates an Oracle Database.	Stop On Error
Initialize Deployment Procedure	Succeeded	Computational	Initializes the current Deployment Procedure execution. Derived variables are set with computations. Do not disable or delete this step.	Stop On Error
Compute the Pool Associated with the Selected DP	Succeeded	Computational	Computes the Pool Associated with the Selected DP. Do not disable or delete this step.	Stop On Error
Validate the Quotas	Succeeded	Computational	Checks for the Quota availability for the target(s).	Stop On Error
Obtain the Target Node List using the Placement Algorithm	Succeeded	Computational	Obtains the Target Node List using the Placement Algorithm. Do not disable or delete this step.	Stop On Error
Configure/Create the Database	Running	Procedure Step	Configures/Creates the Database	Stop On Error
Assign the Target Privileges		Computational	Assigns the Privileges to the Target as well as to the Service Instance to Both the Service User and the Service Admin	Stop On Error

- The Name column lists all the steps for the request. Click the hyperlink for the step with the **Failed** status.

Figure 13–8 Procedure Step Status Page

The screenshot shows the 'Step Status: Initialize Deployment Procedure' page. It displays 'General Information' for the step, including its name, type (Computational), and description. The 'Job Step' section shows the status as 'Failed - Job Failed' and provides options to 'Ignore' or 'Retry'.

- Click the hyperlink next to **Status** field, and view the output log for the step.

Figure 13–9 Output Log Page

The screenshot shows the 'Output Log' page for a failed step. It displays the job name 'COMPUTE_INIT_DP_B35742206138F63AE040E50A278C6600' and the error message: 'ERROR: Master Account username "Sys" cannot be a oracle default username. Error while evaluating expression: \${initDP()}: {Error::: Master Account username cannot be a oracle default username : Sys}'.

13.5.2 Initialization Step Errors

The initialization step focuses mainly on the database master account validation. It checks that the master account user name to ensure that none of the default Oracle user names listed below are used. The user names are case insensitive.

- USER
- SYS
- SYSTEM
- DBSNMP
- OUTLN
- MGMT_VIEW
- FLOWS_FILES
- MDSYS
- ORDSYS
- EXFSYS
- WMSYS
- APPQOSSYS
- APEX_030200
- OWBSYS_AUDIT
- ORDDATA
- CTXSYS
- ANONYMOUS
- SYSMAN
- XDB
- ORDPLUGINS
- OWBSYS
- SI_INFORMTN_SCHEMA
- OLAPSYS
- SCOTT
- ORACLE_OCM
- XS\$NULL
- MDDATA
- DIP
- APEX_PUBLIC_USER
- SPATIAL_CSW_ADMIN_USR
- SPATIAL_WFS_ADMIN_USR

If any of these default user names have been used, a message indicating that the user name is invalid appears.

If this request was made using EMCLI, the Initialize Deployment Procedure step will fail and specify the default Oracle user name in the error message as follows:

Error: The master account user name '`<user specified value>`' cannot be an Oracle default name.

Solution: View the output log for the Initialize Deployment Step by following the steps listed in [Section 13.5.1, "Viewing the Output Log of a Request"](#). Resubmit the request and specify a user name that is not in the list of Oracle default user names.

13.5.3 User Quota Validation Errors

Quota validation errors occur if the user's usage exceeds the ceiling on any of the quota attributes. Users can view their current usage against the quota limits on the Database Self Service Portal Home page. A database request may fail if the user's usage is close to their quota limits.

Solution: View the output log for the User Quota Validation Step by following the steps listed in [Section 13.5.1, "Viewing the Output Log of a Request"](#). Follow these steps to reduce quota:

1. Retire (relinquish) any databases requested through the DBaaS Console to free up resources and lower their current usage. It may take up to 30 minutes before the effect of the retirement is visible.
2. If no databases can be retired, additional databases may need to be granted.
 - A user with the `EM_CLOUD_ADMINISTRATOR` role may consider creating a new temporary `EM_SSA_USER` role, and assigning it a higher quota limit especially for the affected quota attribute. This approach will ensure that the other users will remain unaffected by the quota changes.
 - A user with the `EM_CLOUD_ADMINISTRATOR` role may alter the quota limits for the particular role assigned to the affected user. This increases the quota limits for all users associated with that role.

13.5.4 Placement Logic Failures

Placement logic is used to locate a suitable host, within the selected Database Zone, where a new database can be created. A host is deemed suitable if it satisfies all of the placement constraints. The suitable host must be compliant with the placement constraints and can be used to create the new database. Placement logic will fail if a suitable host cannot be found with a selected zone.

While locating a suitable host, placement logic:

- Uses a 7 day average (lesser if total duration of metric collection for the host is under 7 days) for Memory and CPU.
- Provides instantaneous instance count.
- Places a *predictive* block on any service request under execution to account for concurrent requests.
- Does not delete or alter any databases already running on the host if its usage exceeds the placement constraint's ceiling value.
- Since Placement relies on 7 day average of metrics such as Memory, CPU, and so on, the resources may not be released immediately.

Solution: View the output log for the Obtain the Target Node List using the Placement Algorithm step by following the steps listed in [Section 13.5.1, "Viewing the Output Log of a Request"](#). Follow these steps:

1. Determine the current load on the hosts, and free up resources by shutting down unused or unnecessary applications or databases.
 1. As a user with the EM_SSA_ADMINISTRATOR role, navigate to the Database Zone Home page on which the database request failed. Here you will find the Placement constraint values for the current zone, and the current load on the hosts by way of bar graphs for memory and CPU.
 2. The execution details for the deployment procedure associated with the failed database request, shows the summary for how the hosts under the current zone fared against the given placement constraints.
2. Assuming the resource constraints are due to increasing popularity of this service and not due to a runaway database service, more hosts may have to be added to the zone to increase its capacity.

The following table lists the various scenarios and their sample execution logs.

Table 13–1 Sample Scenarios

Scenario	Sample Execution Log for a Zone Containing 100 Hosts	Comments
	<p>***** Placement Algorithm analyzing Zone "ABC Zone"</p> <p>20 members are currently below the CPU threshold of "80%"</p> <p>20 members are currently below the Memory threshold of "80%"</p> <p>25 members are currently below the DB Instance threshold of "20"</p> <p>Placement Algorithm Found "20" suitable hosts for Placement (Pending the current service requirements).</p> <p>HOST_CREDS=SYSMAN:AIME_NORMAL</p> <p>SID is "DB5ad5a8"</p> <p>Analyzing suitable hosts that would satisfy the current service request requirements ...</p> <p>Analyzing host "unsuitable.us.myorg.com" – NOT suitable</p> <p>Analyzing host "unsuitable1.us.myorg.com" – NOT suitable</p> <p>Analyzing host "unsuitable2.us.myorg.com" – NOT suitable</p> <p>Analyzing host "unsuitable3.us.myorg.com" – NOT suitable</p> <p>Analyzing host "suitable.us.myorg.com" - suitable</p>	<p>There is more than one host in the selected zone that satisfies all of the constraints.</p>

Table 13–1 (Cont.) Sample Scenarios

Scenario	Sample Execution Log for a Zone Containing 100 Hosts	Comments
Failing CPU	<p>***** Placement Algorithm analyzing Zone "ABC Zone" :</p> <p>0 members are currently below the CPU threshold of "80%"</p> <p>10 members are currently below the Memory threshold of "80%"</p> <p>12 members are currently below the DB Instance threshold of "10"</p> <p>***** Placement Algorithm did not find any suitable target</p>	<p>There are no hosts under the selected zone which currently satisfy the CPU Constraint of MAX CPU 80%. Hence no suitable target in the zone satisfies all the placement constraints.</p>
Failing Instance Count	<p>***** Placement Algorithm analyzing Zone " ABC Zone " :</p> <p>12 members are currently below the CPU threshold of "80%"</p> <p>10 members are currently below the Memory threshold of "80%"</p> <p>10 members are currently below the DB Instance threshold of "10"</p> <p>***** Placement Algorithm did not find any suitable target</p>	<p>There are no members under the selected zone which currently satisfy the DB Instance constraint of MAX 10. Hence no suitable target in the zone satisfies all the placement constraints.</p>

13.5.5 Troubleshooting Failures in Other Steps

The Database Creation step, as the name suggests, is responsible for creating the requested database on the host selected by the placement logic. Typically, this step may fail due to conditions beyond the control of Enterprise Manager like hardware, network, storage, and operating system issues. The execution details for the deployment procedure will provide the cause of failure.

Next is the Post Processing step. Most steps will rollback all changes made by them in case of failure; except this step. This step can be retried multiple times once the cause of failure has been fixed. Consequently, the database created in the prior step is not deleted automatically, and instead the DBaaS administrator has the opportunity to fix the issue and retry the Post Processing step.

Solution:

1. If the Database Creation step completed successfully, then find the cause of failure for the Post Processing step by reviewing the execution details for that step, and retry once the issue has been resolved. To do this, navigate to the deployment procedure associated with the request and click **Retry**. This will re-execute the Post Processing steps.
2. If the request cannot be salvaged, and the Post Processing step continues to fail, the database must be deleted.
 - Determine the host, Oracle Home, and SID for the database using the Database Target Home page, or the Placement Logic step of the deployment procedure.
 - Login to the host and navigate to the ORACLE_HOME directory. Execute the following steps to delete the database from the server

```
setenv ORACLE_HOME <ORACLE_HOME>
```



```
$ORACLE_HOME/bin/dbca -silent -deleteDatabase -sourceDB <Database  
SID>
```

Optionally, you can run DBCA in graphical mode and delete the database as well.

3. Next, delete the Database target from Enterprise Manager 12c. This can be done by locating the target using the target search, and then using the right click menu to remove the target.

Part V

Middleware as a Service

This section provides an overview of middleware as a service (MWaaS), describes the configuration steps and the middleware self service portal.

It contains the following chapters:

- [Section 14, "About Middleware as a Service \(MWaaS\)"](#)
- [Section 15, "Setting Up Middleware as a Service"](#)
- [Section 16, "Setting Up the Middleware Self Service Portal"](#)
- [Section 17, "Using the MWaaS Self Service Portal"](#)

About Middleware as a Service (MWaaS)

This chapter provides an overview of Middleware as a Service (MWaaS).

14.1 About Middleware as a Service

Middleware as a Service (MWaaS) provides a complete application development and deployment environment. It delivers a complete runtime environment comprised of all services necessary to deploy and run an enterprise-class application. MWaaS facilitates cheaper and faster deployment of applications as developers need not deal with the complexities of the underlying hardware and software components.

Oracle Enterprise Manager Cloud Control 12c offers a single management tool for creating, monitoring and managing a private or public MWaaS environment. MWaaS delivers a complete runtime environment comprised of all services necessary to deploy and run an enterprise-class application, including services such as application hosting, persistence store, application integration and APIs that enable programmatic access to additional computing services that might be required by an application. Identity services are an example of APIs available within a PaaS environment.

MWaaS in Enterprise Manager extends the Oracle Private Cloud Management solution by:

- **Providing Automated Discovery and Baselineing of Resources:** Enterprise Manager provides an integrated workflow for agent deployment and target discovery on selected auto discovered hosts.
- **Capacity and Consolidation Planning:** Generates consolidation plans based on resource distribution and utilization. Details on consolidation ratio, destination server utilization, server mapping, and so on are provided.
- **Providing a Centrally Managed Software Repository:** Enterprise Manager provides a central repository for storing and retrieving images, scripts, and middleware provisioning profiles. See [Section 15.6, "Loading Provisioning Profiles to the Software Library"](#) for details.
- **Metering for Chargeback and Reporting:** Resource consumption can be tracked and users can be charged based on different Chargeback plans. The Chargeback feature provides in depth historical and on demand reports.
- **Self Service Provisioning:** Enterprise Manager provides an out-of-the-box portal for cloud users. Self Service users can monitor and administer services with minimal or no administrator intervention.
- **Monitoring and Administering:** The Middleware Self Service Portal provides centralized monitoring of the cloud infrastructure and underlying resources. Users can monitor requests and failure rates and identify potential bottlenecks.

Setting Up Middleware as a Service

This chapter describes to set up middleware as a service. It contains the following sections:

- [Getting Started](#)
- [Defining Roles and Assigning Users](#)
- [Adding Hosts](#)
- [Configuring Privilege Delegation Settings](#)
- [Configuring the Software Library](#)
- [Loading Provisioning Profiles to the Software Library](#)
- [Creating a Middleware Home](#)
- [Configuring and Saving the Deployment Procedure](#)

15.1 Getting Started

This section helps you get started by providing an overview of the steps involved in setting up middleware as a service (MWaaS). Before you set up MWaaS, you must download and deploy the required plug-ins. For more details, see Section 2.1, "Installing Virtualization and the Cloud Plug-ins"

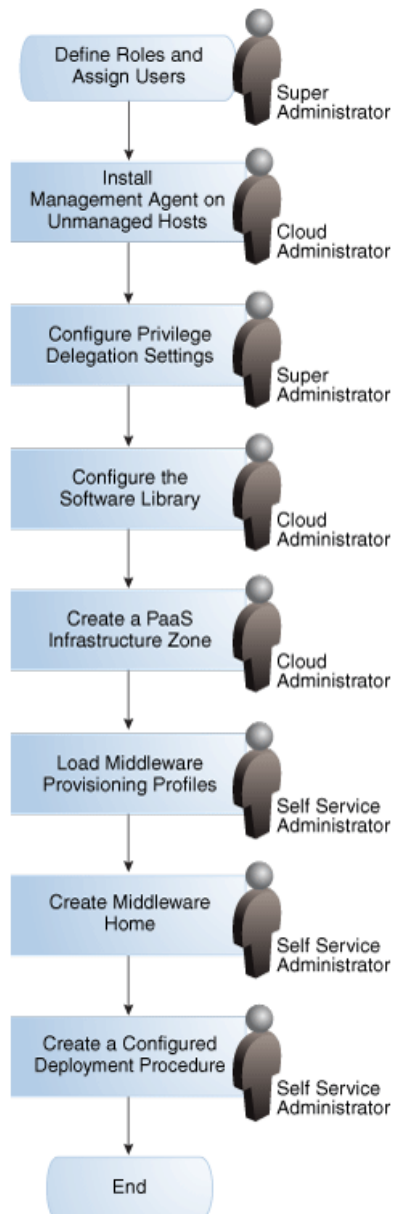
Table 15–1 *Getting Started with MWaaS*

Step	Task	Role
1	Define roles for administrators and self service users. See Section 15.2, "Defining Roles and Assigning Users" .	Super Administrator
2	Install the Management Agent on unmanaged hosts so that they can be monitored by Enterprise Manager. Section 15.3, "Adding Hosts" .	Cloud Administrator
3	Configure Privilege Delegation Settings on your managed hosts. Section 15.4, "Configuring Privilege Delegation Settings" .	Super Administrator
4	Configure the storage location for the Software Library. See Section 15.5, "Configuring the Software Library"	Cloud Administrator
5	Load Middleware Provisioning profiles into the Software Library. See Section 15.6, "Loading Provisioning Profiles to the Software Library" .	Self Service Administrator
6	Create a Middleware Home. See Section 15.7, "Creating a Middleware Home" .	Self Service Administrator

Table 15–1 (Cont.) Getting Started with MWaaS

Step	Task	Role
7	Create a Configured Deployment Procedure. See Section 15.8, "Configuring and Saving the Deployment Procedure" .	Self Service Administrator

Figure 15–1 shows the sequence of steps required to enable middleware as a service.

Figure 15–1 Enabling MWaaS as a Service

15.2 Defining Roles and Assigning Users

Oracle Enterprise Manager supports different classes of Oracle users, depending upon the environment you are managing and the context in which you are using Oracle Enterprise Manager.

User privileges provide a basic level of security in Enterprise Manager. They are designed to control user access to data and to limit the kinds of SQL statements that users can execute. When creating a user, you grant privileges to enable the user to connect to the database, to run queries and make updates, to create schema objects, and more.

A role is a collection of Enterprise Manager resource privileges, or target privileges, or both, which you can grant to administrators or to other roles. These roles can be based upon geographic location (for example, a role for Canadian administrators to manage Canadian systems), line of business (for example, a role for administrators of the human resource systems or the sales systems), or any other model. By creating roles, an administrator needs only to assign the role that includes all the appropriate privileges to his team members instead of having to grant many individual privileges.

To set up a Private Middleware Cloud in Enterprise Manager, you create users for each of these roles:

- **EM_CLOUD_ADMINISTRATOR:** Users with this role are responsible for setting up and managing the cloud infrastructure. This role is responsible for creating the PaaS Infrastructure Zones and for the infrastructure cloud operations for performance and configuration management.
- **EM_SSA_ADMINISTRATOR:** Users have the privilege to define the self service role. They can define quotas and constraints for the self service users and grant them access privileges. Users with this role have provisioning designer privileges, can create and view patch plans, and support the plug-in lifecycle on the Management Agent.
- **EM_SSA_USER:** Users with this role can access the Middleware Self Service Portal.

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

15.2.1 Creating a Custom Role for Self Service Application Users

The Cloud Management Self Service Portal allows self service users to provision and manage their own cloud services. Self service users need access only to the Self Service Portal and the resources allocated to them. These capabilities are inherent in the predefined EM_SSA_USER role.

Since the functions performed by the `EM_CLOUD_ADMINISTRATOR` and `EM_SSA_ADMINISTRATOR` roles are consistent across Enterprise Manager, these ready-to-use roles can be used as they are. On the other hand, the `EM_SSA_USER` role is used for quota assignment, and to limit access to Middleware Pools and Service Templates. Hence the predefined role needs to be modified to create custom SSA user roles that are based on the standard `EM_SSA_ROLE` role.

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.

15.2.2 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.

Note: Repeat these steps to create other users. For the Cloud Administrator and SSA Administrator users, the `EM_USER` and `PUBLIC` roles must not be removed.

15.3 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.

If the host targets have not been discovered, you must install Oracle Management Agents on your unmanaged hosts to monitor them in Enterprise Manager.

Note: If the Management Agent has not been upgraded to 12.1.0.3 then following options will not work:

- Create and Delete Data Sources
 - Start and Stop Applications
 - View List of Data Sources
 - View Application URLs
-
-

To install the Management Agent on an unmanaged server, follow these steps:

1. Log in to Enterprise Manager as an `EM_CLOUD_ADMINISTRATOR` user.
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.

Note:

- You can deploy the Management Agent on a maximum of 16 servers at a time. The total time required will thus increase if the Management Agent is deployed to a large number of servers.
 - Verify if the AS plugin 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**.
-
-

15.4 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.

15.5 Configuring the Software Library

Oracle Software Library (Software Library) is a repository that stores software patches, virtual appliance images, reference gold images, application software and their associated directive scripts. It allows maintaining versions, maturity levels, and states of entities.

The software entities can be automatically mass-deployed to provision software, software updates, and servers using Enterprise Manager Cloud Control in a reliable and repeatable manner. These provisioning operations, which are unattended and can be scheduled, lead to substantial cost savings.

Besides acting as a repository for certified software entities, the Software Library is a logical interface between the deployment models and the automation framework required to perform a large number of patching and provisioning tasks. To configure the storage location for the Software Library, follow these steps:

1. Log in to Enterprise Manager as an `EM_CLOUD_ADMINISTRATOR` user.
2. From the **Setup** menu, select **Provisioning and Patching**, then select **Software Library**.
3. Select **OMS Shared File System** in the Storage Type list and click **Add...**

If you are using the OMS Agent Filesystem location, see the note in [Section 10.1, "Creating a PaaS Infrastructure Zone"](#) for setting up administrator credentials.

4. 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.

5. A job is executed to upload all the ready-to-use content. This may take 15 to 30 minutes depending on your disk speed.

15.6 Loading Provisioning Profiles to the Software Library

The `EM_ADMINISTRATOR` can create provisioning profiles that can be used by operators for mass deployment. Provisioning profiles can be used to standardize deployments and help reduce errors when deployment procedures are configured.

Three out-of-the-box provisioning profiles with different heap size settings (Small: 1.5 GB, Medium: 3 GB, and Large: 6 GB) are provided. To use these profiles, you must import or load them to the Software Library by using the following command:

```
emctl partool deploy -parFile <YOUR_LOCAL_DIR>/<file_name>.par -force
-repPasswd <SYSMAN_PWD>
```

When imported, these profiles are copied under the Oracle Supplied Fusion Middleware Provisioning Profiles/10.3.6.0/linux64 directory in the Software Library. These profiles can then be used to create services templates that can be used provision one or more service instances. See [Section 16.1, "Setting Up the Self Service Portal"](#) for details on creating service templates.

Note:

- The provisioning profiles must be JRF enabled.
 - 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) have been created on the selected destination host(s) when they are provisioned.
 - The maximum heap value of all the servers in the profile must have an appropriate `-Xmx` setting in multiples of 256. For example, `-Xmx512m`, `-Xmx1024m`, 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.
 - Scale up operations are supported only with non-SSL profiles.
-

15.7 Creating a Middleware Home

A Middleware Home consists of the WebLogic Server Home (such as SOA Home, 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 10.1, "Creating a PaaS Infrastructure Zone"](#) for details. If the Middleware Home has not been created on a host, that host cannot be a part of the Middleware Pool.

Note: If you have a Middleware Home installed on your machine and it has been discovered in Enterprise Manager, this step is not required.

To create a Middleware Home in Enterprise Manager, follow these steps:

1. From the **Enterprise** menu, select **Provisioning and Patching**, then select **Middleware Provisioning**. The Middleware Provisioning page appears.
2. The profiles that have been uploaded are listed in the Profile section.
3. Select the Provision Middleware Deployment Procedure from the Deployment Procedures region and click **Launch**.
4. Select **Provision** from Oracle Middleware Home Gold Image option. Click the **Search** icon and select one of the two profiles as the image to use for the oracle home. Click **Next**.
5. In the Destinations page, click **Add Hosts**. Select a host and specify the Host Credentials. You can either use named credentials or specify new credentials.
6. In the Destination Locations section, select **Create New Middleware Home** and specify a non-existing middleware home directory. Click **Next**.
7. Accept the default values in this page and click **Submit**. You will see the Procedure Activity page where you can monitor the job you just submitted to ensure it completes correctly. A middleware home is created in the specified directory.

15.8 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 an `EM_SSA_ADMINISTRATOR` user.
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. Click **Next**.
6. Click **Save**. Enter a name for the customized deployment procedure and click **Next**.

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.

Setting Up the Middleware 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 Self Service Portal](#)
- [Viewing the Middleware Pool Home Page](#)

16.1 Setting Up the Self Service Portal

The EM_SSA_ADMINISTRATOR can configure and set up the middleware self service portal by creating middleware pools, service templates, and so on.

Prerequisites

- MWaaS must be enabled by following the steps described in the [Section 15.1, "Getting Started"](#) section.
- Enterprise Manager is set up to manage the cloud environment.
- One or more PaaS Infrastructure Zones must have been created. See [Section 10.1, "Creating a PaaS Infrastructure Zone"](#) for details.
- A configured Middleware Provisioning Deployment Procedure must have been created.

Follow this procedure to set up the MWaaS Self Service Portal.

1. Log into Enterprise Manager Grid Control as EM_SSA_ADMINISTRATOR.
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. A middleware pool contains a set of resources that can be used to provision a service instance within a PaaS Infrastructure Zone.

Figure 16–1 Middleware Cloud Self Service Portal Setup: Middleware Pools

The screenshot shows the 'Middleware Cloud Self Service Portal Setup' page. On the left is a navigation menu with 'Middleware Pools' selected. The main content area shows a table of existing middleware pools. The table has columns for Pool Name, Owner, PaaS Infrastructure Zone, Member Target Type, Targets, and Description. One pool named 'example_pool' is listed with owner 'SSAADMIN', zone 'my_paas_zone', target type 'Oracle Home', and 1 target.

Pool Name	Owner	PaaS Infrastructure Zone	Member Target Type	Targets	Description
example_pool	SSAADMIN	my_paas_zone	Oracle Home	1	

4. Click **Create**. The Create Software 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 WebLogic Server instances that can be running on each host in the pool.

Figure 16–2 Create Software Pool: General

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 10.1, "Creating a PaaS Infrastructure Zone"](#) for details.
 2. Specify the Version of the Oracle home being added to the middleware pool.
 3. Click **Add** and select one or more Oracle homes to be added to the middleware pool. Click **Submit** to create a software pool. The newly created pool will appear in the Middleware Cloud Self Service Portal Setup: Middleware Pools page. The middleware pool can now be added to the Service Template.
7. Click the **Request Settings** tab. Specify the following and click **Apply**.
 - **Future Reservation:** How far in advance can a request be made.
 - **Request Archive Retention:** Period for which archived requests are stored before they are automatically deleted by the system.
 - **Request Duration:** Maximum duration for which requests can be made.
8. 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 16–3 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

9. 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.

10. 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.
11. Click **Create** to launch the Create New Service Template wizard. The Create Service Template: Name page appears.

Figure 16–4 Create Service Template: General

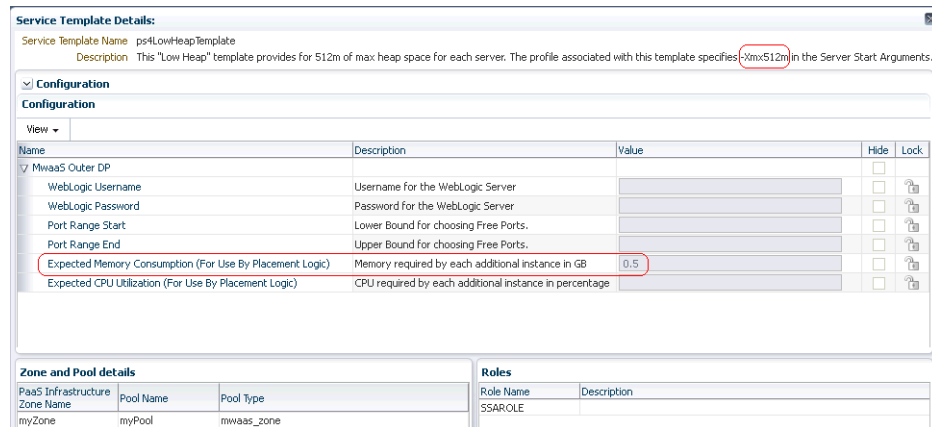
12. Enter a name and description for the service template. The description must be unique and provide information on the type of service template being created.
13. Click the **Search** icon next to the Deployment Procedure field and select a Configured Middleware Provisioning deployment procedure that 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 Configure Middleware Provisioning Deployment Procedure with appropriate values.

14. In the Create New Service Template: Configuration page, enter the WebLogic User Name, Password, and Port Range. The values specified here will be used to provision new WebLogic Server instances. If these values are not specified, default values will be used. It is recommended that you enter appropriate values for the user name and password fields.

Note: You cannot lock any of the fields on this page.

Figure 16–5 Create Service Template: Configuration



The Expected Memory Consumption (For Use By Placement Logic) 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 15.6, "Loading Provisioning Profiles to the Software Library"](#) for more details.

15. Click **Next**. The Create New Service Template: PaaS Infrastructure Zones page appears.

Figure 16–6 Create Service Template: PaaS Infrastructure Zones



16. Click **Add** to select the zones into which that the service instances can be provisioned.
17. Click the **Search** icon in the Software Pool field and select Software Pool that contains 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. Click **Next**.
18. 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.
19. 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.
20. 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 the Chargeback is setup, you can view charges and plan details. For more details on setting up the Chargeback Service, see [Chapter 4, "Chargeback Administration"](#).

16.1.1 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 10.1, "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. Associate the *west coast pool* with the *west coast zone*.
4. Create a service template and associate the west coast zone and the west coast pool. 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.

16.2 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 16–7 *Middleware Pool Home Page*

The screenshot displays the 'ps4_pool' Middleware Pool Home Page. The page is titled 'ps4_pool' and 'Middleware Pool'. It shows the page was refreshed on Jun 25, 2012 at 7:40:34 PM GMT-07:00.

The 'General' section is expanded, showing a 'Summary' table with the following information:

Summary	Placement Policy Constraints	Java Servers
Description: PS4 OH Pool on [redacted]	Maximum Number of Java Servers (per Host): 10	Total Java Servers: 5
Parent PaaS Infrastructure Zone: [redacted]		
Oracle Homes: 3		

The 'Members' section is also expanded, showing a table of members:

Name	Type	Host	Home Location	Middleware Home
WebLogicServer10_3_5_0 [redacted]	[Home icon]	[redacted]	[redacted]	[redacted]
WebLogicServer10_3_5_0 [redacted]	[Home icon]	[redacted]	[redacted]	[redacted]
WebLogicServer10_3_5_0 [redacted]	[Home icon]	[redacted]	[redacted]	[redacted]

- **General:** This region shows the following information:
 - **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.
 - **Total Java Servers:** The total number of WebLogic servers in the middleware pool.
 - **Maximum No of Java Servers (per host):** The maximum number of WebLogic servers that can be running on each host.
- **Incidents and Problems:** Shows the outstanding incidents and problems for the target.
- **Service Template Summary:** The number of service templates with which the software pool is associated.
- **CPU and Memory:** These charts show the average CPU and memory utilizations by all servers (virtual servers and physical hosts) in the middleware pool.
- **Members:** This region shows a list of members in the middleware pool. The name of the target, the target type (Oracle Home), the host on which it is running, the location of the Oracle Home, and the location of the Middleware Home is displayed.

Using the MWaaS Self Service Portal

This chapter describes the procedures you need to follow to monitor and administer the MWaaS Self Service Portal. It contains the following sections:

- [Using the MWaaS Self Service Portal](#)
- [Requesting a 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 Self Service (SSA) users to view, monitor, and request middleware services. SSA users 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 Self Service Administrator.

To view the MWaaS Self Service Portal, follow these steps:

1. Log into Enterprise Manager as a user with `EM_SSA_USER` privileges.
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.
 - **Your Usage:** The total quota (servers and memory) used so far by the SSA user, 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 Service"](#) for details. Click on the **Name** link to drill down to the Request 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 **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**.

17.2 Requesting a 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 the SSA user can access is displayed. Choose a service template and click **Select** to launch the New Service Request page.

Figure 17–2 New Service Request Page

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 lowHeapTemplate
Description Use this template to create "low heap" (500 MB) Middleware Services

* Service Name

* PaaS Infrastructure Zone

Schedule

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

End Date Indefinitely Until

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.

Application Name

Application Component

Deployment Plan

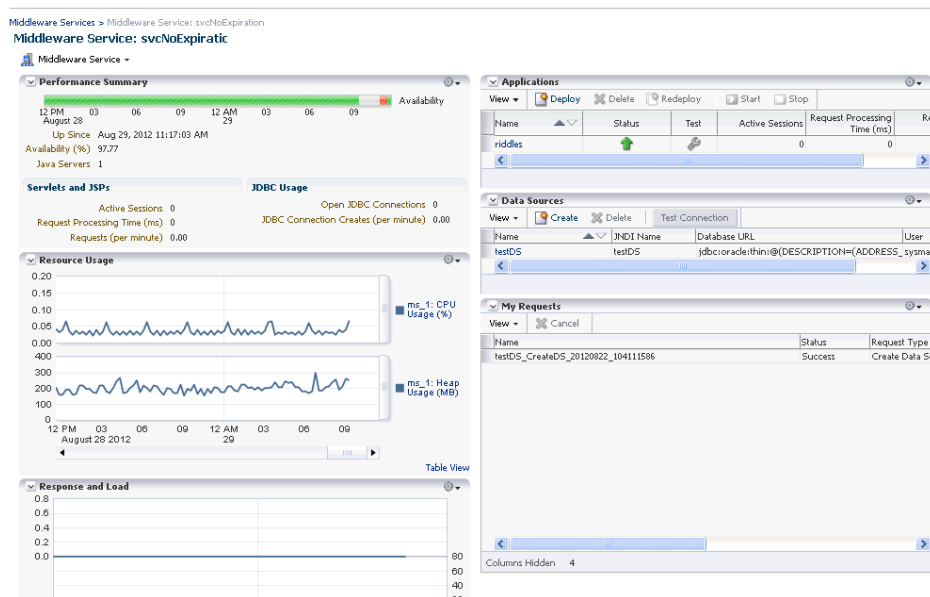
3. Enter a name for the service and select the PaaS Infrastructure Zone in which the new service is to be deployed.
4. You can specify the schedule for the service request.
 - **Start Date:** The service can be started **Immediately** or **Later**. If you select the Later option, specify the date from 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 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
 - **Deployment Plan:** Select an xml file that contains the deployment options. This file overrides the deployment plan (if any) specified with the application component.

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 on the **Name** link in the Middleware Services region in the Middleware Cloud Self Service Portal: Home tab.

Figure 17–3 *Middleware Service Home Page*



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 `ser_1`, after scale up, the subsequent servers are named as `ser_2`, `ser_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 on 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 Deploying 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.
 - **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, JNDI name and the URL for the database is displayed. 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, 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 `EM_SSA_USER` privileges 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. Select an application from the Applications table and click **Deploy**. The Deploy Application page appears.

Figure 17–4 Deploy Application Page

Middleware Cloud Self Service Portal

Deploy Application Cancel Deploy

Specify a unique name for your application.

* Application Name | my_application

Select an application component from the Software Library.

* Application Component | Riddles

Deployment Plan |

4. Enter a name for the application. Click the **Search** icon next to the Application Component field and select a Java EE application component from the Software Library.
5. 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.

- 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 or 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.

- Log into Enterprise Manager as a user with `EM_SSA_USER` privileges and navigate to the Middleware Cloud Self Service Portal page.
- Click **Create** in the Data Sources table in the Middleware Service Home page. The Create Data Source page appears.

Figure 17–5 Create Data Source

The screenshot shows the 'Create Data Source' page in the Middleware Cloud Self Service Portal. The page title is 'Middleware Cloud Self Service Portal' and the sub-header is 'Create Data Source'. Below the sub-header, there is a description: 'Use this page to create a data source on this middleware service.' The form contains the following fields:

- Data Source Name:** my_data_source
- JNDI Name:** (empty)
- Database Driver Class:** com.ddtek.jdbc.oracle.OracleDriver (selected from a dropdown menu)
- Database URL:** (empty)
- Database User Name:** (empty)
- Password:** (empty)
- Confirm Password:** (empty)

At the bottom right of the form is a 'Test Connection' button. At the top right of the page are 'Cancel' and 'Create' buttons.

- Enter the following details:

- Data Source Name
- JNDI Name
- Database Driver Class
- Database URL
- Database User Name
- Password

4. 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 17–6 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 17–7 MWaaS Application Home Page



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.

Part VI

Cloud APIs

This section describes the cloud application programming interfaces (APIs) and the RESTful (Representational State Transfer) API.

It contains the following chapters:

- [Chapter 18, "Introduction to Cloud APIs"](#)
- [Chapter 19, "Cloud Resource Models"](#)
- [Section 20, "Cloud API Examples"](#)
- [Chapter 22, "Chargeback and Metering EMCLI Verbs"](#)

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](#)
- [Common Behaviors](#)

18.1 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.

18.2 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.

18.3 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.

18.3.1 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.

18.3.2 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.

18.3.3 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.

18.3.4 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 18–1 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.

18.3.5 Response Headers

Specific HTTP headers included in the responses returned from the platform are described in the following table:

Table 18–2 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 18–2 (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. This will never be returned on a privileged resource or a resource request that contains an authorization header.

18.3.6 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 18–3 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 18–3 (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. 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.

18.3.7 Common Resource Attributes

All the resource entities in this specification may contain the following common resource attributes.

18.3.7.1 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 18–4 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). When there are vendor extensions, the vendor shall publish and document their semantics.

Table 18–4 (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.

18.3.8 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 18–5 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). 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>>.

18.3.9 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 18–6 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 18–7 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. 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](#)
- [Introduction](#)
- [Resource Data Models](#)

19.1 Enterprise Manager Cloud Resource Model 10001

This version of the Cloud resource model is the culmination of the XaaS implementation for SSA where additional services can be plugged into SSA to support the creation and lifecycle management of service instances.

For the Enterprise Manager for Oracle Cloud 12.1.0.4 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 19–1 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 19–1 (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).

19.2 Introduction

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.Xxxxxxx+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.

Table 19–2 (Cont.) Cloud Data Model

Field	Type	Occurs	Description
service_family_types	Collection <ServiceFamilyType>	0..1	The list of service family types that are supported by the Cloud.
zones	Collection<Zone>	0..1	List of zones that are supported by the cloud and accessible to the user.
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.
media_type	String	1	Value of this media type with the payload format. For example, application/oracle.com.cloud.common.Cloud+json.

19.3.2 ServiceTemplate

A ServiceTemplate represents the definition of the deployable service and has the following media type and payload:

- **Media Type:** application/oracle.com.cloud.common.ServiceTemplate
- **Supported Payload:** json

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 19–3 ServiceTemplate Data Model

Field	Type	Occurs	Description
uri	URI	1	GET against this URI refreshes the client representation of the ServiceTemplate definition to this user.
name	String	1	Human readable name given to the ServiceTemplate. It is a UNICODE string to support different languages.
description	String	0..1	A brief description given to the ServiceTemplate. It is a UNICODE string to support different languages.
type	String	1	The string that describes the type of the service template.
created	Timestamp	1	Date and time, in ISO 8601 format, when the ServiceTemplate was created.
service_family_type	String	0..1	The name of the service family type under which this service template is categorized.

Table 19–3 (Cont.) ServiceTemplate Data Model

Field	Type	Occurs	Description
default_instance_media_type	String	0..1	The default media type of the service instance that can be created using this template. For example, a VMTemplate would have VM for this attribute to indicate that VMTemplate can be deployed into a VM service instance.
resource_state	ResourceState	1	Only a service template with a READY state can be deployed.
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
zones	List of Zones	0..1	Contains the list of Zone resources that this service template can be used to create service instances with
service_instances	Collection<? >	0..1	Contains the list of service instances that are created with this template.

The following table describes the structure of the Deployment Parameters.

Table 19–4 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.
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.

19.3.3 Zone

A Zone represents a logical boundary where the resources may reside and has the following media type and payload:

- **Media Type:** application/oracle.com.cloud.common.Zone
- **Supported Payload:** json

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 19–5 Zone Data Model

Field	Type	Occurs	Description
uri	URI	1	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.
media_type	String	1	The media type of the resource.
resource_state	ResourceState	0..1	The resource state of the resource.
service_family_type	String	0..1	The name of the service family type that this zone is associated with. Each Zone is contextualized in the service family type.
service_templates	Collection <ServiceTemplate>	0..1	Collection of the service templates that this zone supports and can be deployed into.
service_instances	Collection <ServiceInstance>	0..1	Collection of the service instances that are in this zone.

19.3.4 Service Family Type

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. Service Family Type has the following media type and payload:

- **Media Type:** application/oracle.com.cloud.common.ServiceFamilyType
- **Supported Payload:** json

The following table describes the ServiceFamilyType Data Model

Table 19–6 Service Family Type

Field	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the client representation of the ServiceFamilyType definition to this user.
name	String	1	Name of the ServiceFamilyType.
description	String	0..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 <ServiceTemplate>	0..1	The collection of service templates that are of this service family type.
zones	Collection<Zone>	0..1	The collection of zones that support service instances of this service family type.

Table 19–6 (Cont.) Service Family Type

Field	Type	Occurs	Description
instance_types	Collection <InstanceType>	0..1	The collection of instance types that are supported by this service family type.

19.3.5 Service Instance Type

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 Service Instance Type has the following media type and payload:

- **Media Type:** application/oracle.com.cloud.common.InstanceType
- **Supported Payload:** json

The following table describes the ServiceFamilyType Data Model.

Table 19–7 Service Instance Type

Field	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the client representation of the ServiceInstanceType definition to this user.
name	String	1	Name of the ServiceInstanceType.
description	String	0..1	Human readable description of the Service Instance 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.
instance_media_type	String	0..1	The media type of the instance of this type.
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.

19.3.6 Metric

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 Metric has the following media type and payload:

- **Media Type:** application/oracle.com.cloud.common.Metric
- **Supported Payload:** json

The following table describes the Metric Data Model.

Table 19–8 Metric Data Model

Field	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the client representation of the Metric definition to this user.

Table 19–8 (Cont.) Metric Data Model

Field	Type	Occurs	Description
name	String	1	Name of the Metric.
description	String	0..1	Human readable description of the Metric. 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.
type	String	1	The type of the metric value. This is an enumeration of STRING, INTEGER, PERCENTAGE, NUMBER, MIX.
current_value	String	0..1	The latest known value of the metric
current_time	ISO 8601 Date	0..1	The UTC time when the current_value was last observed
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
rollup_unit	String	0..1	The rollup unit for the time series data. Could be one of DAY, DAILY, DAYLY, HOUR, HOURLY, RAW
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. 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.
time_range_epoch	List of 2 Number	0..1	Same as time_range, but in the form of epoch time.

19.3.7 Service Instance

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 Service Instance has the following media type and payload

- **Media Type:** application/oracle.com.cloud.common.ServiceInstance
- **Supported Payload:** json

The following table describes the Service Instance Data Model.

Table 19–9 Service Instance Data Model

Field	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the client representation of the ServiceInstance definition to this user.
name	String	1	Name of the ServiceInstance.
description	String	0..1	Human readable description of the Metric. Shall be a UNICODE string to support different languages.

Table 19–9 (Cont.) Service Instance Data Model

Field	Type	Occurs	Description
resource_state	ResourceState	0..1	The resource state of the resource.
media_type	String	1	The media type of the resource.
metrics	Collection <Metric>	0..1	Collection of metrics that are observed on the resource.
service_family_type	String	1	The name of the ServiceFamilyType that this service instance is in context of.
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.
created	ISO 8601 Date	0..1	The date of which the service instance is created.

19.4 Resource Models for Infrastructure as a Service

The following describes the resource models that are supported by the Infrastructure as a Service family type.

19.4.1 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 19–10 IaaS 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 The media type of the resource. "application/oracle.com.cloud.iaas.Zone"
resource_state	ResourceState	0..1	The resource state of the resource.
service_family_type	String	0..1	The name of the service family type that this zone is associated with. It is "iaas" for this.
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
service_instances	Collection <ServiceInstance>	0..1	Collection of the service instances that are in this zone

Table 19–10 (Cont.) IaaS Zone Data Model

Field	Type	Occurs	Description
vnets	Collection <VNet>	0..1	Collection of the virtual networks that are associated with this zone.
media_type	String	1	The media type of the resource. "application/oracle.com.cloud.iaas.Zone"

19.4.2 AssemblyInstance

Extends: application/oracle.com.cloud.common.ServiceInstance

An AssemblyInstance is a logical grouping of resources from a deployment request of an Assembly template and has the following media type and payload:

- **Media Type:** application/oracle.com.cloud.common.AssemblyInstance
- **Supported Payload:** json

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 19–11 AssemblyInstance Data Model

Field	Type	Occurs	Description
uri	URI	1	GET against this URI refreshes the client representation of the Assembly Instance definition.
name	String	1	Name of the Assembly Instance as given by the user. It is a UNICODE string to support different languages.
description	String	0..1	Human readable description of the Assembly Instance given by the user. It is a UNICODE string to support different languages.
based_on	URI	0..1	URI of the service template of which this AssemblyInstance is based on.
scalability_groups	Collection< ScalabilityGroup>	0..1	List of scalability groups that are included in this Assembly Instance.
servers	Collection< VM>	0..1	List of VMs that are directly included in this Assembly Instance.
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.

Table 19–11 (Cont.) AssemblyInstance Data Model

Field	Type	Occurs	Description
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.
created	Timestamp	0..1	Date and time, in ISO 8601 format, when the Assembly Instance was created.
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.
params	{}	0..1	Vendor specific configuration parameters for this deployment.
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.
media_type	String	1	The media type of the resource.

19.4.3 Scalability Group

Extends: application/oracle.com.cloud.common.ServiceInstance

The Scalability Group is a collection of servers and corresponding virtual networks, and has the following media type and platform:

- **Media Type:** application/oracle.com.cloud.common.ScalabilityGroup
- **Supported Payload:** json

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 Scalability Group Data Model.

Table 19–12 ScalabilityGroup Data Model

Field	Type	Occurs	Description
uri	URI	1	GET against this URI refreshes the client representation of the scalability group definition.
name	String	1	Name of the scalability group as given by the user. It is a UNICODE string to support different languages.
description	String	0..1	Human readable description of the scalability group given by the user. It is a UNICODE string to support different languages.
nodes	Collection<URI>	1	List of URIs that represent the entities making up this scalability group.
count	Integer	1	Count of the nodes that are in the scalability group.

Table 19–12 (Cont.) ScalabilityGroup Data Model

Field	Type	Occurs	Description
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.
resource_state	ResourceState	1	Validity of the other scalability group fields is guaranteed only when the resource state is READY. Otherwise, the client shall not assume the validity of the fields.
created	Timestamp	1	Date and time, in ISO 8601 format, when the scalability group is created.
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".
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.
homogenous	BOOLEAN	1	TRUE when the scalability group contains homogenous entities and FALSE otherwise.

19.4.4 VM

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 and has the following media type and platform:

- **Media Type:** application/oracle.com.cloud.common.VM
- **Supported Payload:** json

The following table describes the Server Data Model.

Table 19–13 VM Data Model

Field	Type	Occurs	Description
uri	URI	1	GET against this URI refreshes the client representation of the VM definition.
name	String	1	Name of the VM as given by the user or generated by the platform.
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.
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.
based_on	URI	0..1	URI of the service template on which this Server is based.
hostname	String	0..1	Qualified host name of this Server when available.

Table 19–13 (Cont.) VM Data Model

Field	Type	Occurs	Description
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.
memory	Number	1	Numeric sizing of the RAM in MByte.
disks	{String,Number}[]	0..1	Name and size in GB of local disks.
interfaces	Collection<NetworkInterface>	1	Network interfaces associated with this Server.
params	{ }	0..1	Vendor specific configuration parameters for this Server.
resource_state	ResourceState	1	Validity of the other Server fields is guaranteed only when the resource state is READY. Otherwise, the client should not assume the validity of the fields.
created	Timestamp	1	Date and time, in ISO 8601 format, when the Server was created.
type	String	1	Type of server that represents this computational container (PHYSICAL, VIRTUAL).
os	String	0..1	OS information of the Server where appropriate.
media_type	String	1	The media type of the resource.

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

19.4.5 VNet

A VNet is a service that is capable of providing network addresses, routing rules, security constraints, and access limits. It has the following media type and platform:

- **Media Type:** application/oracle.com.cloud.common.VNet
- **Supported Payload:** json

The following table describes the VNet Data Model.

Table 19–14 VNet Data Model

Field	Type	Occurs	Description
uri	URI	1	GET against this URI refreshes the client representation of the VNet definition.
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.
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.
created	Timestamp	1	Date and time, in ISO 8601 format, when this VNet was created.
interfaces	Collection<NetworkInterface>	0..1	List of NetworkInterface resources that are part of the VNet.

19.4.6 NetworkInterface

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. It has the following media type and platform:

- **Media Type:** application/oracle.com.cloud.common.NetworkInterface
- **Supported Payload:** json

The following table describes the NetworkInterface Data Model.

Table 19–15 NetworkInterface Data Model

Field	Type	Occurs	Description
uri	URI	1	GET against this URI refreshes the client representation of the NetworkInterface definition.
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.
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.
address	String	1	Address of this interface. If an IP based protocol, specified as a dotted notation IPv4 or IPv6 address.
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.
mac	String	0..1	The MAC address of the network interface where appropriate.

19.4.7 VMTemplate

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. It has the following media type and platform:

- **Media Type:** application/oracle.com.cloud.common.VMTemplate
- **Supported Payload:** json

VMTemplate is a preconfigured deployable entity that realizes a VM resource.

Table 19–16 VMTemplate Data Model

Field	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the client representation of the VMTemplate definition to this user.
name	String	1	A human readable name, given to the VMTemplate.
description	String	0..1	A brief description, given to the ServiceTemplate. SHALL be a UNICODE string to support different languages.
type	String	1	The String that describes the type of the VM Template. This has value "Template".
created	Timestamp	1	Date and time, in ISO 8601 format, when the VMTemplate was created.
service_family_type	String	0..1	The name of the service family type that this service template is categorized under. This has value "iaas".
default_instance_media_type	String	0..1	The default media type of the service instance that can be created using this template. This has value "application/oracle.com.cloud.common.VM".
resource_state	ResourceState	0..1	Only a service template with READY state can be deployed.
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.
zones	List of Zones	0..1	Contains the list of Zone resources that this service template can be used to create service instances with.
service_instances	Collection<?>	0..1	Contains the list of service instances that are created with this template.

19.4.8 AssemblyTemplate

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. It has the following media type and platform:

- **Media Type:** application/oracle.com.cloud.common.AssemblyTemplate
- **Supported Payload:** json

AssemblyTemplate is a deployable entity that realizes an AssemblyInstance resource that may contain multiple resources that are interconnected.

Table 19–17 AssemblyTemplate Data Model

Field	Type	Occurs	Description
uri	URI	1	A GET against this URI refreshes the client representation of the AssemblyTemplate definition to this user.
name	String	1	A human readable name, given to the AssemblyTemplate.
description	String	0..1	A brief description, given to the AssemblyTemplate. SHALL be a UNICODE string to support different languages.
type	String	1	The String that describes the type of the Assembly Template. This has value "Assembly".
created	Timestamp	1	Date and time, in ISO 8601 format, when the AssemblyTemplate was created.
service_family_type	String	0..1	The name of the service family type that this service template is categorized under. This has value "iaas".
default_instance_media_type	String	0..1	The default media type of the service instance that can be created using this template. This has value "application/oracle.com.cloud.common.AssemblyInstance".
resource_state	ResourceState	0..1	Only a service template with READY state can be deployed.
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.
zones	List of Zones	0..1	Contains the list of Zone resources that this service template can be used to create service instances with.
service_instances	Collection<?>	0..1	Contains the list of service instances that are created with this template.
definition	String	0..1	Contains the definition of the assembly template. This would be the OVF xml metadata representing the assembly template.
ovfdescriptor	String	0..1	Same as definition.
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.

The following table describes the structure of the Deployment Plan.

Table 19–18 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.

Table 19–18 (Cont.) Deployment Plan Structure

Field	Type	Occurs	Description
xml	String	1	The deployment plan XML.

Cloud API Examples

The following sections provide examples of interaction with Enterprise Manager Cloud Control 12.1.0.2 with the Enterprise Manager for Oracle Cloud 12.1.0.4 plug-in. The following topics are covered:

- [Introspecting Resources](#)
- [Creating a Service Instance](#)
- [Updating a Service Instance](#)
- [Deleting a Service Instance](#)
- [Listing Service Instances](#)
- [Version 10000 Support](#)

20.1 Introspecting Resources

The following examples illustrates the shape of the various resources. In this version, default x-specification-version is assumed at 10001.

20.1.1 Cloud Resource

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" ,
  "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%3Aacmp%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%3Aprovisioning%3A1%3Aacmp%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%3Aprovisioning%3A1%3Aacmp%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 as the "/em/cloud" address uniquely identifies the Cloud as the default resource to be returned.

The following shows interaction with the Cloud resource picking up specific attributes

- URL
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"
    }
  ]
}

```

```

    } ,
    {
      "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 will be gathered and returned to the client.

20.1.2 Service Family Type Resource

The following describes the Service Family Type resource.

- URL
https://example.us.oracle.com/em/cloud/service_family_type/iaas
- 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.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%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%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%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"
        } ,
        {
            "uri" : "/em/cloud/instance_type/assembly%40iaas" ,
            "name" : "assembly" ,
            "media_type" : "application/oracle.com.cloud.common.InstanceType+json"
        } ,
        {
            "uri" : "/em/cloud/instance_type/assemblyinstance%40iaas" ,
            "name" : "assemblyinstance" ,
            "media_type" : "application/oracle.com.cloud.common.InstanceType+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 media type `application/oracle.com.cloud.iaas.IaaSServiceFamilyType+json` is returned automatically as the most detailed resource. Similarly, you can interact with the system by accepting the common type:

- URL
 - `https://example.us.oracle.com/em/cloud/service_family_type/iaas`
- Headers
 - Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=
 - Accept-Type: application/oracle.com.cloud.common.ServiceFamilyType+json
- Method
 - GET

The Web service will return the same as described above.

20.1.3 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/vm%40iaas`
- Headers
 - Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=
- Method
 - GET

The following shows 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"
    }
  ]
}

```

20.1.4 Zone

This describes the Zone resource.

- URL

<https://example.us.oracle.com/em/cloud/iaas/zone/D73AF0D42C8459E11419862797D1F37D>

- Headers

Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=

- Method

GET

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" :
"oracle:defaultService:em:provisioning:1:netConfig:C76F741AFD7EB760E040578CDA812CD8" ,
        "media_type" : "application/oracle.com.cloud.common.VNet+json" ,
        "id" :
"oracle:defaultService:em:provisioning:1:netConfig:C76F741AFD7EB760E040578CDA812CD8"
    }
} ,
"delegated_credentials" : {
    "media_type" :
"application/oracle.com.cloud.iaas.ExalogicZoneCredential+json" ,
    "total" : "0" ,
    "elements" :
    [
    ]
}
}

```

20.1.5 Service 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/iaas/servicetemplate/vm/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp%3AVirtualization%3ATemplate%3AC76CEB5563EA5E13E040578CDA817FAF%3A0.1>
- Headers

Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=
- Method

GET

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" :
"oracle:defaultService:em:provisioning:1:cmp:Virtualization:Template:C76CEB5563EA5E13E040578CDA817FAF:0.1" ,
  "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"
}

```

20.1.6 Metric resource

A metric contains some time series information, and depending on the type of the metric, rollout information can be retrieved. For example,

- URL
 - https://example.oracle.com/em/cloud/metric/iaas%3Avm%3ACPU_Utilization%3A523CAE80A305928C9C5BE8A67A4181FD
- Headers
 - Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=
- Method
 - GET

This 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:

- URL


```
https://example.oracle.com/em/cloud/metric/iaas%3Avm%3ACPU_
Utilization%3A523CAE80A305928C9C5BE8A67A4181FD?time_range_
epoch=1345595894965~1345598894965
```
- Headers


```
Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=
```
- Method


```
GET
```

The query parameter "time_range_epoch" contains two 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 different rollup unit, the following interaction shows that HOURLY rollup for the last 7 hours:

- URL


```

        "min" : "0" ,
        "max" : "52.8489990234375" ,
        "std" : "1.80268512452981440717409537961069439254E01"
    } ,
    {
        "time_utc" : "2012-08-21T19:00:00+0000" ,
        "average" : "0" ,
        "min" : "0" ,
        "max" : "0" ,
        "std" : "0"
    } ,
    {
        "time_utc" : "2012-08-21T18:00:00+0000" ,
        "average" : "0" ,
        "min" : "0" ,
        "max" : "0" ,
        "std" : "0"
    }
    ] ,
    "key" : "iaas:vm:CPU_Utilization:523CAE80A305928C9C5BE8A67A4181FD"
}

```

20.2 Creating a Service Instance

This section describes how to create a service instance using the Cloud resources.

20.2.1 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.

20.2.1.1 POST to the "iaas" Zone

The following shows the configuration for POSTing to "iaas" Zone:

- 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=Content-Type:
application/oracle.com.cloud.common.VM+json

```

- Body

```

{
  "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"
        }
    }
}

```

- Method

POST

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"
}

```

20.2.1.2 POST to the VMTemplate

The following describes the configurations for POSTing to the VMTemplate:

- 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=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

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:

- 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=Content-Type: application/oracle.com.cloud.common.VM+json

- Body

```
{
  "cpu" : [2,0],
  "memory" : "512",
  "params":{
    "server_prefix":"STPOST",
    "vnc_password":"welcome1",
    "root_password":"welcome1"
  }
}
```

- 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.RestControllerException: '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
t
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.AbstractRestController.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"
  }
]
}

```

20.2.1.3 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.

- 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" : "dadvm0052.us.oracle.com" ,
"cpu" :
[
  "2" ,
  "0"
] ,
"memory" : "512" ,
"disks" :
[
  [
    "default_disk" ,
    "10240"
  ] ,
  [
    "7e8a5d25-7995-40c9-9b16-9274cc2d64f3.img (2)" ,
    "6918"
  ] ,
  [
    "fcdf377c-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"
}

```

20.2.2 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:

- URL
 - <https://example.oracle.com/em/cloud/iaas/servicetemplate/assembly/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp%3AVirtualization%3AAsssembly%3AC75E88B04D7FEDEDE040578CDA810E49%3A0.1>
- Headers
 - Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=
- Method
 - GET

The following is the response from the Web service:

```
{
  "uri" :
"/em/cloud/iaas/servicetemplate/assembly/oracle%3AdefaultService%3Aem%3Aprovisioni
ng%3A1%3Acmp%3AVirtualization%3AAssembly%3AC75E88B04D7FEDEDE040578CDA810E49%3A0.1"
,
  "name" : "sidb_assembly" ,
  "resource_state" : {
    "state" : "READY"
  } ,
  "context_id" :
"oracle:defaultService:em:provisioning:1:cmp:Virtualization:Assembly:C75E88B04D7FE
DEDE040578CDA810E49:0.1" ,
  "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-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<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
</disk>\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%3Aacmp%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:

- URL

```

https://example.oracle.com/em/cloud/iaas/servicetemplate/assembly/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Aacmp%3AVirtualization%3AAssembly%3AC75E88B04D7FEDEDE040578CDA810E49%3A0.1

```

- Headers

Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=Content-Type: application/oracle.com.cloud.common.AssemblyInstance+json

- 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'>\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:C75E
88B04D7FEDEDE040578CDA810E49:0.1'>\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>fals
e</Required>\n<Value>>false</Value>\n<Secret>>false</Secret>\n
</Property>\n
<Property
id='ocm.metalinkCsiRegistration.CSI'>\n<Name>ocm.metalinkCsiRegistration.CSI<
/Name>\n<Required>>false</Required>\n<Value></Value>\n<Secret>>false</Secret>\n
</Property>\n
<Property
id='ocm.metalinkCsiRegistration.metalinkId'>\n<Name>ocm.metalinkCsiRegistrati
on.metalinkId</Name>\n<Required>>false</Required>\n<Value></Value>\n<Secret>fals
e</Secret>\n
</Property>\n
<Property
id='ocm.proxyPassword'>\n<Name>ocm.proxyPassword</Name>\n<Required>>false</Req
uired>\n<Value>
</Value>\n<Secret>>true</Secret>\n
</Property>\n
```



```

<Property
id=\"ocm.metalinkCsiRegistration.countryCode\">\n<Name>ocm.metalinkCsiRegistrat
ion.countryCode</Name>\n<Required>>false</Required>\n<Value></Value>\n<Secret>fa
lse</Secret>\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>fal
se</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>fal
se</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>fal
se</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</
Secret>\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</
Secret>\n
</Property>\n
<Property
id=\"asm-password\">\n<Name>asm-password</Name>\n<Required>>true</Required>\n<Va
lue>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</
Secret>\n
</Property>\n
<Property
id=\"ocm.metalinkEmailRegistration.metalinkEmailId\">\n<Name>ocm.metalinkEmailR
egistration.metalinkEmailId</Name>\n<Required>>false</Required>\n<Value></Value>
\n<Secret>false</Secret>\n
</Property>\n
<Property
id=\"input|listener-1|global-db-name\">\n<Name>input|listener-1|global-db-name<
/Name>\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>fal
se</Secret>\n
</Property>\n
<Property
id=\"ocm.proxyHost\">\n<Name>ocm.proxyHost</Name>\n<Required>>false</Required>\n
<Value></Value>\n<Secret>false</Secret>\n
</Property>\n
<Property
id=\"ocm.repeaterURI\">\n<Name>ocm.repeaterURI</Name>\n<Required>>false</Require

```

```

d>\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>fal
se</Secret>\n
</Property>\n
<Property
id=\"input|listener-1|port\">\n<Name>input|listener-1|port</Name>\n<Required>fa
lse</Required>\n<Value>1521</Value>\n<Secret>>false</Secret>\n
</Property>\n
<Property
id=\"ocm.anonymousEmailRegistration.emailId\">\n<Name>ocm.anonymousEmailRegistr
ation.emailId</Name>\n<Required>>false</Required>\n<Value></Value>\n<Secret>fals
e</Secret>\n
</Property>\n
<Property
id=\"db-account-password\">\n<Name>db-account-password</Name>\n<Required>>true</
Required>\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</
Secret>\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</
Secret>\n
</Property>\n
<Property
id=\"ocm.proxyPort\">\n<Name>ocm.proxyPort</Name>\n<Required>>false</Required>\n
<Value></Value>\n<Secret>>false</Secret>\n
</Property>\n
<Property
id=\"ocm.proxyUsername\">\n<Name>ocm.proxyUsername</Name>\n<Required>>false</Req
uired>\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</
Secret>\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</
Secret>\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</
Secret>\n
</Property>\n
<Property
id=\"ocm.metalinkEmailRegistration.metalinkPassword\">\n<Name>ocm.metalinkEmail
Registration.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"
}

```

- Method

POST

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:

- URL

<https://example.oracle.com/em/cloud/iaas/assemblyrequest/223>

- Headers

Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=

- Method

GET

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.

- URL

<https://example.oracle.com/em/cloud/iaas/assemblyrequest/223>

- Headers

Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=

- Method

GET

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
<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-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<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=\"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"
}

```

20.3 Updating a Service Instance

A service instance may be resized, its status updated, or additional resources added.

20.3.1 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.

20.3.1.1 Change VM Status

The following shows the GET of the VM's status:

- URL
`https://example.oracle.com/em/cloud/iaas/server/byrequest/101?status,uri`
- Headers
Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=
- Method
GET

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.

- URL
`https://example.oracle.com/em/cloud/iaas/server/vm/C5CA59AF9386975E8AEA45B0F040F095`
- Headers
Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=Content-Type: application/oracle.com.cloud.common.VM+json
- Body

```
{
  "status" : "STOPPED"
}
```
- Method
PUT

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:

- URL
`https://example.oracle.com/em/cloud/iaas/server/byrequest/101?status,uri`
- Headers
Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=
- Method
GET

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:

- URL
`https://example.oracle.com/em/cloud/iaas/server/vm/C5CA59AF9386975E8AEA45B0F040F095`
- Headers
Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=Content-Type: application/oracle.com.cloud.common.VM+json
- Body

```
{
  "status" : "STOPPING"
}
```

- Method

PUT

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
      stController.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:

- URL

https://example.oracle.com/em/cloud/iaas/server/vm/C5CA59AF9386975E8AEA45B0F040F095?resource_state

- Headers

Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=Content-Type: application/oracle.com.cloud.common.VM+json

- Body

```
{
  "status" : "STOPPED"
}
```

- Method

PUT

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:

- URL

https://example.oracle.com/em/cloud/iaas/server/vm/C5CA59AF9386975E8AEA45B0F040F095?resource_state

- Headers

Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=Content-Type: application/oracle.com.cloud.common.VM+json

- Body

```
{
  "status" : "STARTED"
}
```

- Method

PUT

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'"
      }
    ]
  }
}
```

20.3.1.2 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:

- URL


```
https://example.oracle.com/em/cloud/iaas/server/vm/C5CA59AF9386975E8AEA45B0F040F095?resource_state
```
- Headers


```
Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=Content-Type: application/oracle.com.cloud.common.VM+json
```
- Body


```
{
  "cpu": [1, 0],
  "memory": "1024"
}
```
- Method


```
PUT
```

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:

- 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"
}

```

20.4 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:

- URL
 - https://example.oracle.com/em/cloud/iaas/server/vm/C5CA59AF9386975E8AEA45B0F040F095
- Headers
 - Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=
- Method
 - DELETE

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:

- 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"
}
```

Please note the "status" of the VM has the "DELETED" value.

20.5 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.

- URL

https://example.oracle.com/em/cloud/iaas/zone/D73AF0D42C8459E11419862797D1F37D?service_instances

- Headers

Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=

- Method

GET

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:

- URL

`https://example.oracle.com/em/cloud/iaas/servicetemplate/vm/oracle%3AdefaultService%3Aem%3Aprovisioning%3A1%3Acmp%3AVirtualization%3ATemplate%3AC76CEB5563EA5E13E040578CDA817FAF%3A0.1?service_instances`
- Headers

Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=
- Method

GET

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"
      }
    ]
  }
}

```

20.6 Version 10000 Support

The 12.1.0.4 Cloud Service Portal plug-in supports the specification version 10000. To start, issue the following HTTP request on the Enterprise Manager:

- URL

`https://example.oracle.com/em/cloud`
- Headers

Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=
X-Specification-Version: 10000
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 the X-specification-version must be included in the header whenever working with 10000 version resources).

- URL

```
https://example.oracle.com/em/cloud/vdc/default/D73AF0D42C8459E11419862797D1F37D
```

- Headers

```
Authorization: basic ZGVtb3VzZXI6ZGVtb3VzZXI=
```

```
X-Specification-Version: 10000
```

```
Accept: application/oracle.com.cloud.common.Cloud+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" :
"oracle:defaultService:em:provisioning:1:netConfig:C76F741AFD7EB760E040578CDA812CD
8"
    }
}
} ,
"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" : "FD4F3945CDD6BD6EEEEFFAB064735ECDD"
        } ,
        {
            "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 is not specified, the following will 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"
    }
  ]
}
```

Only Infrastructure Service resources are supported on the 10000 version. Features supported are the same as the previous version.

Database as a Service APIs

Note: DBaaS APIs are supported only with the Enterprise Manager for Oracle Cloud 12.1.0.3 plug-in.

This chapter describes the Database as a Service (DBaaS) APIs. It contains the following sections:

- [Operations on ServiceTemplate Resources](#)
- [Operations on DbPlatformInstance Resources](#)
- [Operations on DbPlatformTemplate Resources](#)

21.1 Operations on ServiceTemplate Resources

This section describes the operations that can be performed on a ServiceTemplate resource, such as viewing the attributes of the service template and the definition of the service template. It contains the following sections:

- [Retrieving Service Template Details](#)
- [Deleting a Service Template](#)

21.1.1 Retrieving Service Template Details

The GET command has the following characteristics:

- Synopsis
GET {URI of a Service template instance}
- Request Headers:
Host, X-Cloud-Client-Specification-Version, Authorization, Accept
- Request Message Body:
N/A
- Response Headers
Content-Length, Content-Type, Content-Location, Cache-Control
- Response Message Body
ServiceTemplate
- Response Status

200, 400, 401, 403, 404

Example Request: Retrieve the Information About a Service Template

```
GET /templates/items/t833
Host: cloudcompany.com
Authorization: Basic xxxxxxxxxxxx
Accept: application/oracle.com.cloud.common.ServiceTemplate+json
X-Cloud-Client-Specification-Version: 0.1
```

Example Response:

```
HTTP/1.1 200 OK
Content-Type:
application/oracle.com.cloud.common.application/oracle.com.cloud.common.ServiceTem
plate+json
Content-Location: /templates/items/t833
Content-Length: nnn
{
  "uri" : "...",
  "name" : "Oracle EM NextGen Demo in a box",
  "description" : "Template that can be deployed into complete EM
  Next Generation Demo system with..." ,
  "type" : "VMTemplate" ,
  "created" : "February 22, 2010 15:34:34 GMT" ,
  "definition" : "<xml><some_definition/>...</xml>" ,
  "tags" : ["ABC", "EM", "Next Gen", "My note", ... ] ,
  "resource_state" : [
    "state" : "READY"
  ]
}
```

21.1.2 Deleting a Service Template

The cloud service provider may allow users to delete the service templates, based on the provider specified accessibility rules. However, since there may be active instances referencing to the service template, the cloud service provider may reserve the right to remove only the visibility of the service template from any menus. In that case, subsequent GET requests to a deleted service template may still be valid.

The command has the following characteristics:

- Synopsis
 - DELETE {URI of a service template}
- Request Headers:
 - Host, X-Cloud-Client-Specification-Version, Authorization
- Request Message Body:
 - N/A
- Response Headers
 - N/A
- Response Message Body
 - N/A
- Response Status
 - 200, 401, 403, 404, 410

Example Request: Delete a Service Template

```
DELETE /templates/items/t903
Host: cloudcompany.com
Authorization: Basic xxxxxxxxxxxx
X-Cloud-Client-Specification-Version: 0.1
```

Example Response:

```
HTTP/1.1 200 OK
```

21.2 Operations on DbPlatformInstance Resources

This chapter describes the operations that can be performed on a DbPlatformInstance resource, such as creating a new DbPlatformInstance, retrieving all the DbPlatformInstances from a virtual data center, and deleting a DbPlatformInstance.

This chapter contains the following sections:

- [Creating a New DbPlatformInstance](#)
- [Retrieving all the DbPlatformInstances in a Virtual Data Center](#)
- [Deleting a DbPlatformInstance](#)

21.2.1 Creating a New DbPlatformInstance

A DbPlatformInstance can be provisioned based on a DbPlatformTemplate. The following example request uses POST on ZONE_1 URI with the based_on attribute set to the DbPlatformTemplate URI.

Example Request: POST the following payload to /em/cloud/dbzone/D7EEE339C1249F83FC4587C722EB20B5

```
POST /em/cloud/dbzone/D7EEE339C1249F83FC4587C722EB20B5
Host: cloudcompany.com
Authorization: Basic xxxxxxxxxxxx
Content-Type: application/oracle.com.cloud.common.DbPlatformInstance+json
Accept: application/oracle.com.cloud.common.DbPlatformInstance +json
{
  "name" : "Database 11.2.0.3.0 Instance for SSA_USER_1",
  "description" : "Instance 2 of Oracle Database 11.2.0.1.0" ,
  "based_on" : "/em/cloud/dbplatformtemplate/B59EDC65951039E8E040E50A8F5B0435",
  "params" : {
    "username" : "xxx",
    "password" : "yyy"
  }
}
```

Example Response:

```
HTTP/1.1 200 OK
Content-Type: application/oracle.com.cloud.common.DbPlatformInstance +json
Content-Location: /em/cloud/dbplatforminstance/byrequest/41
Content-Length: nnn
{
  "uri" : "/em/cloud/dbplatforminstance/byrequest/41",
  "name" : "SSA_USER_1 - Thu Jan 05 02:53:24 PST 2012",
  "status" : "SCHEDULED",
  "resource_state" : {
    "state" : "INITIATED"
  }
}
```

```
}

```

21.2.2 Retrieving all the DbPlatformInstances in a Virtual Data Center

The following shows an example request and response to retrieve the information about a Small Database Service Template for an SSA_USER_ROLE.

Example Request:

```
GET /em/cloud/vdc/default/D7EEE339C1249F83FC4587C722EB20B5
Host: cloudcompany.com
Authorization: Basic xxxxxxxxxxxx
Accept: application/oracle.com.cloud.common.VDC+json

```

Example Response:

```
HTTP/1.1 200 OK
Content-Type: application/oracle.com.cloud.common.VDC+json
Content-Location: /em/cloud/vdc/default/D7EEE339C1249F83FC4587C722EB20B5
Content-Length: nnn
{
  "uri" : "/em/cloud/vdc/default/D7EEE339C1249F83FC4587C722EB20B5",
  "name" : "Oracle Enterprise Manager Virtual Data Center on Zone ZONE_1",
  "dbplatform_instances" : {
    "uri" : "/em/cloud/vdc/default/D7EEE339C1249F83FC4587C722EB20B5/dbplatform_
instances",
    "name" : "DbPlatform Instances in ZONE_1",
    "total" : "1",
    "elements" : [
      "uri" : "/em/cloud/dbplatforminstance/52A5283827BEDAC79AFEE2A0BFFF531A",
      "name" : "Dbf54e97.sta00434.us.oracle.com",
      "based_on": "/em/cloud/dbplatformtemplate/B59EDC65951039E8E040E50A8F5B0435",
      "destination_zone" : "/em/cloud/dbzone/D7EEE339C1249F83FC4587C722EB20B5",
      "connect_string" : "(DESCRIPTION=(ADDRESS_
LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=sta00434.us.oracle.com)(PORT=1521)))(CONNECT_
DATA=(SID=DBf54e97)))",
      "status" : "RUNNING",
      "created" : "2012-01-04 01:35:12"
    "resource_state" : {
      "state" : "READY"
    }
  ]
}
}

```

21.2.3 Deleting a DbPlatformInstance

The following shows an example request and response to deleting a DbPlatformInstance.

Example Request: Delete a DbPlatformInstance

```
DELETE /em/cloud/dbplatforminstance/52A5283827BEDAC79AFEE2A0BFFF531A
Host: cloudcompany.com
Authorization: Basic xxxxxxxxxxxx
Accept: application/oracle.com.cloud.common.DbPlatformInstance +json

```

```
{
  "uri" : "/em/cloud/dbplatforminstance/52A5283827BEDAC79AFEE2A0BFFF531A"
  "name" : "Dbf54e97.sta00434.us.oracle.com",
  "resource_state" : {
"state" : "DESTROYING"
  }
}
```

Example Response:

```
HTTP/1.1 200 OK
Content-Type: application/oracle.com.cloud.common.DbPlatformInstance +json
Content-Location: /em/cloud/dbplatforminstance/52A5283827BEDAC79AFEE2A0BFFF531A
Content-Length: nnn
{
  "uri" : "/em/cloud/dbplatforminstance/52A5283827BEDAC79AFEE2A0BFFF531A",
  "name" : "Dbf54e97.sta00434.us.oracle.com",
  "resource_state" : {
    "state" : "DESTROYING"
  }
}
```

21.3 Operations on DbPlatformTemplate Resources

This section describes the operations that can be performed on a DbPlatformTemplate resource, such as getting all the DbPlatformTemplates in the cloud and retrieving information about DbPlatformTemplates. It contains the following sections:

- [Retrieving all the DbPlatformTemplates in the Cloud](#)
- [Retrieving Information about a DbPlatformTemplate](#)

21.3.1 Retrieving all the DbPlatformTemplates in the Cloud

This operation is performed by issuing a GET request on Cloud resource and filtering Collection<ServiceTemplates> by type = "dbaas".

Example Request: Retrieve the information about all DbPlatformTemplates

```
GET /em/cloud
Host: cloudcompany.com
Authorization: Basic xxxxxxxxxxxx
Accept: application/oracle.com.cloud.common.Cloud+json
```

Example Response:

```
HTTP/1.1 200 OK
Content-Type: application/oracle.com.cloud.common.Cloud+json
Content-Location: cloudcompany.com
Content-Length: nnn
{
  "uri" : "/em/cloud",
  "name" : "Oracle Cloud Service Provider",
  "service_templates" : {
    "uri" : "/em/cloud/collection/servicetemplates",
    "total" : "1" ,
  "elements" : [
    "name" : "Small Database Service Template for SSA_USER_ROLE_1",
    "uri": "/em/cloud/dbplatformtemplate/B59EDC65951039E8E040E50A8F5B0435",
    "type" : "dbaas"
  ]
}
```

```

    },
    "resource_state" : {
"state" : "READY"
    }
}

```

21.3.2 Retrieving Information about a DbPlatformTemplate

The following shows an example request and response to retrieve the information about Small Database Service Template for an SSA_USER_ROLE.

Example Request:

```

GET /em/cloud/dbplatformtemplate/B59EDC65951039E8E040E50A8F5B0435
Host: cloudcompany.com
Authorization: Basic xxxxxxxxxx
Accept: application/oracle.com.cloud.common.DbPlatformTemplate+json

```

Example Response:

```

HTTP/1.1 200 OK
Content-Type: application/oracle.com.cloud.common.DbPlatformTemplate+json
Content-Location: /em/cloud/dbplatformtemplate/B59EDC65951039E8E040E50A8F5B0435
Content-Length: nnn
{
  "uri" : "/em/cloud/dbplatformtemplate/B59EDC65951039E8E040E50A8F5B0435",
  "name" : "Small Database Service Template for SSA_USER_ROLE_1",
  "description" : "Small Database Service Template for SSA_USER_ROLE_1"
  "created" : "February 22, 2011 15:34:34 GMT" ,
  "db_version" : "11.2.0.3.0",
  "db_type" : "dbTypeSI",
  "storage_type" : "File System",
  "backup_enabled" : "false",
  "total_memory" : "3660",
  "block_size" : "8192",
  "cpu" : "4",
  "processes" : "150 ",
  "type" : "dbaas",
  "instance_configurable_params" : ["username", "password"],
  "zones" : {
    "uri" : "/em/cloud/dbplatformtemplate/zones/
      B59EDC65951039E8E040E50A8F5B0435"
    "type" : "dbZone" ,
    "total" : "1" ,
    "elements" : [
      {
        "uri" : "/em/cloud/dbzone/D7EEE339C1249F83FC4587C722EB20B5",
        "name" : "ZONE_1"
      }
    ]
  }
}

```

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:

- [When to Use Chargeback APIs](#)
- [Retrieving Metering Data from Enterprise Manager](#)

22.1 When to Use Chargeback APIs

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.

22.2 Retrieving Metering Data from Enterprise Manager

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 specified a target type in the `target_type` field, enter the name of the target here. 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.

22.2.1 `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", ""
```

22.2.1.1 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.

A

Accept, 18-3
activating and deactivating policies, 8-7
adding target
 Chargeback, 4-14
administering resources and requests, 6-6
assemblies
 cloud setup, 5-18
assembly
 binary file prerequisite, 7-21
 components, creating, 7-21
 definition of, 7-21
 deploying, 7-28
 instance, definition of, 7-28
AssemblyInstance, 19-10
assigning users, 12-2
Authorization, 18-3

B

bare metal provisioning, definition of, 7-44
BI Publisher
 integrate with Chargeback, 4-21
 integrate with Consolidation Planner, 3-10
business hierarchy
 Chargeback
 Chargeback
 business hierarchy, 4-2
 import from LDAP server, 4-13
business hierarchy. See also cost center

C

Cache-Control, 18-4
charge plan
 assign to target, 4-16
charge trends
 Chargeback, 4-18
Chargeback
 add target, 4-14
 assign charge plan to target, 4-16
 assign cost center to target, 4-15
 BI Publisher, 4-21
 cloud computing, 4-2
 cost center, 4-2

 currency, 4-16
 report generation, 4-20
 self service applications, 4-4
 shared resources, 4-1
 target resource usage, 4-19
 universal metrics, 4-2
 usage and charge trends, 4-18
 usage metering, 4-2
Chargeback metrics
 configuration-based, 4-3
 usage-based, 4-3
Chargeback target
 dedicated, 4-3
 shared, 4-3
cloud
 defining performance policies, 8-3
 infrastructure, 5-1
 installing plug-ins, 2-1
 Oracle Management Agent and Cloud
 Control, 12-4, 15-4
 policies
 monitoring, 8-2
 policy groups, 8-8
 viewing, 8-2
 viewing for a target, 8-2
 Private Cloud Management, 11-1
 setting up private cloud, 2-7
 Super Administrator, 2-6
 system security, 2-6
Cloud computing, 18-1
cloud computing
 Chargeback, 4-2
Cloud Control and Management Agent, 12-4, 15-4
Cloud Policy Home page, 8-2
Cloud Resource Models, 19-3
Cloud Self Service Portal Home page, 6-6
Common Behaviors, 18-2
Common Resource Attributes, 18-5
configuring
 listener running from same Oracle Home, 12-8
 listener running from SIHA Oracle Home, 12-8
 Oracle Software Library, 12-6, 15-6
 privilege delegation settings, 12-5, 15-5
 VLAN groups, 5-9
consolidation planner
 using, 3-4

- consolidation constraints, 3-3
- Consolidation Planner
 - BI Publisher, 3-10
- consolidation planner
 - overview, 3-1
- consolidation project
 - creating, 3-4
 - definition, 3-2
- consolidation scenario
 - custom, 3-7
 - definition, 3-2
 - evaluating, 3-7, 3-9
 - pre-configured, 3-6
- Content-Length, 18-3
- Content-Type, 18-3
- cost center
 - assign to target, 4-15
 - Chargeback, 4-2
 - create, 4-12
- Cost Center target property, 4-11
- Create New Service Template Wizard, 13-6
- Create Zone page, 5-25
- creating
 - assembly components, 7-21
 - ISO component, 7-24
 - named credentials, 12-7
 - networks, Configure Ports and VLAN Interfaces
 - page, 5-11
 - networks, prerequisites for, 5-10
 - PXE bootable guest virtual machines, 7-42
 - roles, 2-7
 - service template, prerequisite, 13-6
 - storage repository, 5-16
 - template components, 7-23
 - users, 2-7, 12-4, 15-4
 - virtual disk components, 7-26
 - virtual machines, 7-28
 - virtual server pool, 5-22
 - Virtual IP, 5-23
 - zones, 5-24, 7-7
- creating a cost center, 4-12
- creating a universal charge plan, 4-7
- creating an extended charge plan, 4-8
- credentials
 - setting for OMS requirement, 7-22
 - setting up for provisioning, 12-7
- currency
 - Chargeback, 4-16
- customizing
 - database deployment procedure, 13-1

D

- data sharing, 4-22
- Database as a service (DBaaS), 11-1
- Database Cloud Self Service Portal, 13-9
- Database Self Service Portal, 13-7
- databases
 - deployment procedure, customizing, 13-1
 - requesting, 13-9

- software, provisioning, 12-7
- dedicated target
 - Chargeback, 4-3
- defining
 - performance policies, 8-3
 - policy group, 8-8
 - roles
 - EM_CLOUD_ADMINISTRATOR, 12-3
 - EM_SSA_ADMINISTRATOR, 12-3
 - EM_SSA_USER, 12-3
- Deleting a Service Template, 21-2, 21-4, 21-6
- deploying
 - an assembly, 7-28
 - ISO image, 7-40
 - Oracle Chargeback and Trending, 2-1
 - Oracle Cloud Application, 2-1
 - Oracle Virtualization, 2-1
 - templates, 7-36
 - virtual machine provisioning, 7-20
 - virtual machined infrastructure
 - requirements, 7-21
- Deployment Configuration page, 7-37
- discovering Oracle VM Server, 5-6
- Distributed Power Management (DPM)
 - policies, 8-11
- Dynamic Resource Scheduling (DRS) policies, 8-11

E

- editing
 - existing service template, 13-6
 - virtual server, 7-11
 - zones, 7-7
- EM_SSA_ADMINISTRATOR, and provisioning
 - database software, 12-7
- enable metric collection
 - for shared database, 4-15
- Error Response Message, 18-6
- extended charge plan
 - Chargeback
 - Chargeback
 - extended charge plan, 4-2
 - create, 4-8
 - update, 4-11

G

- generating reports
 - Chargeback, 4-20
- guest virtual machine
 - definition of, 7-13
 - live migrating, 7-13
 - prerequisites for creating, 7-12
 - restarting, 7-14
 - starting, 7-14
- Guest Virtual Machine Home page, 7-13

H

- Host, 18-3

HTTP Status Codes, 18-4
hypervisor, 7-10

I

IaaS Self Service portal, 11-2
installing
 Management Agent on an unmanaged server, 12-4
 plug-ins, 2-1
ISO component, creating, 7-24
ISO Deployment page, 7-40
ISOs, 5-18

J

JavaScript Object Notation, 18-2

L

LDAP server
 import business hierarchy, 4-13
list of OVM Managers, 7-5
live migrating guest virtual machine, 7-13
Location, 18-3

M

machine sizes, setting up, 5-26
maintaining virtual servers, 7-12
managing
 virtualization targets
 OVM Manager Home page, 7-3
 zones, 7-6
monitoring
 cloud policies, 8-2
 resources and requests, 6-6

N

network bootable (PXE boot) virtual machine, 7-42
network profile
 domain name, 5-13
 Netmask, 5-13
 setting up, 5-12
 specifying IP addresses, 5-13
network roles
 Cluster Heartbeat, 5-11
 Live Migrate, 5-11
 Server Management, 5-10
 storage, 5-11
 Virtual Machine, 5-11
network type, setting up, 5-26
NetworkInterface, 19-14
networks
 creating, 5-10
 setting up, 5-7
NFS, 15-7

O

OCFS2, 15-7
OMS, setting credentials, 7-22
Oracle Software Library
 configuring, 12-6, 15-6
 setting up, 2-3
Oracle VM Manager, 5-2
 registering, 5-3
Oracle VM Network, 5-7
 Virtual Area Network (VLAN) Group, 5-9
Oracle VM Server
 Additional OS Details page, 7-45
 creating virtual server pool, 5-22
 discovering, 5-6
 file systems, 5-13
 image, 7-44
 storage servers, setting up, 5-13
 virtual server, 5-6
 virtual server pool, 5-22
 virtualization software (hypervisor), 5-6
OVM Manager Home page, 7-3
OVM Managers, list of, 7-5

P

performance policies, defining, 8-3
plug-ins
 deploying Oracle Chargeback and Trending, 2-1
 deploying Oracle Cloud Application, 2-1
 deploying Oracle Virtualization, 2-1
 installing, 2-1
policies
 activating and deactivating, 8-7
 defining groups, 8-8
 Distributed Power Management (DPM), 8-11
 Dynamic Resource Scheduling (DRS), 8-11
 schedule-based, 8-4
PowerBroker, 12-5, 15-5
Preboot Execution Environment (PXE), 7-42
prerequisites
 assembly binary file, 7-21
 creating a zone, 5-24
 creating networks, 5-10
 creating PXE bootable guest virtual machines, 7-42
 provisioning guest virtual machine using ISO image, 7-40
 provisioning guest virtual machine using templates, 7-37
 setting up self-service portal, 13-3
 virtual server pools, 5-22
Private Cloud Management, 11-1
Private Cloud, setting up, 2-7
privilege delegation settings, configuring, 12-5, 15-5
privilege delegation tools, Sudo and
 PowerBroker, 12-5, 15-5
Provision Oracle VM Server
 Deployment page, 7-45
 Disk Layout page, 7-46
provisioning

- credentials, setting up, 12-7
- guest virtual machine
 - using ISO image, 7-40
- guest virtual machine using ISO image, 7-40
- guest virtual machine using templates
 - prerequisites, 7-37
- guest virtual machines, 7-20
- virtual management servers, 7-44

Provisioning VM Server Basic OS Details page, 7-45

PXE bootable guest virtual machines, creating, 7-42

Q

QoS, setting up, 5-25

R

- reference gold images, 2-3
- Representational State Transfer API, 18-1
- requesting or reserving servers, 6-12
- resources and requests, monitoring and administering, 6-6
- RESTARTING, 19-13
- restarting guest virtual machine, 7-14
- RESTful API, 18-1
- RESTful APIs, 18-2
- RESUMING, 19-13
- Retrieve Information About the Service
 - Template, 21-1, 21-3, 21-5
- roles
 - creating, 2-7
 - defining, 12-2
 - definition of, 12-2
 - EM_CLOUD_ADMINISTRATOR, 2-7, 12-3, 12-4, 15-5
 - EM_SSA_ADMINISTRATOR, 2-7, 12-3
 - provisioning database software, 12-7
 - EM_SSA_USER, 2-7, 12-3

S

- Scalability Group, 19-11
- schedule-based policy, 8-4
- self service applications
 - Chargeback, 4-4
- Self Service Portal
 - defining mapping between roles and other entities, 6-3
 - defining set of templates or assemblies, 6-4
 - Machines Sizes, 6-1
 - requesting or reserving servers, 6-12
 - setting up, 13-3
 - using for database, 13-7
- Server, 19-12
- server data model, 19-12
- servers, requesting or reserving, 6-12
- Service Template, 19-4
- service templates
 - editing existing template, 13-6
 - prerequisite for creating, 13-6
 - setting credentials for OMS, 7-22

- setting up
 - database cloud self-service portal, 13-3
 - listener for database zone hosts, 12-8
 - machine sizes, 5-26
 - network profile, 5-12
 - networks, 5-7
 - network profile, 5-12
 - network type, 5-12
 - private cloud, 2-7
 - provisioning credentials, 12-7
 - QoS, 5-25
 - QoS for network, 5-26
 - Self Service Portal, 13-3
 - Software Library, 2-3
 - storage servers, 5-13
 - virtualization infrastructure, 5-2
- shared database
 - enable metric collection, 4-15
- shared resources
 - Chargeback, 4-1
- shared target
 - Chargeback, 4-3
- sharing Chargeback data, 4-22
- Software Library
 - configuring, 12-6, 15-6
 - setting up, 2-3
- STARTED, 19-13
- STARTING, 19-13
- STOPPED, 19-13
- STOPPING, 19-13
- stopping the guest virtual machine, 7-14
- Storage Quality of Service (QoS), setting up, 5-25
- storage repository
 - assemblies, 5-19
 - creating, 5-16
 - presenting, 5-17
 - Virtual Disks, 5-20
 - VM Templates, 5-20
- Sudo, 12-5, 15-5
- SUSPENDED, 19-13
- SUSPENDING, 19-13
- system security, 2-6
 - Super Administrator, 2-6

T

- target property, for cost center, 4-11
- target resource usage
 - Chargeback, 4-19
- template components
 - creating, 7-23
 - virtualization types, 7-23
- templates
 - deploying, 7-36
 - provisioning guest virtual machine, 7-37
- topology of a zone, 5-24
- Transport Protocol, 18-2

U

- universal charge plan
 - Chargeback
 - Chargeback
 - universal charge plan, 4-2
 - create, 4-7
 - update, 4-8
- universal metrics
 - Chargeback, 4-2
- updating a universal charge plan, 4-8
- updating an extended charge plan, 4-11
- URI Space, 18-2
- usage metering
 - Chargeback, 4-2
- usage trends
 - Chargeback, 4-18
- users
 - assigning, 12-2
 - creating, 2-7, 12-4, 15-4

V

- viewing
 - cloud policies, 8-2
 - cloud policies for a target, 8-2
- virtual
 - appliance images, 2-3
 - disk component, creating, 7-26
 - machine provisioning, 7-20
 - machines, creating, 7-28
- Virtual Server Home page, 7-10
- virtual server pool
 - zone, creating, 5-24
- Virtual Server Pool Home page, 7-8
- virtual server pools
 - creating or editing, 5-21
 - definition, 7-8
 - prerequisites, 5-22
- virtual servers
 - definition, 7-10
 - editing, 7-11
 - maintaining, 7-12
 - Oracle VM Server, 5-6
- virtualization
 - infrastructure, 5-1
 - infrastructure, setting up, 5-2
 - managing members, 7-5
 - managing targets
 - OVM Manager Home page, 7-3
 - software (hypervisor), 5-6
 - types, template components, 7-23
 - viewing members, 7-5
- VLAN groups, configuring, 5-9
- VM files, 5-18
- VNet, 19-13

Z

- Zone, 19-5
- Zone Home page, 7-6

zones

- Create Zone page, 5-25
- creating, 7-7
- definition, 5-24
- editing, 7-7
- managing, 7-6
- prerequisites to creating, 5-24
- topology of, 5-24

