

Oracle® Fusion Applications

Administrator's Guide

11g Release 1 (11.1.3)

E14496-04

March 2012

Documentation for system administrators that describes how to administer the Oracle Fusion Applications environment, including how to start and stop components, change ports and passwords, monitor the environment, logging, back up and recovery, configure high availability, and move components across environments.

Oracle Fusion Applications Administrator's Guide, 11g Release 1 (11.1.3)

E14496-04

Copyright © 2011, 2012, Oracle and/or its affiliates. All rights reserved.

Primary Author: Deborah Steiner (lead), Shelly Butcher, Thom Chumley, Christine Ford, Helen Grembowicz, Vinaye Misra, Leslie Studdard, Carlos Subi, Karen Summerly

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 America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

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

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

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

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

Contents

Preface	xxi
Audience	xxi
Documentation Accessibility	xxi
Related Documents	xxi
Conventions	xxii

What's New in This Guide for Release 11.1.3	xxiii
--	-------

Part I Understanding Oracle Fusion Applications

1 Introduction to Oracle Fusion Applications for Systems Administrators

1.1	Architecture of Oracle Fusion Applications	1-1
1.1.1	Oracle Fusion Applications Product Families	1-3
1.1.2	Oracle Fusion Middleware Infrastructure Components.....	1-4
1.1.3	Oracle Fusion Middleware Components	1-5
1.1.4	Oracle Database	1-9
1.1.5	Oracle Enterprise Manager Fusion Applications Control	1-9
1.1.6	Oracle Enterprise Manager Cloud Control.....	1-9
1.1.7	Oracle Enterprise Manager Database Control.....	1-9
1.2	Key Oracle Fusion Applications Concepts.....	1-10
1.2.1	Provisioning and the Installation Process	1-10
1.2.1.1	Installation Options	1-10
1.2.1.2	Installing an Oracle Database and Loading Content	1-11
1.2.2	Oracle WebLogic Server Domains Configuration	1-11
1.2.3	Provisioned Oracle Fusion Applications Home Directories	1-13
1.3	Roadmap for Administering Oracle Fusion Applications.....	1-17

Part II Basic Administration

2 Getting Started with Administering Oracle Fusion Applications

2.1	Introduction to the Tools Used to Administer Oracle Fusion Applications	2-1
2.2	Key Differences Between Fusion Applications Control and Cloud Control.....	2-3
2.3	When to Use Fusion Applications Control Tool, When to Use Oracle WebLogic Server Administration Console	2-4
2.4	Finding the Administration Server Port and Other Administrative URLs for Domains.	2-5

2.5	Using Fusion Applications Control.....	2-7
2.5.1	Understanding How Fusion Applications Control Compares with Fusion Middleware Control	2-8
2.5.2	Starting Fusion Applications Control	2-8
2.5.3	Using Fusion Applications Control Help.....	2-9
2.5.4	Navigating within Fusion Applications Control	2-9
2.5.5	Understanding the Targets in the Target Navigation Pane.....	2-12
2.5.6	Navigating to the Product Family and Product Administration Pages	2-12
2.5.7	Viewing the Performance of Targets	2-14
2.5.8	Understanding Users and Roles for Fusion Applications Control.....	2-15
2.6	Using Oracle WebLogic Server Administration Console	2-15
2.6.1	Starting Oracle WebLogic Server Administration Console.....	2-16
2.6.2	Locking the WebLogic Server Configuration	2-18
2.7	Using Cloud Control	2-18
2.7.1	Starting Cloud Control and Discovering Oracle Fusion Applications Targets	2-18
2.7.2	Rediscovering Oracle Fusion Applications Targets	2-20
2.7.3	Removing Oracle Fusion Applications Targets.....	2-20
2.7.4	Navigating the Fusion Applications Target Home Page in Cloud Control.....	2-21
2.7.5	Using Cloud Control Help	2-22
2.8	Using Oracle Fusion Functional Setup Manager.....	2-22
2.9	Using Database Tools to Modify Oracle Fusion Applications Data.....	2-23

3 Performing Routine Administrative Tasks

3.1	Introduction to Performing Routine Administrative Tasks	3-1
3.2	Running Administration Servers and Managed Servers from Shared and Local Disks..	3-2
3.3	Starting and Stopping.....	3-2
3.3.1	Starting and Stopping Components in the Oracle Fusion Applications Environment.....	3-2
3.3.2	Starting and Stopping the Entire Oracle Fusion Applications Environment	3-4
3.3.2.1	Starting an Oracle Fusion Applications Environment	3-4
3.3.2.2	Stopping an Oracle Fusion Applications Environment	3-9
3.3.3	Starting and Stopping the Oracle Fusion Applications Middle Tier Using the fastartstop Utility	3-13
3.3.3.1	fastartstop Syntax	3-13
3.3.3.2	Registry Entries.....	3-15
3.3.3.3	Starting an Oracle Fusion Applications Environment with the fastartstop Utility for the Middle Tier	3-15
3.3.3.4	Stopping an Oracle Fusion Applications Environment with the fastartstop Utility ..	3-18
3.3.4	Starting and Stopping a Product Family Oracle WebLogic Server Domain.....	3-19
3.3.4.1	Stopping an Oracle WebLogic Server Domain for a Product Family	3-20
3.3.4.2	Starting an Oracle WebLogic Server Domain for a Product Family	3-20
3.3.5	Starting and Stopping an Oracle WebLogic Server Cluster for a Configuration Change.....	3-21
3.3.5.1	Starting and Stopping an Oracle WebLogic Server Cluster Using WLST and Node Manager	3-21
3.3.5.2	Starting and Stopping an Oracle WebLogic Server Cluster Using Fusion Applications Control	3-21

3.3.6	Starting and Stopping Specific Applications	3-22
3.3.6.1	Starting and Stopping Java EE Applications Using WLST	3-22
3.3.6.2	Starting and Stopping an Individual Application Instance for a Cluster Using Fusion Applications Control or Cloud Control	3-22
3.3.6.3	Starting and Stopping All Application Instances for a Cluster Using Fusion Applications Control or Cloud Control	3-23
3.4	Viewing and Changing Ports for Components	3-23
3.5	Changing Passwords	3-25
3.6	Managing the Oracle Database	3-27
3.6.1	Viewing and Modifying Initialization Parameters	3-27
3.6.2	Managing Database Storage Structures.....	3-28
3.6.3	Managing Memory	3-29
3.6.4	Administering User Accounts	3-30
3.7	Patching.....	3-30
3.8	Managing Oracle Fusion Applications-Specific Labels in the Oracle Metadata Repository	3-31
3.9	Modifying Oracle Application Development Framework Connections in Oracle Fusion Applications	3-32

4 Securing Oracle Fusion Applications

4.1	Introduction to Security	4-1
4.2	About the Enterprise Identity Store	4-2
4.2.1	Supported LDAP Identity Providers	4-3
4.2.2	Configuring the Identity Store.....	4-3
4.3	Provisioning Identities	4-3
4.3.1	Identity Provisioning Concepts	4-4
4.3.1.1	Administrators For Fusion Applications	4-4
4.3.1.2	Two Types of Users During Provisioning	4-5
4.3.2	WebLogic Authenticators and the Primary Identity Store	4-6
4.3.3	Provisioning Steps	4-6
4.3.4	Best Practices for the Administrator Groups	4-7
4.3.5	Managing Identities after Deployment	4-7
4.4	Managing Authorization Policies	4-7
4.4.1	Managing Oracle Fusion Application Policies	4-8
4.4.2	Managing System Policies	4-8
4.4.3	Reconciling GUIDs	4-9
4.4.3.1	Prerequisites to Running DSDDataMigrator.....	4-9
4.4.3.2	DSDDataMigrator Syntax.....	4-11
4.4.3.3	Examples of Use.....	4-12
4.4.4	Managing Data Security	4-12
4.5	Configuring Roles	4-13
4.5.1	Configuring Oracle Fusion Application Roles	4-13
4.5.2	Configuring Enterprise Roles.....	4-14
4.6	Configuring Audit Trail.....	4-14
4.7	Configuring SSL for Oracle Fusion Applications.....	4-15
4.7.1	SSL Configuration in Oracle Fusion Middleware.....	4-15
4.7.1.1	SSL and Infrastructure Hardening.....	4-15

4.7.1.2	Communication in the Three-Tier Model	4-15
4.7.2	SSL Configuration for Oracle Fusion Applications	4-16
4.7.2.1	Basic Network Topology	4-16
4.7.2.2	Provisioned SSL Connections	4-18
4.7.3	Implementing SSL for Identity Management Configuration	4-18
4.7.4	Additional SSL Configuration	4-19
4.7.4.1	SSL-enable Oracle Identity Management.....	4-20
4.7.4.2	SSL-enable Oracle Business Intelligence	4-20
4.7.4.3	SSL-enable ECM	4-20
4.7.4.4	SSL to External Web Services.....	4-20
4.7.5	Enabling Secure Sockets Layer on ECSF	4-21
4.8	Managing Wallets, Keystores, Credentials, and Certificates.....	4-23
4.8.1	Wallets and Keystores.....	4-23
4.8.1.1	JKS Keystore and Truststore	4-23
4.8.1.2	Oracle Wallet.....	4-23
4.8.1.3	Keystore Types Used by Products	4-24
4.8.2	Management Tools	4-24
4.8.3	Managing Wallets and their Contents	4-25
4.8.3.1	Wallet Lifecycle.....	4-25
4.8.3.2	Wallet Operations.....	4-25
4.8.3.3	Certificate Lifecycle	4-25
4.8.3.4	Certificate Operations	4-26
4.8.4	Managing Keystores and their Contents.....	4-26
4.8.4.1	Keystore Lifecycle.....	4-26
4.8.4.2	Keystore Operations.....	4-26
4.8.4.3	Certificate Lifecycle	4-26
4.8.4.4	Certificate Operations	4-27
4.8.5	Managing Credentials.....	4-27
4.8.5.1	Changing App ID Passwords	4-27
4.8.5.2	Changing the Oracle Fusion Middleware Administrative User Password	4-28
4.9	Data Masking.....	4-29
4.9.1	Introduction to Data Masking.....	4-29
4.9.1.1	Masking Terminology.....	4-29
4.9.1.2	Types of Sensitive Data.....	4-30
4.9.1.3	The FAST Methodology.....	4-30
4.9.1.4	Administration Tools	4-31
4.9.2	Data Masking in Oracle Fusion Applications.....	4-31
4.9.2.1	Requirements for Data Masking	4-31
4.9.2.2	Sensitive Data in Oracle Fusion Applications	4-35
4.9.2.3	Masking Definitions	4-35
4.9.3	Managing the Masking Definitions.....	4-35
4.9.3.1	Viewing and Modifying Data Masking Definitions	4-36
4.9.3.2	Generating the Masking Script.....	4-38
4.9.3.3	Customizing Mask Formats.....	4-39
4.9.4	Data Masking in Oracle Enterprise Manager Cloud Control.....	4-40
4.9.5	Best Practices when Masking Test Databases.....	4-42
4.9.6	References	4-42

4.10	Securing Web Services	4-42
4.10.1	Local Policy Attachment.....	4-43
4.10.2	Global Policy Attachment.....	4-43
4.10.3	Web Services Security Profiles.....	4-44
4.10.4	Key Exchange with the Domain Hosting Oracle Identity Manager	4-44
4.10.5	Web Services Security Hardening	4-45
4.11	Securing Oracle Fusion Middleware Products.....	4-45
4.11.1	Administrative Tasks and Features Specific to the IDCCS Stripe	4-45
4.11.1.1	Grants Supported by UCM	4-45
4.11.1.2	Security Groups and Accounts Associated with Documents	4-46
4.12	Extracting Data from an LDAP- Based Store to a File.....	4-46
4.13	Customizing Security from Installation to Deployment	4-47
4.13.1	Installing a New Oracle Fusion Application	4-47
4.13.2	Customizing and Testing Security with Oracle JDeveloper.....	4-47
4.13.3	Migrating to a Staging Environment	4-48
4.13.4	Migrating to a Production Environment.....	4-49

5 Managing Oracle Enterprise Scheduler Service and Jobs

5.1	Introduction to Managing Oracle Enterprise Scheduler	5-2
5.2	Navigating to the Oracle Enterprise Scheduler Administration Pages	5-2
5.3	The Scheduling Service Home Page.....	5-5
5.3.1	Top Ten Long Running Requests and Top Ten Ready Job Requests Regions.....	5-6
5.3.2	Scheduler Components Region	5-7
5.3.3	Completed Requests Region	5-7
5.3.4	Response and Load Region	5-7
5.3.5	Performance Region	5-7
5.4	Configuring Oracle Enterprise Scheduler	5-7
5.5	Searching for Configuration Changes to Oracle Enterprise Scheduler in Cloud Control	5-8
5.6	Viewing General Oracle Enterprise Scheduler Instance Information.....	5-9
5.7	Starting and Stopping Oracle Enterprise Scheduler Service Components.....	5-9
5.8	Managing Oracle Enterprise Scheduler Requests	5-10
5.8.1	Submitting an Oracle Enterprise Scheduler Job Request.....	5-10
5.8.1.1	Creating a Job Request.....	5-10
5.8.1.2	Specifying a Job Request Schedule	5-11
5.8.2	Searching for Oracle Enterprise Scheduler Job Requests.....	5-12
5.8.2.1	Searching for a Job Request Using Simple Search	5-12
5.8.2.2	Searching for a Job Request Using the Advanced Search Feature	5-13
5.8.2.3	Configuring the Display of Columns in the Search Results Table	5-15
5.9	Monitoring Oracle Enterprise Scheduler Service Activity.....	5-15
5.9.1	Viewing Top Ten Long Running Oracle Enterprise Scheduler Requests	5-16
5.9.2	Viewing Top Ten Oracle Enterprise Scheduler Jobs Ready to Be Executed	5-16
5.10	Managing Purge Policies in Fusion Applications Control.....	5-16
5.11	Managing Job Metadata	5-17
5.11.1	Managing Job Definitions.....	5-17
5.11.1.1	Viewing Job Definitions.....	5-19
5.11.1.2	Creating a Job Definition	5-20
5.11.2	Managing Job Sets	5-24

5.11.2.1	Viewing Job Sets	5-25
5.11.2.2	Creating or Editing a Job Set.....	5-25
5.11.2.3	Deleting a Job Set.....	5-31
5.11.3	Managing Incompatibilities	5-31
5.11.3.1	Viewing Incompatibilities	5-31
5.11.3.2	Creating or Editing an Incompatibility	5-32
5.11.3.3	Deleting an Incompatibility	5-33
5.12	Managing Work Assignments and Workshifts	5-33
5.13	Managing Oracle ADF Connections	5-34
5.14	Managing Application Properties	5-35
5.15	Managing Application and Role Policies	5-36
5.16	Managing Oracle Enterprise Scheduler Web Services	5-37
5.17	Managing Logging for Oracle Enterprise Scheduler	5-37
5.17.1	Finding Log Information for a Job Request	5-37
5.17.2	Setting Oracle Enterprise Scheduler Log Levels	5-37
5.17.3	Saving Job Request Logs.....	5-39
5.18	Managing Alerts for Oracle Enterprise Scheduler in Cloud Control.....	5-39
5.18.1	Configuring Metric Alerts for Oracle Enterprise Scheduler in Cloud Control	5-40
5.18.2	Viewing Alerts for Oracle Enterprise Scheduler in Cloud Control.....	5-41

6 Managing Report Delivery Servers

6.1	Introduction to Managing Report Delivery Servers	6-1
6.2	Navigating to the Oracle BI Publisher Administration Page	6-2
6.3	Configuring Report Delivery Servers	6-3
6.4	Understanding the Report and Delivery Processors	6-4
6.5	Managing Report Processing and Delivery Server Load	6-4
6.6	Scaling Out Report Processors	6-5
6.7	Diagnosing Report Processing Issues	6-5
6.7.1	Using the Scheduler Diagnostics Page	6-6
6.7.2	Using Fusion Applications Control	6-6
6.8	Configuring System Properties for Reporting.....	6-6

7 Managing Search with Oracle Enterprise Crawl and Search Framework

7.1	Introduction to Oracle Fusion Applications Search.....	7-1
7.1.1	Oracle Fusion Applications Search Use Case Example.....	7-1
7.1.2	Oracle Fusion Applications Search UI.....	7-2
7.1.3	Oracle Enterprise Crawl and Search Framework	7-4
7.1.3.1	ECSF Management Features	7-5
7.1.3.2	Key ECSF Features	7-5
7.1.3.3	ECSF Process Flow	7-6
7.1.4	Item Business Object and Data Level Security	7-6
7.2	Validating the Environment for ECSF	7-7
7.3	Getting Started with Managing ECSF with Fusion Applications Control	7-9
7.3.1	Navigating to the ECSF Administration Pages	7-10
7.3.2	Understanding the Enterprise Crawl and Search Framework Configuration Settings Page	7-11
7.4	Administering Search.....	7-12

7.4.1	Making Seeded Searchable Objects Available for Search	7-13
7.4.2	Managing Index Schedules	7-18
7.4.2.1	Creating the Index Schedules	7-19
7.4.2.2	Associating the Searchable Objects with Index Schedules	7-20
7.4.2.3	Deploying the Index Schedules	7-21
7.4.2.4	Starting the Index Schedules.....	7-21
7.4.2.5	Stopping Index Schedules	7-21
7.4.2.6	Disassociating Searchable Objects from Index Schedules	7-22
7.4.2.7	Starting Full Indexing	7-22
7.4.2.8	Undeploying Index Schedules.....	7-23
7.4.2.9	Deleting Index Schedules	7-23
7.4.3	Changing the Association of Searchable Objects with Search Categories.....	7-24
7.4.4	Renaming Search Categories.....	7-25
7.4.5	Deleting Search Categories.....	7-26
7.4.6	Making External Search Categories Available for Federated Search.....	7-26
7.4.7	Deleting External Search Categories.....	7-29
7.4.8	Modifying the Display Name of Deployed Searchable Objects.....	7-29
7.4.9	Undeploying Searchable Objects.....	7-30
7.4.10	Unregistering Searchable Objects from the Oracle Fusion Applications Database	7-31
7.4.11	Modifying Search Application Service Component Instance Parameters	7-33
7.5	Performing Advanced Administration Tasks.....	7-34
7.5.1	Modifying the Title, Body, Keyword, and Action Title of Searchable Objects.....	7-34
7.5.2	Making New Searchable Objects Available for Search	7-37
7.5.3	Defining Search Engine Instances	7-43
7.5.3.1	Creating Search Engine Instances	7-43
7.5.3.2	Configuring Search Engine Instance Parameters.....	7-44
7.5.3.3	Modifying Search Engine Instances.....	7-45
7.5.3.4	Modifying Search Engine Instance Parameters	7-46
7.5.3.5	Deleting Search Engine Instances	7-47
7.5.4	Enabling Search on External Data Sources	7-48
7.5.5	Adding Product Families for Federated Search	7-50

8 Managing Oracle WebLogic Communication Services for Click-to-Dial Functionality

8.1	Introduction to Third Party Call Web Service	8-1
8.2	Modifying Third Party Call Web Service Default Settings Using the System MBean Browser	8-1
8.3	Call Routing Through a Session Border Controller	8-3
8.3.1	Setting Up an Outbound Proxy Server.....	8-3
8.3.2	Configuring the PreLoadedRouteUri Attribute	8-4

Part III Monitoring

9 Monitoring Oracle Fusion Applications

9.1	Introduction to Monitoring Oracle Fusion Applications	9-1
9.2	The Product Family Home Page and Sections in Fusion Applications Control.....	9-2

9.2.1	Applications and Infrastructure Section	9-3
9.2.2	Key Performance Indicators Section	9-3
9.2.3	Topology Section	9-3
9.2.4	Oracle Enterprise Scheduler Sections	9-4
9.2.5	SOA Composite Section	9-4
9.3	The Product Home Page and Sections in Fusion Applications Control	9-5
9.3.1	Fusion Applications Section	9-5
9.3.2	Servers Section	9-6
9.4	The Fusion Instance Overview Page and Sections in Cloud Control	9-6
9.4.1	Summary Section	9-7
9.4.1.1	Monitoring and Diagnostics Region	9-7
9.4.1.2	Database System Region	9-7
9.4.2	Member Target Summary Section	9-7
9.4.3	Status Section	9-7
9.4.4	Top Long Running Job Requests	9-7
9.5	The Product Family Home Page and Sections in Cloud Control	9-8
9.6	The Product Home Page and Sections in Cloud Control	9-8
9.7	Monitoring Product Families and Products	9-9
9.8	Monitoring the Detailed Status of Oracle Fusion Applications	9-11
9.9	Monitoring Business Performance Metrics for a Product Family	9-12
9.10	Monitoring Business Performance Metrics for a Product	9-14
9.11	Viewing the Routing Topology of an Oracle Fusion Applications Instance, Product Family, or Product	9-15
9.12	Viewing the Configuration Topology of an Oracle Fusion Applications Instance, Product Family, or Product	9-16

10 Monitoring the Oracle Fusion Applications Middle Tier

10.1	Introduction to Monitoring the Middle Tier	10-1
10.2	Monitoring and Viewing Performance Status of Oracle Fusion Middleware	10-2
10.2.1	Viewing the General Status of Oracle Fusion Middleware Components for a Product Family	10-2
10.2.2	Viewing the Detailed Status of Oracle Fusion Middleware Components for a Product Family	10-3
10.2.3	Viewing the Routing Topology of Oracle Fusion Middleware	10-4
10.2.3.1	Viewing the Routing Topology of Oracle Fusion Middleware Using Fusion Applications Control	10-4
10.2.3.2	Viewing the Routing Topology of Oracle Fusion Middleware Using Cloud Control	10-6
10.2.4	Viewing the Performance of an Oracle WebLogic Server	10-6
10.2.5	Monitoring the System Performance of Oracle Fusion Applications Deployed to a Cluster	10-7
10.2.5.1	Monitoring the Performance of Oracle Fusion Applications Cluster Applications Using Fusion Applications Control	10-7
10.2.5.2	Monitoring the Performance of Oracle Fusion Applications Cluster Applications Using Cloud Control	10-8
10.2.5.3	Monitoring the Performance of Oracle Fusion Application Instances Using Fusion Applications Control	10-9
10.2.5.4	Monitoring the Performance of Oracle Fusion Application Instances Using Cloud Control	10-9

10.3	Monitoring System Performance Metrics for a Fusion Instance, Product Family, or a Product	10-10
10.3.1	Monitoring System Performance Metrics for a Product Family or a Product with Fusion Applications Control	10-10
10.3.2	Monitoring System Performance Metrics for a Fusion Instance, Product Family, or a Product with Cloud Control	10-11
10.4	Monitoring Performance for Key Oracle Fusion Middleware Components with Fusion Applications Control	10-12
10.4.1	Monitoring Oracle Enterprise Crawl and Search Framework	10-14
10.4.2	Monitoring the Performance of a Single Oracle Enterprise Scheduler Instance ...	10-14
10.4.2.1	Viewing a Summary of Oracle Enterprise Scheduler Activity	10-14
10.4.2.2	Viewing Current Oracle Enterprise Scheduler Activity	10-15
10.4.2.3	Viewing Historical Oracle Enterprise Scheduler Reports	10-15
10.4.2.4	Viewing Oracle Enterprise Scheduler Metrics	10-16
10.4.2.5	Managing Chart Sets	10-17
10.4.2.6	Selecting the Performance Time Period	10-18
10.4.3	Monitoring Oracle Identity Manager Performance Metrics	10-18
10.4.4	Locating Key Performance Metrics for Specific Components	10-18
10.4.4.1	Monitoring Key Performance Metrics for Oracle WebLogic Server	10-19
10.4.4.2	Monitoring Key Performance Metrics for Oracle SOA Suite	10-20
10.4.4.3	Monitoring Key Performance Metrics for Oracle Enterprise Scheduler	10-22
10.5	Performing Proactive Performance Monitoring with Cloud Control	10-23
10.5.1	Monitoring Middle Tier Targets	10-23
10.5.2	Creating Monitoring Templates	10-24
10.5.3	Creating Blackouts	10-25
10.5.4	Setting Up Notifications	10-25
10.6	Tuning Middle Tiers for Oracle Fusion Applications	10-25

11 Diagnosing the Oracle Fusion Applications Middle Tier Performance

11.1	Introduction to Java Diagnostics in the Middle Tier	11-1
11.2	Diagnosing the Oracle Fusion Applications Middle Tier	11-1
11.2.1	Viewing JVM Diagnostics Summary Information	11-2
11.2.2	Finding the Top Java Methods	11-4
11.2.3	Finding the Top SQL Queries	11-5
11.2.4	Analyzing Stuck Threads	11-6
11.2.5	Drilling Down from JVM Diagnostics to SQL Instances	11-7
11.2.6	Analyzing Potential Memory Leaks	11-8

12 Monitoring and Tuning Oracle Database for Oracle Fusion Applications

12.1	Monitoring Oracle Database	12-1
12.1.1	Database Statistics	12-1
12.1.2	Monitoring Database Alerts	12-2
12.1.2.1	Configuring Database Alerts	12-3
12.1.2.2	Responding to Database Alerts	12-3
12.1.3	Monitoring User and Session Activity	12-3
12.1.4	Monitoring the Database Instance	12-5
12.1.5	Monitoring the Host System	12-6

12.2	Tuning Oracle Database.....	12-6
12.2.1	Creating a Baseline of the Database.....	12-6
12.2.2	Diagnosing and Tuning Database Performance Problems Using ADDM.....	12-7
12.2.3	Diagnosing and Tuning Database Performance Problems Reactively.....	12-7
12.2.3.1	Comparing Database Performance Over Time Using AWR Reports.....	12-7
12.2.3.2	Identifying Transient Performance Problems Using ASH Reports.....	12-8
12.2.3.3	Identifying Wait Time Using Wait Events.....	12-8
12.2.4	Tuning SQL Statements.....	12-9
12.2.4.1	Collecting Optimizer Statistics.....	12-9
12.2.4.2	SQL Tuning Sets.....	12-10
12.2.4.3	SQL Profiles.....	12-10
12.2.4.4	Using SQL Tuning Advisor.....	12-10
12.2.4.5	Using SQL Access Advisor.....	12-10
12.2.4.6	Using SQL Tracing.....	12-11
12.3	Testing Oracle Database Changes.....	12-12
12.3.1	Testing SQL Changes.....	12-12
12.3.2	Testing Database Changes.....	12-12

13 Managing Oracle Fusion Applications Log Files and Diagnostic Tests

13.1	Introduction to Incidents, Log File Management, QuickTrace, and Diagnostic Tests ..	13-1
13.2	Relationships Between Log Files and Diagnostic Tests.....	13-2
13.3	Standard Log File Administration Tasks and Tools.....	13-3
13.4	Log Files and Settings.....	13-4
13.5	Log Message Structure.....	13-4
13.5.1	Log Message Attributes Supplied by the Oracle Fusion Middleware and Oracle Database Layers.....	13-4
13.5.2	Log Message Attributes Supplied by Logging APIs for Oracle Fusion Applications.....	13-5
13.5.3	Log Message Attributes Supplied by the Application Session.....	13-6
13.5.4	Log Message Attributes Supplied by Oracle Enterprise Scheduler Job Requests...	13-6
13.5.5	Log Message Attributes Supplied by SOA.....	13-7
13.6	Viewing and Searching Log Files During Normal Operation.....	13-7
13.6.1	Viewing Logs from Fusion Applications Control.....	13-7
13.6.2	Searching for Specific Information in Log Files within a Single Domain.....	13-9
13.6.3	Searching for Specific Information in Log Files Across Multiple Domains.....	13-12
13.6.4	Typical Log File Locations.....	13-13
13.6.5	Location of Logged Information When Using Multiple Managed Servers.....	13-16
13.7	Configuring Settings for Log Files During Normal Operation.....	13-16
13.7.1	Default System Log Settings.....	13-18
13.7.2	When Changes to Log File Configuration Settings Take Effect.....	13-22
13.7.3	Standard Logging Levels.....	13-23
13.7.4	Managing Log File Size and Disk Space Usage.....	13-25
13.7.4.1	Managing Rotating Log File Space Usage for Java and SOA Applications....	13-26
13.7.4.2	Managing Rotating Log File Space Usage for PL/SQL Applications.....	13-27
13.7.4.3	Managing Log File Space Usage for C Applications.....	13-28
13.7.5	Using Profile Options to Configure Standard Log Settings.....	13-29
13.7.6	Using Additional Settings to Configure Logs for Selected Components.....	13-31

13.7.6.1	Configuring Additional Log Settings for Oracle Fusion Incentive Compensation Batch Jobs	13-32
13.7.6.2	Configuring Additional Log Settings for Oracle Fusion General Ledger	13-32
13.7.6.3	Configuring Additional Log Settings for Oracle Fusion Receivables AutoInvoice.....	13-32
13.7.7	Configuring Access to Logs for Fusion Applications Control	13-33
13.8	Standard Diagnostic Testing Administration Tasks and Tools	13-35
13.9	Configuring the Diagnostic Testing Framework for Normal Operation.....	13-36
13.9.1	Controlling Access to Diagnostic Testing Functionality	13-37
13.9.2	Navigating to the Diagnostic Dashboard.....	13-38
13.9.2.1	Navigating to the Diagnostic Dashboard from an Oracle Fusion Application.....	13-38
13.9.2.2	Navigating to the Diagnostic Dashboard from Cloud Control:	13-39
13.10	Using Diagnostic Tests to Monitor Normal System Health	13-39
13.10.1	Running Diagnostic Tests.....	13-40
13.10.1.1	Using the Diagnostic Dashboard to Run Diagnostic Tests.....	13-40
13.10.1.2	Using the diagctl Command Line Interface to Run Diagnostic Tests.....	13-44
13.10.2	Searching for Diagnostic Tests by Name, Categorization Tag, or Module	13-47
13.10.3	Checking Diagnostic Test Availability	13-47
13.10.4	Providing Input Parameters for Diagnostic Tests.....	13-48
13.10.5	Running Diagnostic Tests Immediately	13-48
13.10.6	Scheduling Diagnostic Tests to Run Later	13-48
13.10.7	Checking the Status of a Diagnostic Test	13-49
13.10.7.1	Using the Diagnostic Dashboard to Check the Status of a Diagnostic Test....	13-49
13.10.7.2	Using the diagctl Command Line Interface to Check the Status of a Diagnostic Test	13-50
13.10.8	Canceling Diagnostic Test Runs	13-51
13.10.9	Viewing the Results of Diagnostic Tests	13-53
13.10.9.1	Using the Diagnostic Dashboard to View the Results of Diagnostic Tests....	13-54
13.10.9.2	Using the diagctl Command Line Interface to View the Results of Diagnostic Tests	13-54
13.10.10	Identifying Diagnostic Test Launch Methods from Test Run Names	13-55
13.10.11	Purging the Results of Selected Diagnostic Test Runs	13-56

14 Troubleshooting Oracle Fusion Applications Using Incidents, Logs, QuickTrace, and Diagnostic Tests

14.1	Introduction to Troubleshooting Using Incidents, Logs, QuickTrace, and Diagnostic Tests..	14-1
14.1.1	Relationships Between Incidents, QuickTrace Dumps, Log Files, and Diagnostic Tests .	14-2
14.1.2	Introduction to the Oracle Fusion Middleware Diagnostic Framework	14-4
14.1.3	Standard Incident Administration Tasks and Tools.....	14-4
14.2	Configuring Log, QuickTrace, and Incident Settings for Troubleshooting	14-5
14.2.1	Assisting Users in Gathering Data Using Troubleshooting Options	14-5
14.2.2	Configuring the Oracle Fusion Middleware Diagnostic Framework	14-8
14.2.3	Precedence and Log Settings for Troubleshooting	14-8
14.2.4	Default System Settings for Incident Creation and QuickTrace	14-8
14.2.5	Adjusting Incident Settings for Troubleshooting.....	14-10

14.2.6	When Changes to Incident Configuration Settings Take Effect.....	14-11
14.2.7	Adjusting QuickTrace Configuration Settings for Troubleshooting.....	14-11
14.2.8	Adjusting Standard Log Levels for Troubleshooting.....	14-14
14.2.9	Adjusting Log Settings for Oracle Fusion Incentive Compensation Batch Jobs for Troubleshooting	14-14
14.2.10	Adjusting Oracle Fusion General Ledger Log Settings for Troubleshooting	14-15
14.2.11	Adjusting Oracle Fusion Receivables AutoInvoice Log Settings for Troubleshooting	14-16
14.2.12	Disabling Logging of a Particular Message	14-18
14.2.13	Logging More Detailed Information for a Particular Code Module.....	14-19
14.2.14	Configuring and Using Profile Options for Troubleshooting.....	14-19
14.3	Configuring the Diagnostic Testing Framework for Troubleshooting	14-20
14.4	Investigating, Reporting, and Solving a Problem	14-20
14.4.1	Process for Investigating, Reporting, and Solving a Problem.....	14-20
14.4.1.1	Process for Investigating, Reporting, and Solving a Problem Using Cloud Control .	14-21
14.4.1.2	General Process for Investigating, Reporting, and Solving a Problem.....	14-22
14.4.2	Managing Oracle Fusion Applications Problems and Incidents	14-24
14.4.2.1	Viewing Incidents Using Support Workbench	14-25
14.4.2.2	Viewing Incidents Using the Automatic Diagnostic Repository Command Interpreter (adrci)	14-28
14.4.2.3	Recovering from Incidents Generated During SOA Operations.....	14-30
14.4.2.4	Creating Incidents Manually Using Support Workbench.....	14-31
14.4.2.5	Packaging Incidents Using Support Workbench.....	14-33
14.4.2.6	Creating and Packaging Incidents Using the Automatic Diagnostic Repository Command Interpreter (adrci)	14-41
14.4.3	Working with Automatically Created QuickTrace Dumps.....	14-42
14.4.3.1	Viewing Automatically Created QuickTrace Dump Files Using Oracle WebLogic Scripting Tool	14-42
14.4.3.2	Viewing Automatically Created QuickTrace Dump Files Using the Automatic Diagnostic Repository Command Interpreter (adrci)	14-43
14.4.3.3	Viewing Automatically Created QuickTrace Dump Files Using Support Workbench	14-43
14.4.4	Working with Manually Created QuickTrace Dumps.....	14-44
14.4.5	Working with Other Diagnostic Dumps	14-44
14.4.6	Working with Cloud Control Problem Analysis and the Analyze Logs Options.	14-45
14.5	Troubleshooting Using Log Settings.....	14-46
14.5.1	Troubleshooting Using Standard Log Files	14-46
14.5.2	Viewing Special Log Output for Oracle Fusion Incentive Compensation.....	14-46
14.5.3	Correlating Log Messages Across Log Files and Components.....	14-47
14.5.3.1	Values for Correlating Log Messages.....	14-48
14.5.3.2	Correlating Log Messages for Oracle Enterprise Scheduler Jobs.....	14-48
14.5.3.3	Correlating Log Messages for SOA Services and BPEL Processes.....	14-49
14.5.3.4	Correlating Log Messages for Other Process Types.....	14-50
14.5.4	Downloading Log Files.....	14-51
14.6	Troubleshooting Using the Diagnostic Testing Framework	14-51
14.6.1	Working with Automatically Launched Diagnostic Tests	14-52
14.6.2	Registering Diagnostic Tests	14-52
14.6.3	Sending Diagnostic Test Results to Oracle Support	14-52

14.6.4	Purging the Results of Selected Diagnostic Test Runs	14-53
--------	--	-------

Part IV Advanced Administration

15 Backing Up and Recovering Oracle Fusion Applications

15.1	Introduction to Backup and Recovery	15-1
15.2	Overview of Backing Up Your Environment	15-2
15.2.1	Tools to Use to Back Up Your Environment.....	15-2
15.2.2	Modes of Backup	15-3
15.2.3	Types of Backups	15-3
15.2.4	Recommended Backup Strategy.....	15-4
15.3	Overview of Recovering Your Environment	15-6
15.3.1	Tools to Use to Recover Your Environment	15-6
15.3.2	Types of Recovery.....	15-6
15.3.3	Recommended Recovery Strategies	15-7
15.4	Prerequisites for Using Cloud Control to Back Up or Restore Your Environment.....	15-7
15.5	Backup and Recovery Recommendations for Oracle Fusion Applications.....	15-8
15.5.1	Backup and Recovery Recommendations for Oracle Fusion Customer Relationship Management	15-9
15.5.2	Backup and Recovery Recommendations for Oracle Fusion Financials	15-10
15.5.3	Backup and Recovery Recommendations for Oracle Fusion Human Capital Management	15-11
15.5.4	Backup and Recovery Recommendations for Oracle Fusion Supply Chain Management	15-12
15.5.5	Backup and Recovery Recommendations for Oracle Fusion Project.....	15-13
15.5.6	Backup and Recovery Recommendations for Oracle Fusion Procurement.....	15-14
15.5.7	Backup and Recovery Recommendations for Oracle Fusion Incentive Compensation....	15-15
15.5.8	Backup and Recovery Recommendations for Oracle Fusion Applications Technology ..	15-16
15.5.9	Backup and Recovery Recommendations for Oracle Fusion Setup	15-17
15.5.10	Backup and Recovery Recommendations for Oracle Enterprise Scheduler	15-18
15.5.11	Backup and Recovery Recommendations for Oracle Enterprise Crawl and Search Framework	15-18
15.5.12	Backup and Recovery Recommendations for Oracle Authorization Policy Manager	15-19
15.5.13	Backup and Recovery Recommendations for Java Servers for Oracle E-Mail and Web Marketing	15-20
15.6	Performing a Backup	15-20
15.6.1	Performing a Backup Using Cloud Control.....	15-21
15.6.1.1	Configuring Cloud Control Backups.....	15-21
15.6.1.2	Backing Up Oracle Fusion Applications Using Cloud Control.....	15-23
15.6.2	Performing a Full Offline Backup Using the Command Line.....	15-25
15.6.3	Performing an Online Backup of Configuration Files Using the Command Line	15-27
15.7	Creating a Record of Your Oracle Fusion Applications Configuration.....	15-27
15.8	Recovering Using Cloud Control	15-29
15.9	Recovering After Data Loss, Corruption, or Media Failure	15-30
15.9.1	Recovering the Applications Base Directory	15-31

15.9.2	Recovering a Middleware Home.....	15-31
15.9.3	Recovering an Oracle WebLogic Server Domain.....	15-32
15.9.4	Recovering an Oracle Home	15-32
15.9.5	Recovering an Oracle Instance Home.....	15-32
15.9.5.1	Recovering After Oracle Instance Home Deleted from File System.....	15-32
15.9.5.2	Recovering After Oracle Instance Home Deregistered.....	15-33
15.9.6	Recovering the Administration Server Configuration.....	15-33
15.9.7	Recovering a Managed Server	15-34
15.9.8	Recovering the Databases.....	15-35
15.9.9	Reconciling the Data.....	15-35
15.9.9.1	Recovering the Oracle Identity Manager Database and Reconciling It with the LDAP Database	15-35
15.9.9.2	Recovering the Oracle Fusion Applications Database and Reconciling It with the LDAP Database	15-36
15.9.9.3	Recovering the LDAP Database Using Multimaster Replication.....	15-37
15.9.10	Recovering Oracle Fusion Applications.....	15-37
15.9.10.1	Recovering Oracle Fusion Customer Relationship Management	15-37
15.9.10.2	Recovering Oracle Fusion Procurement.....	15-37
15.9.11	Recovering Components Related to Oracle Fusion Applications	15-38
15.9.11.1	Recovering Oracle HTTP Server	15-38
15.9.11.2	Recovering Oracle Enterprise Scheduler	15-38
15.9.11.3	Recovering Oracle Enterprise Crawl and Search Framework	15-38
15.9.11.4	Recovering Oracle Authorization Policy Manager.....	15-38
15.9.11.5	Recovering Oracle E-Mail and Web Marketing	15-39
15.10	Recovering After Loss of Host	15-39
15.10.1	Recovering the Applications Base Directory After Loss of Host.....	15-39
15.10.2	Recovering After Loss of Administration Server Host	15-39
15.10.2.1	Recovering the Administration Server to the Same Host.....	15-40
15.10.2.2	Recovering the Administration Server to a Different Host.....	15-40
15.10.3	Recovering After Loss of Managed Server Host.....	15-42
15.10.3.1	Recovering a Managed Server to the Same Host.....	15-42
15.10.3.2	Recovering a Managed Server to a Different Host.....	15-43
15.10.4	Recovering the Databases After Loss of Host	15-46
15.10.5	Additional Actions for Recovering Entities After Loss of Host.....	15-46
15.10.5.1	Changing the Host Name in the targets.xml File for Fusion Middleware Control	15-46
15.10.5.2	Modifying the FusionVirtualHost_x.conf File.....	15-46
15.10.5.3	Creating a New Machine for the New Host Name	15-47
15.10.5.4	Updating Oracle Inventory	15-47
15.10.6	Recovering Oracle Fusion Applications After Loss of Host.....	15-48
15.10.6.1	Recovering Oracle Fusion Supply Chain Management After Loss of Host ...	15-48
15.10.6.2	Recovering Oracle Fusion Procurement After Loss of Host	15-48
15.10.7	Recovering Components Related to Oracle Fusion Applications	15-48
15.10.7.1	Recovering Oracle HTTP Server After Loss of Host	15-49
15.10.7.2	Recovering Oracle Enterprise Scheduler After Loss of Host	15-49
15.10.7.3	Recovering Oracle Enterprise Crawl and Search Framework After Loss of Host..... 15-49	
15.10.7.4	Recovering Oracle Authorization Policy Manager After Loss of Host	15-50

15.10.7.5	Recovering Oracle E-Mail and Web Marketing After Loss of Host.....	15-50
15.10.7.6	Recovering Oracle Essbase In Clustered Environment After Loss of Host	15-50
15.11	A Case Study: Recovering Oracle Fusion Customer Relationship Management.....	15-50
15.11.1	The Recovery Case Study Scenario	15-50
15.11.2	Recovering the Middleware Home Containing Java EE Components	15-51
15.11.3	Recovering the Web Tier Middleware Home.....	15-52
15.11.4	Recovering the Oracle Fusion Customer Relationship Management Domain.....	15-53
15.11.5	Recovering Servers When the Installation Directory Is Shared Between Hosts ...	15-53
15.11.6	Recovering Servers When the Managed Server Configuration Is Local	15-54
15.11.7	Recovering an Oracle Instance in an Oracle Fusion Customer Relationship Management Installation	15-55
15.11.8	Recovering an Oracle Fusion Customer Relationship Management Cluster	15-56
15.11.9	Recovering Databases for Oracle Fusion Customer Relationship Management...	15-57

16 Configuring High Availability and Scaling Out Oracle Fusion Applications

16.1	Introduction to Oracle Fusion Applications High Availability Architecture	16-1
16.2	Oracle WebLogic Cluster Configuration for Oracle Fusion Applications.....	16-2
16.3	Oracle Fusion Applications High Availability	16-2
16.3.1	Oracle Fusion Applications Single Instance Architecture	16-2
16.3.1.1	Oracle Fusion Applications External Dependencies	16-3
16.3.1.2	Oracle Fusion Applications Run-Time Processes	16-3
16.3.1.3	Oracle Fusion Applications Request Flow.....	16-4
16.3.1.4	Oracle Fusion Applications Configuration Artifacts	16-4
16.3.1.5	Oracle Fusion Applications Deployment Artifacts	16-4
16.3.1.6	Oracle Fusion Applications Log File Location	16-4
16.3.2	Oracle Fusion Applications High Availability Architecture and Failover Considerations	16-4
16.3.2.1	Starting and Stopping the Oracle Fusion Applications Cluster	16-6
16.3.2.2	Oracle Fusion Applications Cluster-Wide Configuration Changes.....	16-6
16.3.2.3	Oracle Fusion Applications Failures and Expected Behaviors	16-6
16.4	Oracle RAC Database Configuration for Oracle Fusion Applications Repositories.....	16-7
16.4.1	Oracle Real Application Clusters	16-8
16.4.2	Providing High Availability for Oracle Fusion Applications Repositories	16-9
16.4.2.1	Preparing the System to Use Oracle RAC.....	16-9
16.4.2.2	Providing High Availability Through Multi Data Sources	16-9
16.4.2.3	Failover Processing and Recovery for Oracle RAC Node or Instance Failures.....	16-10
16.4.2.4	Load Balancing Across Oracle RAC Nodes.....	16-10
16.4.2.5	Retry Logic to Protect Read-Only Operations.....	16-10
16.4.3	Best Practices for Deploying JDBC Multi Data Sources on Servers and Clusters.	16-10
16.5	Scaling Out Oracle Fusion Applications.....	16-11
16.5.1	Scaling Up the Topology (Adding Managed Servers to Existing Machines)	16-12
16.5.1.1	Testing the Routing from Oracle HTTP Server to the New Managed Server	16-15
16.5.2	Scaling Out the Topology (Adding Managed Servers to New Machines).....	16-15
16.5.3	Scaling Out Oracle RAC Databases.....	16-17
16.5.3.1	Cloning Oracle RAC Nodes to Scale Out Oracle Fusion Applications	16-17
16.5.3.2	Configuring the Multi Data Source to Include New Oracle RAC Nodes	16-17

17 Moving Components for Oracle Fusion Applications Across Environments

17.1	Introduction to Moving Oracle Fusion Applications	17-1
17.2	Understanding the Full-Movement Tasks	17-3
17.3	Getting Started with Moving Components.....	17-5
17.3.1	Preparing the Source Environment.....	17-5
17.3.2	Preparing the Target Environment	17-5
17.4	Moving the Oracle Identity Management Domain Component Artifacts	17-7
17.4.1	Moving the Security Artifacts in the Identity Store.....	17-8
17.4.2	Moving the Policy and Credential Store by Seeding the Oracle ADF Credentials	17-12
17.4.3	Moving Data Security Policies	17-14
17.4.4	Modifying the CSF Entries in the Target Environment	17-14
17.4.5	Moving Oracle Access Manager Artifacts.....	17-18
17.4.6	Configure Oracle Identity Manager for Oracle Fusion Applications	17-19
17.4.7	Moving Oracle Identity Federation Artifacts	17-20
17.5	Moving Binary Files to a Production Target Environment.....	17-20
17.5.1	Creating Binary Archives	17-21
17.5.2	Applying the Binary Files to the Target Environment.....	17-23
17.5.3	Moving the Oracle Database Client Installation	17-26
17.6	Moving Configurations to a Production Target Environment.....	17-27
17.6.1	Understanding the Configuration Scripts.....	17-28
17.6.1.1	copyConfig Script	17-29
17.6.1.2	extractMovePlan Script.....	17-30
17.6.1.3	pasteConfig Script	17-30
17.6.2	Setting the JVM System Variable on Windows	17-30
17.6.3	Moving the Webgate Installation	17-30
17.6.3.1	Moving the Webgate Installation on UNIX.....	17-30
17.6.3.2	Moving the Webgate Installation on Windows	17-31
17.6.4	Creating Configuration Archives.....	17-32
17.6.4.1	Running copyConfig for Oracle WebLogic Server Domains.....	17-33
17.6.4.2	Running copyConfig for System Components	17-35
17.6.4.3	Running copyConfig for Node Manager	17-36
17.6.4.4	Running copyConfig for Oracle BI Enterprise Edition	17-37
17.6.4.5	Running copyConfig for the Oracle Global Order Promising Server Instance.....	17-38
17.6.5	Extracting the Move Plan	17-39
17.6.6	Modifying the Move Plan.....	17-41
17.6.6.1	Editing Oracle Fusion Applications FusionAppsTopology Properties.....	17-42
17.6.6.2	Editing Oracle Fusion Applications Patching Framework Properties	17-44
17.6.6.3	Editing Additional Oracle Fusion Applications Environment Properties	17-44
17.6.6.4	Editing Oracle BI EE Properties.....	17-46
17.6.6.5	Editing Oracle Data Integrator Properties	17-56
17.6.6.6	Editing Oracle Enterprise Crawl and Search Framework Properties	17-57
17.6.6.7	Editing Oracle Enterprise Scheduler Properties	17-58
17.6.6.8	Editing Oracle Global Order Promising Properties.....	17-59
17.6.6.9	Editing Oracle SES Properties.....	17-59
17.6.6.10	Editing Oracle Topology Manager Properties	17-61
17.6.6.11	Editing Oracle WebCenter Content: Imaging Properties	17-63

17.6.7	Applying the Configurations to the Target Environment	17-64
17.6.7.1	Running pasteConfig for Oracle WebLogic Server Domains	17-65
17.6.7.2	Running pasteConfig for System Components.....	17-66
17.6.7.3	Running pasteConfig for Node Manager	17-68
17.6.7.4	Running pasteConfig for Oracle BI Enterprise Edition	17-70
17.6.7.5	Running pasteConfig for Oracle Global Order Promising Server Instance....	17-71
17.6.8	Completing the Component-Specific Configuration Move	17-73
17.6.8.1	Moving Oracle BI Enterprise Edition Components.....	17-73
17.6.8.2	Moving Oracle BI Applications Components	17-75
17.6.8.3	Moving the Oracle Governance, Risk and Controls Client Configuration	17-76
17.6.8.4	Moving Oracle SES.....	17-76
17.6.8.5	Moving Oracle WebLogic Communication Services	17-76
17.6.9	Applying Domain Configurations to Other Machines	17-76
17.7	Moving Binary and Configuration Files to a Non-Distributed Target Environment..	17-77
17.8	Moving Functional Setup Data	17-83
17.9	Completing Post-Movement Tasks	17-88
17.9.1	Completing Post-Movement Tasks for a Target Distributed Environment.....	17-88
17.9.1.1	Completing Post-Movement Tasks for a Target Distributed Environment....	17-88
17.9.1.2	Completing Post-Movement Tasks on a Target Non-Distributed Environment	17-89
17.9.2	Tips for Completing the Oracle RAC Configuration in the Target Environment.	17-91
17.10	A Case Study: Moving Oracle Fusion Supply Chain Management Components to a Target Distributed Environment	17-91

Part V Appendixes

A High Availability for Oracle Fusion Middleware Extensions for Applications

A.1	How Oracle Fusion Middleware Extensions Components Use Fusion Middleware Components for High Availability and Failover	A-1
A.1.1	MDS	A-1
A.1.2	ADF.....	A-2
A.1.3	WebLogic Server Failover for Session Content	A-2

B High Availability for Oracle Enterprise Scheduler

B.1	Introduction to High Availability for Oracle Enterprise Scheduler	B-1
B.2	Oracle Enterprise Scheduler Concepts	B-2
B.3	Configuring High Availability for Oracle Enterprise Scheduler	B-2
B.3.1	Oracle Enterprise Scheduler Configuration and Deployment Artifacts.....	B-2
B.3.2	Oracle Enterprise Scheduler Logging	B-3
B.3.3	Backup and Recovery	B-3
B.4	Managing an Oracle Enterprise Scheduler Cluster.....	B-3

C Seeded Searchable Objects, Search Categories, Index Schedules, and Facets

C.1	Searchable Objects by Product Family.....	C-1
C.2	Search Categories by Product Family	C-5
C.3	Index Schedules by Product Family.....	C-8

C.4 Facets by Product Family..... C-12

Preface

This guide describes how to administer the Oracle Fusion Applications environment, including how to start and stop components, change ports and passwords, manage the Oracle Database, monitor the environment, logging, back up and recover, configure high availability, move components across environments, and troubleshoot runtime issues.

Audience

This document is intended for administrators of the Oracle Fusion Applications environment. The Oracle Fusion Applications administrator performs tasks to manage the Oracle Fusion Applications environment. These tasks are performed to the Oracle Fusion applications, Oracle Fusion Middleware components, and the Oracle Database.

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 Oracle Fusion Applications documentation set:

- *Oracle Fusion Applications Administrator and Implementor Roadmap*
- *Oracle Fusion Applications Administrator's Troubleshooting Guide*
- *Oracle Fusion Applications Common Implementation Guide*
- *Oracle Fusion Applications Customer Relationship Management Enterprise Deployment Guide*
- *Oracle Fusion Applications Master Glossary*
- *Oracle Fusion Applications Post-Installation Guide*

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.

UNIX is used as a generic reference to all UNIX-based platforms. Where information for a particular UNIX platform is different, this is noted in the text.

What's New in This Guide for Release 11.1.3

For Release 1 (11.1.3), this guide has been updated in several ways. The following table lists the chapters or sections that have been added or changed.

Chapter/Section	Change Description
Chapter 2, "Getting Started with Administering Oracle Fusion Applications"	
Section 2.7.1, "Starting Cloud Control and Discovering Oracle Fusion Applications Targets"	Revised this section to include information on enabling business metrics in Cloud Control.
Chapter 3, "Performing Routine Administrative Tasks"	
Section 3.3, "Starting and Stopping"	Revised this section to clarify starting and stopping Node Manager, Administration Servers, and Managed Servers.
Section 3.8, "Managing Oracle Fusion Applications-Specific Labels in the Oracle Metadata Repository"	Added this section to describe how to manage labels for Oracle Fusion Applications in the Oracle Metadata Repository.
Section 3.9, "Modifying Oracle Application Development Framework Connections in Oracle Fusion Applications"	Added this section to describe how to modify the connections for all the applications in an Oracle WebLogic Server domain using Fusion Applications Control.
Chapter 5, "Managing Oracle Enterprise Scheduler Service and Jobs"	Removed redundant information already captured in the <i>Oracle Fusion Middleware Administrator's Guide for Oracle Enterprise Scheduling Service</i> .
Chapter 10, "Monitoring the Oracle Fusion Applications Middle Tier"	Modified this chapter to reflect performing procedures through Cloud Control
Section 10.5.2, "Creating Monitoring Templates"	Removed metric tables already present in the <i>Oracle Fusion Applications Performance and Tuning Guide</i> .
Chapter 11, "Diagnosing the Oracle Fusion Applications Middle Tier Performance"	Modified this chapter to reflect performing procedures through Cloud Control
Chapter 13, "Managing Oracle Fusion Applications Log Files and Diagnostic Tests"	

Chapter/Section	Change Description
Section 13.6.2, "Searching for Specific Information in Log Files within a Single Domain"	Revised section to add details about selected targets within product families or clusters, available search operators, and use of wildcard characters.
Section 13.6.3, "Searching for Specific Information in Log Files Across Multiple Domains"	Added this section to describe how to use Cloud Control to search log files in multiple domains within a single Fusion instance.
Section 13.10.1.1, "Using the Diagnostic Dashboard to Run Diagnostic Tests"	Revised section to correct detailed instructions on how to supply input parameters in the Diagnostic Dashboard.
Section 13.10.1.2, "Using the diagctl Command Line Interface to Run Diagnostic Tests"	Revised this section to add instructions for setting two environment variables that the <code>diagctl</code> command line interface requires.
Section 13.10.2, "Searching for Diagnostic Tests by Name, Categorization Tag, or Module"	Revised this section to clarify the available ways to search for diagnostic tests that have specific characteristics.
Section 13.10.4, "Providing Input Parameters for Diagnostic Tests"	Revised this section to add information about the message that the Diagnostic Dashboard displays when all required inputs for a diagnostic test have been supplied.
Section 13.10.7.2, "Using the diagctl Command Line Interface to Check the Status of a Diagnostic Test"	Revised this section to clarify the location of the <code>diagctl</code> command line interface in UNIX and Windows environments.
Section 13.10.9.2, "Using the diagctl Command Line Interface to View the Results of Diagnostic Tests"	Revised this section to clarify the location of the <code>diagctl</code> command line interface in UNIX and Windows environments.
Section 13.10.10, "Identifying Diagnostic Test Launch Methods from Test Run Names"	Revised this section to clarify the circumstances in which it may be useful to know test run naming conventions.
Chapter 14, "Troubleshooting Oracle Fusion Applications Using Incidents, Logs, QuickTrace, and Diagnostic Tests"	
Section 14.4.1.1, "Process for Investigating, Reporting, and Solving a Problem Using Cloud Control"	Revised this section to clarify how to work with diagnostic tests during the troubleshooting process in an environment that includes Cloud Control.
Section 14.4.2.1, "Viewing Incidents Using Support Workbench"	Revised this section to add information about locating and using the Support Workbench rollup page to work with problems and incidents.

Chapter/Section	Change Description
Section 14.4.2.2, "Viewing Incidents Using the Automatic Diagnostic Repository Command Interpreter (adrci)"	Revised this section to clarify the location of the Automatic Diagnostic Repository Command Interpreter (adrci).
Section 14.4.2.4, "Creating Incidents Manually Using Support Workbench"	Added this section to describe how to create incidents manually in an environment that includes Cloud Control.
Section 14.4.2.5, "Packaging Incidents Using Support Workbench"	Added this section to describe how to package incidents for transmittal to Oracle Support in an environment that includes Cloud Control.
Section 14.4.2.6, "Creating and Packaging Incidents Using the Automatic Diagnostic Repository Command Interpreter (adrci)"	Revised this section to clarify when to use <code>adrci</code> to create and package incidents, rather than using Support Workbench to do so.
Section 14.4.5, "Working with Other Diagnostic Dumps"	Revised this section to clarify the behavior of incidents generated from SOA code.
Chapter 17, "Moving Components for Oracle Fusion Applications Across Environments"	Updated this chapter with inclusion of automated scripts, streamlining the movement process, especially for non-distributed environments.

Part I

Understanding Oracle Fusion Applications

This part provides an overview to Oracle Fusion Applications and its concepts as they relate to administering Oracle Fusion Applications.

Part I contains the following chapters:

- [Chapter 1, "Introduction to Oracle Fusion Applications for Systems Administrators"](#)

Introduction to Oracle Fusion Applications for Systems Administrators

Oracle Fusion Applications is Oracle's next-generation applications suite built on a service-oriented platform. It brings together next-generation enterprise technologies, applications, and services, including Oracle Fusion Applications and Oracle Fusion Middleware, to fundamentally change the dynamics in the applications marketplace and revolutionize business. This chapter provides an introduction to the architecture, components, and concepts of Oracle Fusion Applications.

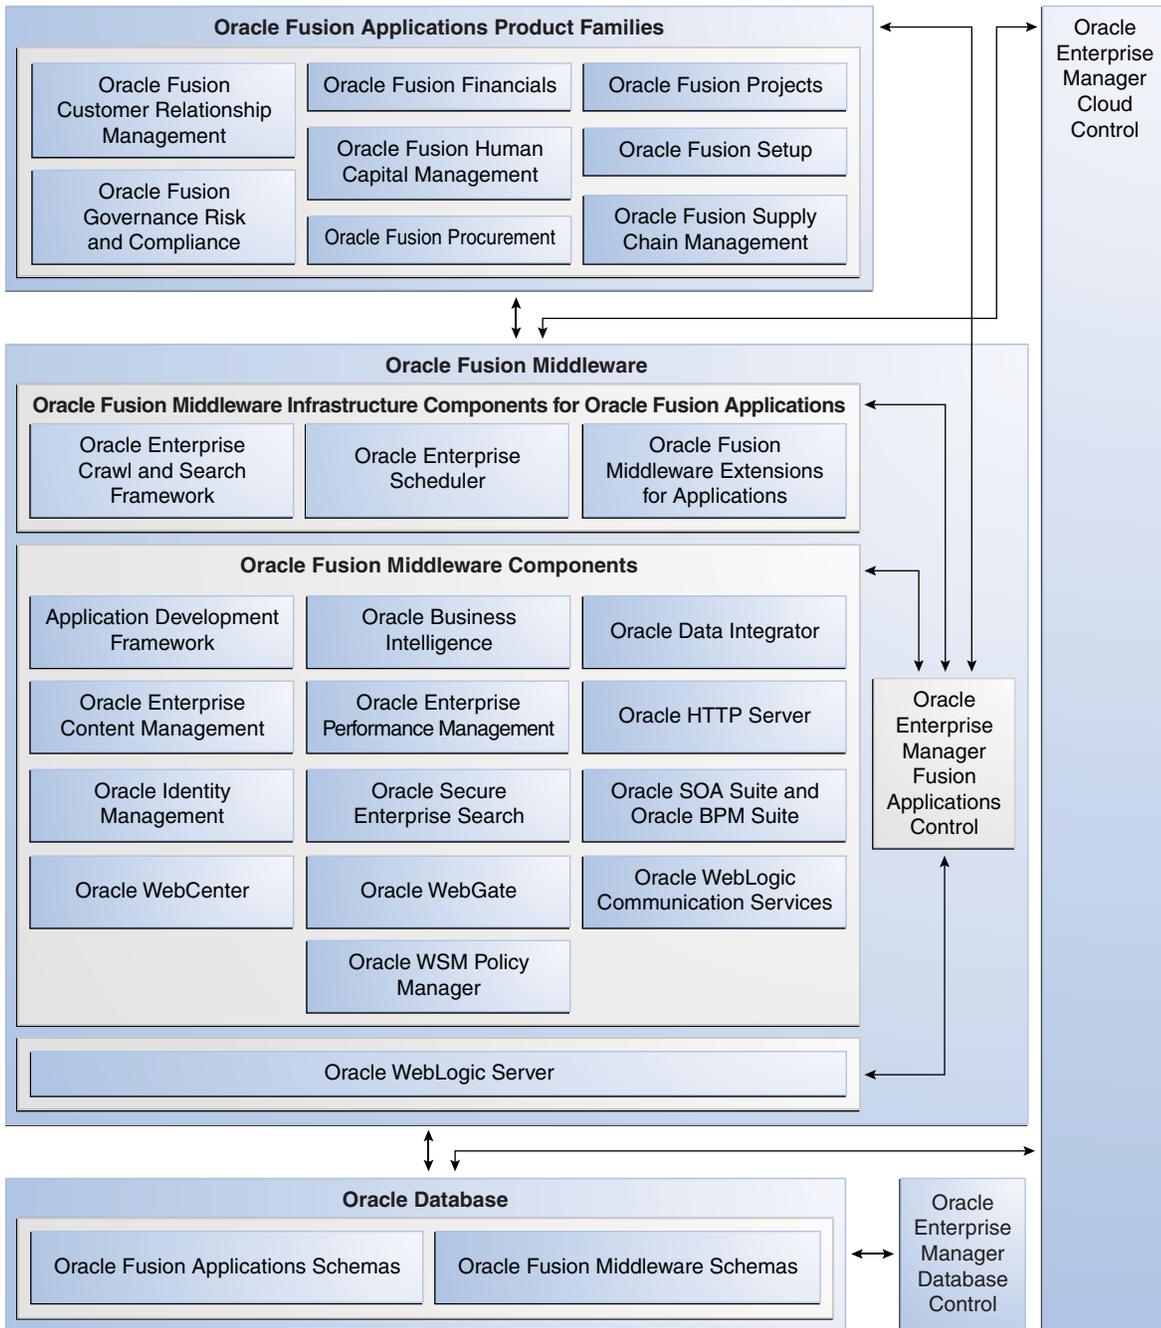
This chapter contains the following topics:

- [Architecture of Oracle Fusion Applications](#)
- [Key Oracle Fusion Applications Concepts](#)
- [Roadmap for Administering Oracle Fusion Applications](#)

1.1 Architecture of Oracle Fusion Applications

Oracle Fusion applications are built on the Oracle Fusion Middleware stack and utilize the Oracle Database. This architecture is depicted in [Figure 1-1](#).

Figure 1-1 Oracle Fusion Applications Architecture



The section contains the following topics

- [Oracle Fusion Applications Product Families](#)
- [Oracle Fusion Middleware Infrastructure Components](#)
- [Oracle Fusion Middleware Components](#)
- [Oracle Database](#)
- [Oracle Enterprise Manager Fusion Applications Control](#)
- [Oracle Enterprise Manager Cloud Control](#)

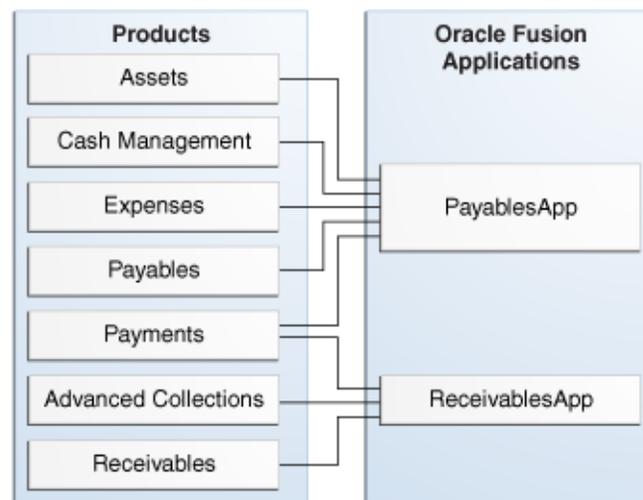
- [Oracle Enterprise Manager Database Control](#)

1.1.1 Oracle Fusion Applications Product Families

An installation of Oracle Fusion Applications is logically broken up into groups of features known as **product offerings**, which represent the highest-level collection of functionality that you can license and implement. A provisioning configuration is a collection of one or more product offerings. During installation, you select a product offering or a combination of offerings as a way to install the **product families**. Product families are comprised of one or more Java EE applications specific to Oracle Fusion Applications.

An application can contain multiple products, and a product can also span multiple applications. A product typically has a one-to-one correspondence with an EAR file. [Figure 1-2](#) shows two applications, PayablesApp and ReceivablesApp, each of which is mapped to the same product, Payments.

Figure 1-2 Relationship of Products and Applications



Oracle Fusion Applications includes the product families described in [Table 1-1](#).

Table 1-1 Product Families

Oracle Fusion Applications Product Families	Description
Oracle Fusion Customer Relationship Management	Manages customers, contacts, and resources, including data quality configuration.
Oracle Fusion Financials	Manages financial flows, including assets, ledgers, cash cycle, invoices and payments, accounts receivable, collections, and setup of subledger accounting and tax configuration.

Table 1–1 (Cont.) Product Families

Oracle Fusion Applications Product Families	Description
Oracle Fusion Governance, Risk, and Compliance	Provides critical business controls to manage risk, multi-regulatory compliance, and controls enforcement. The connector for Oracle Fusion Applications provides a prebuilt solution for managing Separation of Duties (SoD) within and across product families. You can also utilize Oracle Fusion Governance, Risk, and Compliance to analyze suspect transactions and configuration settings based on user defined conditions. This allows organizations to actively determine the risk that exists within their application that can materially impact the reliability of the information that exists for reporting and decision making purposes. Finally, Oracle Fusion Governance, Risk, and Compliance can apply preventive controls that will limit what a user can see and do within an Oracle Fusion Applications user interface according to user-defined conditions. The objective is to pro-actively mitigate the risk of extraneous access or improper transactions from existing.
Oracle Fusion Human Capital Management	Provides employee management for an organization.
Oracle Fusion Procurement	Manages the procurement process including requisitions, purchase orders, and supplier negotiations.
Oracle Fusion Project	Manages projects, including how to plan, budget, forecast, collect costs, bill customers, and report performance.
Oracle Fusion Supply Chain Management	Integrates and automates all key supply chain processes, from design, planning and procurement to manufacturing and fulfillment, providing a complete solution set to enable companies to power information-driven value chains.
Oracle Fusion Setup	<p>Supports the other product families.</p> <p>In addition to Oracle Fusion Functional Setup Manager for setting up functional data, this product family includes applications to assist application users:</p> <ul style="list-style-type: none"> ■ The Oracle Fusion Home page provides a Welcome dashboard with a collection of portlets and task flows for answering common questions. ■ Oracle Fusion Applications Help delivers content users need to complete their tasks. You can optionally install a local version of Oracle Fusion Applications Help, enabling you to extend and customize the help. For more information about using and customizing a local install of the help, see the <i>Oracle Fusion Applications Installation Guide</i> and the <i>Oracle Fusion Applications Common Implementation Guide</i>.

To aid with performing current-state analysis of the applications, you can use Oracle Transactional Business Intelligence, which is built on Oracle BI Enterprise Edition. For more information about configuring Oracle Transactional Business Intelligence, see the *Oracle Fusion Administrator's Guide for Transactional Business Intelligence*.

1.1.2 Oracle Fusion Middleware Infrastructure Components

The product families use the following common core framework and infrastructure for Oracle Fusion Applications described in [Table 1–2](#).

Table 1–2 Oracle Fusion Middleware Infrastructure Components

Oracle Fusion Middleware Infrastructure Components	Description
Oracle Fusion Middleware Extensions for Applications (Applications Core)	<p>Provides design time and runtime infrastructure to help standardize complex development patterns for Oracle Fusion Applications. It simplifies the development process of these patterns and provides a consistent user experience. Examples of these patterns include extensibility (Flexfields), hierarchical relationships (Trees), data security, and UI patterns. Applications Core creates simplified methods of implementing these complex requirements by providing robust metadata and comprehensive UI components and services. All of the Applications Core components have been intricately integrated with the rest of the Oracle Fusion Middleware infrastructure so they are available across every layer of the Oracle Fusion Applications platform.</p> <p>Applications Core provides shared libraries referenced by all the Oracle Fusion Applications, a standalone application for application setup and configuration, an Oracle JDeveloper extension to seamlessly integrate our components with the rest of the Oracle Fusion Applications technology stack, PLSQL API's, C libraries, and common seed data.</p> <p>See the "Oracle Fusion Middleware Extensions for Applications" section in the <i>Oracle Fusion Applications Concepts Guide</i>.</p>
Oracle Enterprise Scheduler	<p>Enables you to manage and schedule jobs for Oracle Fusion Applications.</p> <p>See Chapter 5, "Managing Oracle Enterprise Scheduler Service and Jobs."</p>
Oracle Enterprise Crawl and Search Framework (ECSF)	<p>Oracle Enterprise Crawl and Search Framework (ECSF) enables Oracle Fusion Applications Search for performing full-text searches securely and simultaneously against multiple logical business objects. Any application that connects to multiple data sources or manages a significant amount of unstructured (non-database) information—or both—needs advanced search capabilities so that application users can easily locate and take action on data that is relevant to them.</p> <p>See Chapter 7, "Managing Search with Oracle Enterprise Crawl and Search Framework."</p>

1.1.3 Oracle Fusion Middleware Components

Oracle Fusion Middleware includes the following components describes in [Table 1–3](#).

Table 1–3 Oracle Fusion Middleware Components

Oracle Fusion Middleware Components	Description
Oracle Application Development Framework (Oracle ADF)	<p>Provides an end-to-end application framework that builds on Java Platform, Enterprise Edition (Java EE) standards and open-source technologies to simplify and accelerate implementing service-oriented applications.</p> <p>See the <i>Oracle Fusion Middleware Getting Started with Oracle Data Integrator</i>.</p>

Table 1–3 (Cont.) Oracle Fusion Middleware Components

Oracle Fusion Middleware Components	Description
Oracle Business Intelligence	<p>Oracle Business Intelligence provides a complete, integrated solution of analytics and reporting for Oracle Fusion Applications.</p> <p>See the following documentation on Oracle Business Intelligence components in Oracle Fusion Applications:</p> <ul style="list-style-type: none"> ■ Oracle BI Enterprise Edition: <i>Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition</i> ■ Oracle Business Intelligence Publisher: Chapter 6 ■ Oracle Essbase: <i>Oracle Essbase Database Administrator's Guide</i> and <i>Oracle Essbase Technical Reference</i> ■ Oracle Enterprise Performance Management: http://www.oracle.com/technetwork/middleware/performance-management/documentation/index.html ■ Oracle Real-Time Decisions: <i>Oracle Fusion Middleware Platform Developer's Guide for Oracle Real-Time Decisions</i> ■ Oracle Transactional Business Intelligence: <i>Oracle Fusion Administrator's Guide for Transactional Business Intelligence</i> <p>After the Oracle Fusion applications are running, you can configure the Oracle Business Intelligence Applications for analyzing historical data. See the <i>Oracle Fusion Middleware Installation and Configuration Guide for Oracle Business Intelligence Applications</i>.</p>
Oracle Data Integrator	<p>Extracts, transforms, and loads data for the product families.</p> <p>See the <i>Oracle Fusion Middleware Developer's Guide for Oracle Data Integrator</i>.</p>
Oracle HTTP Server	<p>Provides a web listener for applications and the framework for hosting static and dynamic pages and applications over the web. Based on the proven technology of the Apache HTTP Server, Oracle HTTP Server includes significant enhancements that facilitate load balancing, administration, and configuration.</p> <p>See the <i>Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server</i>.</p>

Table 1–3 (Cont.) Oracle Fusion Middleware Components

Oracle Fusion Middleware Components	Description
Oracle Identity Management	<p>Provides a shared infrastructure for all applications, enabling developers to incorporate identity management into applications.</p> <p>For more information about documentation for Oracle Identity Management components in Oracle Fusion Applications, see the following guides:</p> <ul style="list-style-type: none"> ■ Oracle Platform Security Services (OPSS): <i>Oracle Fusion Middleware Application Security Guide</i> and <i>Oracle Fusion Middleware Security Overview</i> ■ Oracle Access Manager: <i>Oracle Fusion Middleware Administrator's Guide for Oracle Access Manager with Oracle Security Token Service</i> ■ Oracle Authorization Policy Manager: <i>Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)</i> ■ Oracle Identity Manager: <i>Oracle Fusion Middleware Administrator's Guide for Oracle Identity Manager</i> ■ Oracle Internet Directory: <i>Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory</i> ■ Oracle Virtual Directory: <i>Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory</i>
Oracle SOA Suite	<p>Provides a complete set of service infrastructure components for designing, deploying, and managing composite applications. Oracle SOA Suite enables services to be created, managed, and orchestrated into composite applications and business processes. Composites enable you to easily assemble multiple technology components into one SOA composite application.</p> <p>An important component of Oracle SOA Suite is Oracle WSM Policy Manager. Oracle WSM Policy Manager provides the infrastructure for enforcing global security and auditing policies. By securing various endpoints and setting and propagating identity, it secures applications. Oracle WSM Policy Manager provides a standard mechanism for signing messages, performing encryption, performing authentication, and providing role-based access control. You also can change a policy without having to change the endpoints or clients for this endpoints, providing greater flexibility and security monitoring for your enterprise.</p> <p>The Oracle Business Process Management (Oracle BPM) Suite provides an integrated environment for developing, administering, and using business applications centered around business processes. The Oracle BPM Suite is layered on the Oracle SOA Suite and shares many of the same product components.</p> <p>See the <i>Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite</i>.</p>

Table 1–3 (Cont.) Oracle Fusion Middleware Components

Oracle Fusion Middleware Components	Description
Oracle Secure Enterprise Search (Oracle SES)	<p>Provides a search engine for Oracle Fusion Applications Search. For more information on Oracle SES, see the following guides:</p> <ul style="list-style-type: none"> ■ <i>Oracle Secure Enterprise Search Administrator's Guide</i> ■ <i>Oracle Secure Enterprise Search Administration API Guide</i> ■ <i>Oracle Secure Enterprise Search Java API Reference</i> <p>For more information on how Oracle SES and ECSF provide Oracle Fusion Applications Search, see Chapter 7, "Managing Search with Oracle Enterprise Crawl and Search Framework."</p>
Oracle WebCenter Content	<p>Provides a comprehensive suite of digital content management tools. These tools can be used across the enterprise to cohesively track, manage, and dispose of content whether written, in digital images, or as email.</p> <p>For more information about documentation for Oracle Enterprise Content Management components in Oracle Fusion Applications, see:</p> <ul style="list-style-type: none"> ■ <i>WebCenter Content: Oracle WebCenter Content System Administrator's Guide for Content Server</i> ■ <i>Oracle WebCenter Content: Imaging: Oracle WebCenter Content Administrator's Guide for Imaging</i>
Oracle WebCenter Portal	<p>Enables you to create social applications, enterprise portals, collaborative communities, and composite applications, built on a standards-based, service-oriented architecture. Oracle WebCenter Portal combines dynamic user interface technologies with which to develop rich internet applications, the flexibility and power of an integrated, multichannel portal framework, and a set of horizontal Enterprise 2.0 capabilities delivered as services that provide content, collaboration, presence, and social networking capabilities. Based on these components, Oracle WebCenter Portal also provides an out-of-the-box, enterprise-ready customizable application, WebCenter Spaces, with a configurable work environment that enables individuals and groups to work and collaborate more effectively.</p> <p>See the <i>Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal</i>.</p>
Oracle WebGate	<p>Acts as a communicator plug-in that accepts users requests through Oracle HTTP Server and communicates with Oracle Access Manager.</p> <p>See the <i>Oracle Fusion Middleware Administrator's Guide for Oracle Access Manager with Oracle Security Token Service</i>.</p>
Oracle WebLogic Communication Services	<p>Provides click-to-dial functionality for applications primarily through contextual actions. Contextual actions provide related information and actions to users within the immediate context of the object instances upon which they act.</p> <p>For more information for enabling these services, see Chapter 8, "Managing Oracle WebLogic Communication Services for Click-to-Dial Functionality."</p>
Oracle WebLogic Server	<p>Supports the deployment of mission-critical applications in a robust, secure, highly available, and scalable environment. Oracle WebLogic Server is an ideal foundation for building applications based on service-oriented architecture (SOA).</p> <p>See the <i>Oracle Fusion Middleware Introduction to Oracle WebLogic Server</i>.</p>

1.1.4 Oracle Database

The Oracle Database contains the schemas and tablespaces required for both the Oracle Fusion Applications and for your applications. Oracle Fusion Applications does not support other databases.

Oracle Fusion Applications encryption APIs mask data such as credit card numbers in application user interface fields. For encryption and masking beyond that, Transparent Data Encryption (TDE) and Oracle Database Vault (ODV) are certified but optional with Oracle Fusion Applications.

TDE and ODV provide information lifecycle protections, such as the following:

- Data access restrictions on database administrators and other privileged users
- Sensitive data at rest in database files and file backups
- Sensitive data in transit
- Sensitive attributes in non-production databases

ODV establishes limitations on the power of privileged users to access sensitive data through segregation of duties policies on DBA roles and by securely consolidating application data in the database. These limitations prevent DBAs and other privileged users from overriding the protections placed on sensitive data by the Virtual Private Database (VPD). Oracle Fusion Applications deploys with the ODV enabled when it is installed.

TDE prevents access to PII in the file system or on backups or disk. TDE protects confidential data, such as credit card and social security numbers. TDE encrypts sensitive table data stored in data files at the tablespace level.

For more information about ODV and TDE, see the *Oracle Fusion Applications Security Guide*.

1.1.5 Oracle Enterprise Manager Fusion Applications Control

Oracle Enterprise Manager Fusion Applications Control (Fusion Applications Control) enables you to manage a single product family in an Oracle WebLogic Server domain for the Oracle Fusion Applications environment, including the products, applications, and Oracle Fusion Middleware components. As a part of management, you can monitor the runtime performance metrics for the various Oracle Fusion Applications and Oracle Fusion Middleware components. See [Section 2.5](#) for more information about Fusion Applications Control and [Section 2.2](#) for a comparison with Cloud Control.

1.1.6 Oracle Enterprise Manager Cloud Control

Oracle Enterprise Manager Cloud Control (Cloud Control) enables you to monitor and manage the complete IT infrastructure for Oracle Fusion Applications from a single console. You can monitor all the product families, Oracle Fusion Middleware components, and the Oracle Database. For example, you can monitor all the Oracle WebLogic Server domains for all the product families from one console. See [Section 2.7](#) for more information about Fusion Applications Control and [Section 2.2](#) for a comparison with Fusion Applications Control.

1.1.7 Oracle Enterprise Manager Database Control

Oracle Enterprise Manager Database Control (Database Control) enables you to manage the Oracle Database.

Using Database Control, you can perform administrative tasks such as creating schema objects (tables, views, indexes, and so on), managing user security, managing database memory and storage, backing up and recovering your database, and importing and exporting data. You can also view performance and status information about your database.

For more information, see the "Introduction to Oracle Enterprise Manager Database Control" section in the *Oracle Database 2 Day DBA*.

1.2 Key Oracle Fusion Applications Concepts

Before you begin administering Oracle Fusion Applications, it is important to understand the basics of installation, how Oracle WebLogic Server domains are configured, and the basic directory structure.

This section contains the following topics:

- [Provisioning and the Installation Process](#)
- [Oracle WebLogic Server Domains Configuration](#)
- [Provisioned Oracle Fusion Applications Home Directories](#)

1.2.1 Provisioning and the Installation Process

Provisioning is the entire set of operations required to install, configure, and deploy applications product offerings from a system point of view. It performs these operations:

- Installation provides the operations related to laying down all the component needed to create an Oracle Fusion Applications environment.
- Configuration tailors components based on the applications topology, the creation of Oracle WebLogic Server Managed Servers and clusters, and the updating of endpoints and virtual hosts.
- Deployment starts the Managed Servers and clusters and facilitates the actual use of product offerings.

This orchestration by a single processing engine ensures that all components interact smoothly and consistently in the applications environment.

For detailed information about provisioning and installing an Oracle Fusion Applications environment, see *Oracle Fusion Applications Installation Guide*.

This section contains the following topics:

- [Installation Options](#)
- [Installing an Oracle Database and Loading Content](#)

1.2.1.1 Installation Options

You choose options from the Provisioning Wizard menu to initiate one installation-related processes:

- Install an empty, single-instance Enterprise database
- Create a new provisioning plan
- Install, configure, and deploy the product offerings in a provisioning plan
- Uninstall an existing environment

1.2.1.2 Installing an Oracle Database and Loading Content

You must have installed and configured a transaction database before you install product offerings. You can use the Provisioning Wizard to create an empty, single-instance database instance. This is a discrete and separate task from the other provisioning options. Alternatively, you can install the database manually without using the wizard.

In either case, you finish the database installation by running the Oracle Fusion Applications Repository Creation Utility (Applications RCU) to load applications and middleware content into the database. This process creates the applications and middleware schemas, loads seed data, and creates the tablespaces, as well all other required packages.

For more information about creating a database with Applications RCU, see the *Oracle Fusion Applications Installation Guide*.

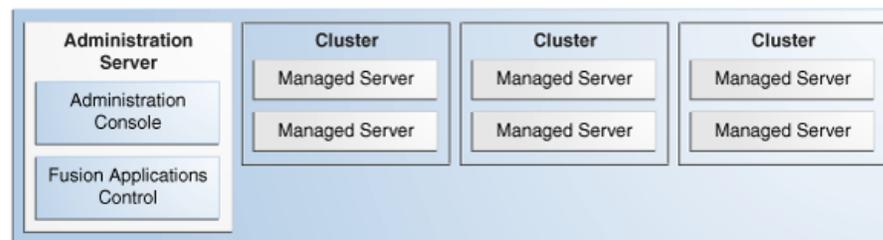
1.2.2 Oracle WebLogic Server Domains Configuration

See "What Is an Oracle WebLogic Server Domain?" in the *Oracle Fusion Middleware Administrator's Guide* to understand the following concepts for Oracle WebLogic Server:

- Oracle WebLogic Server domains
- Managed Servers and the Administration Server

During installation, the applications for a single product family are deployed to one Oracle WebLogic Server domain. Within a single domain, an Administration Server hosts the application for Fusion Applications Control, and the Java EE applications for the product family deploy to the Managed Servers. The Managed Servers are grouped together into a cluster for each application of the product family. [Figure 1–3](#) shows a domain with an Administration Server and three Managed Servers in clusters. If you want to scale out the Java EE application, you simply add new servers to the cluster. In an Oracle Fusion Applications environment, the domains are predefined so that the product families and their dependencies are always stored in a standardized arrangement.

Figure 1–3 Oracle WebLogic Server Domain



As an example of greater detail for a product family, [Figure 1–4](#) shows a portion of the domain for the Oracle Fusion Financials product family with an Administration Server and eight Managed Servers in a cluster. The applications are each deployed to a Managed Server within a cluster. In addition to the applications for Oracle Fusion Financials, the domain also contains an Oracle Enterprise Scheduler cluster and a SOA cluster. The Oracle Enterprise Scheduler (ESSAPP) application manages job requests for the product family. The SOA Infrastructure (*soa-infra*) application hosts the SOA composites for the applications in the product family.

Figure 1–4 Oracle WebLogic Server Domain for the Oracle Fusion Financials Family

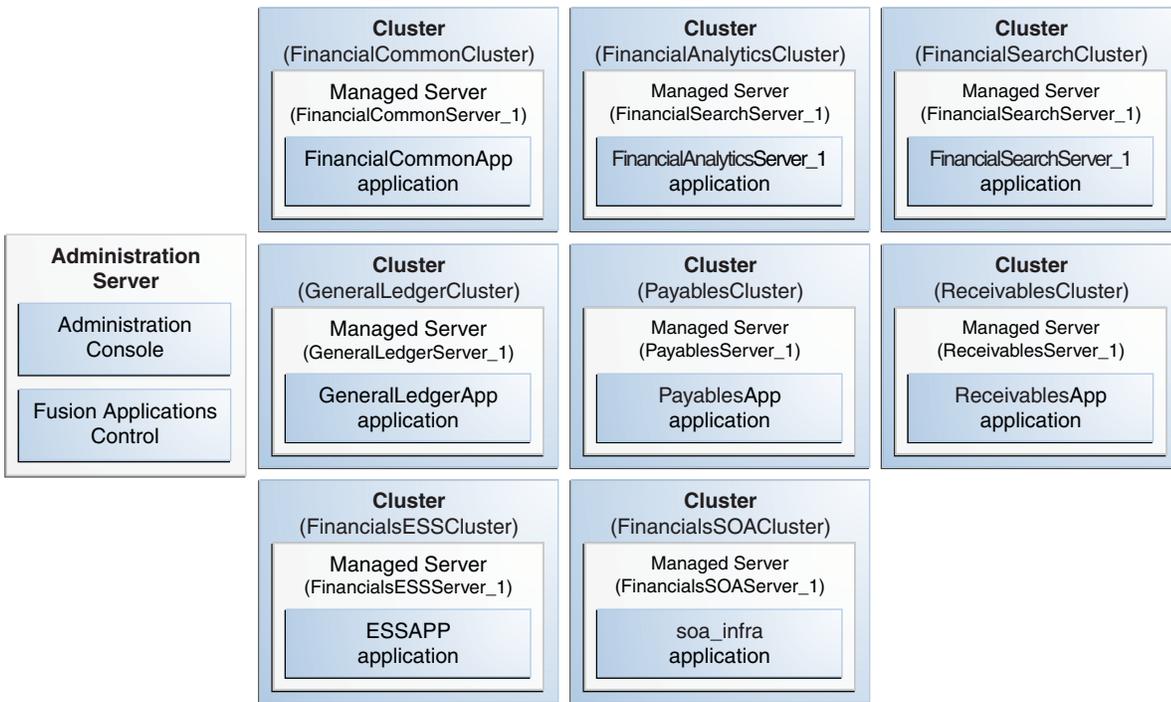
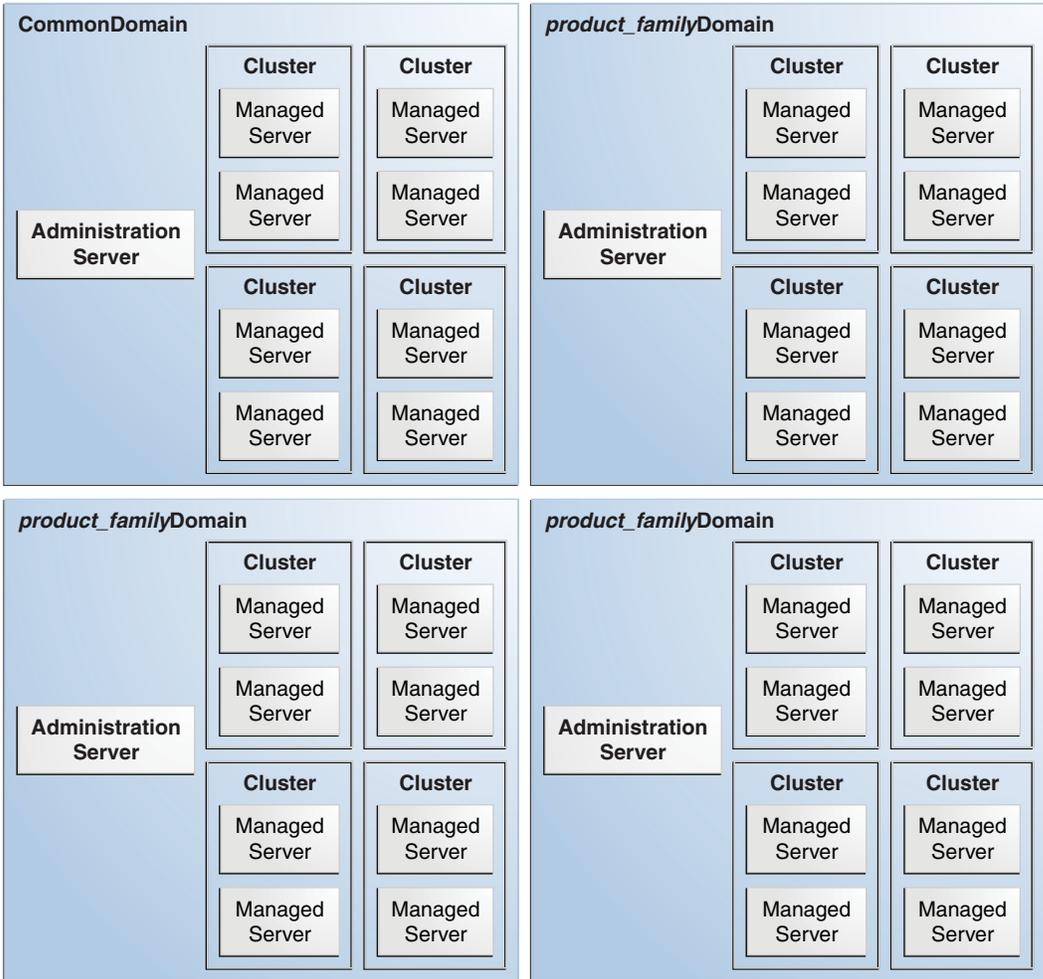


Figure 1–5 shows an Oracle Fusion Applications environment with multiple product families, each in separate Oracle WebLogic Server domains. The `CommonDomain` is the domain in the Oracle Fusion Setup product family.

Figure 1-5 Multiple Product Families



1.2.3 Provisioned Oracle Fusion Applications Home Directories

An installation of Oracle Fusion Applications contains two mount point directories, *APPLICATIONS_BASE* for the binaries and *APPLICATIONS_CONFIG* for the configuration files. The *APPLICATIONS_BASE* and *APPLICATIONS_CONFIG* directories contain various types of home directories. The home directories described in [Table 1-4](#).

Table 1–4 Home Directories

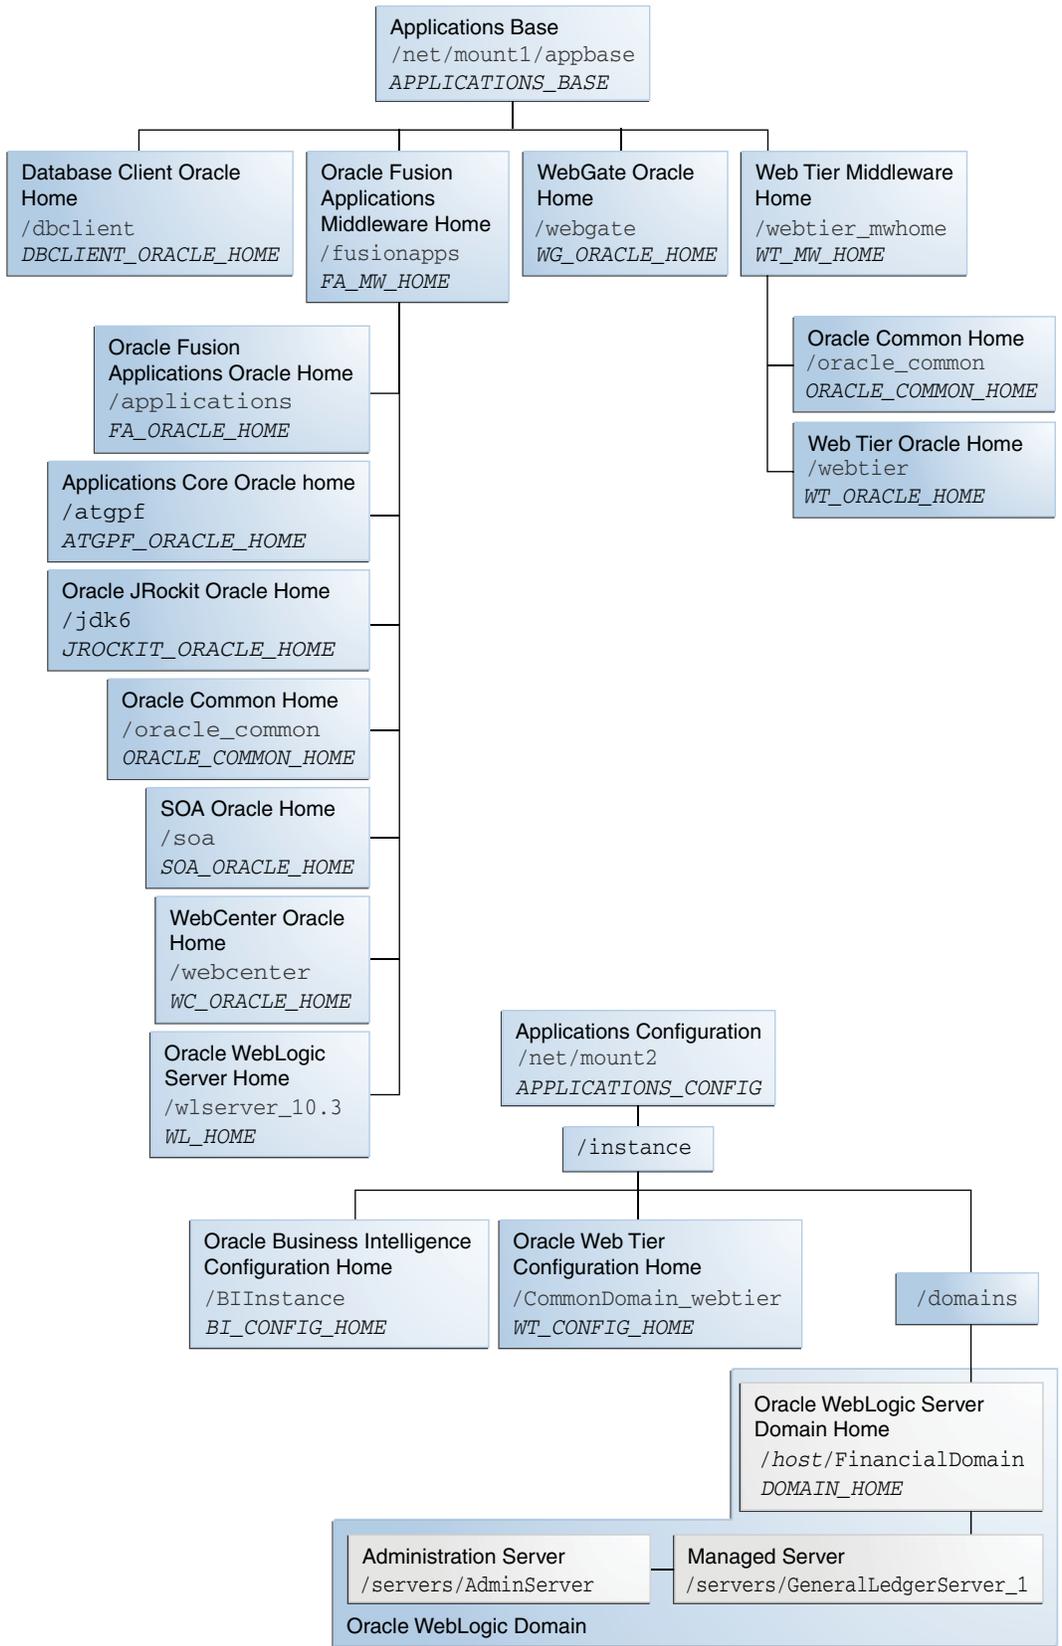
Home Directory	Description	Variable Representation
Applications base	This directory is the top-level directory for the Oracle Fusion Applications binaries. You specify the location of this top-level directory at the time of provisioning. Figure 1–6 represents this directory as <code>/net/mount1/appbase</code> for the purpose of showing an example directory structure.	<code>APPLICATIONS_BASE</code>
Applications configuration	This directory is the top-level directory for the Oracle Fusion Applications configuration files. You specify a location for this top-level directory at the time of provisioning. Figure 1–6 represents this directory as <code>/net/mount2</code> for the purpose of showing an example directory structure. Provisioning automatically creates a subdirectory named <code>instance</code> as a container for the configuration files.	<code>APPLICATIONS_CONFIG</code>
Middleware home	<p>This directory provides a container for the Oracle WebLogic Server home, an Oracle Common home, and one or more Oracle homes for components. A Middleware home can reside on a local file system or on a remote shared disk that is accessible through Network File System (NFS).</p> <p>An installation of Oracle Fusion Applications creates two Middleware homes:</p> <ul style="list-style-type: none"> fusionapps for the Oracle WebLogic Server home and Oracle homes for Oracle Fusion Middleware components, such as Oracle Business Intelligence, Oracle SOA Suite, Oracle Data Integrator, WebCenter, Oracle Enterprise Content Management, Oracle Secure Enterprise Search, and Oracle Fusion Middleware Extensions for Applications (Applications Core) webtier_mwhome for the system components, such as Oracle HTTP Server <p>For operations on specific Middleware home directories, this guide refers to each Middleware home by its product name or specifies the directory. For example, the Middleware home for Oracle Fusion Applications is referred to as the Oracle Fusion Applications Middleware home or the fusionapps Middleware home directory and its variable is <code>FA_MW_HOME</code>, while the Middleware home for Oracle Web Tier is referred to as the Web Tier Middleware home or the webtier_mwhome Middleware home directory and its variable is <code>WT_MW_HOME</code>. In cases where differentiation is not necessary, this guide refers to the Middleware home and uses the <code>MW_HOME</code> variable.</p>	<p><code>MW_HOME</code> <code>FA_MW_HOME</code> <code>WT_MW_HOME</code></p>
Oracle Business Intelligence configuration home	This directory contains the configuration files for the Oracle Business Intelligence installation.	<code>BI_CONFIG_HOME</code>
Oracle Common home	This directory contains the binary and library files required for Fusion Applications Control, test to production, and Java Required Files (JRF). The Java Required Files include the Oracle ADF libraries, a number of other components shared by the product suites included in a given installation, and components such as Oracle Metadata Service (MDS) and Oracle Platform Security Services (OPSS).	<code>ORACLE_COMMON_HOME</code>

Table 1–4 (Cont.) Home Directories

Home Directory	Description	Variable Representation
Oracle Database client home	This directory provides a container for the Oracle Database client software files.	<i>DBCLIENT_ORACLE_HOME</i>
Oracle WebLogic Server home	This directory contains installed files necessary to host Oracle WebLogic Server. The Oracle WebLogic Server home directory is a peer of Oracle home directories and resides within the directory structure of the Middleware home.	<i>WL_HOME</i>
Oracle home	<p>This directory contains installed files necessary to host a specific component or software suite. For example, the files for Oracle SOA Suite are stored in an Oracle home directory named <i>soa</i>, and the application artifacts for Oracle Fusion Applications are stored in an Oracle home directory named <i>applications</i>.</p> <p>An Oracle home resides within the directory structure of the Middleware home. There can be multiple Oracle homes within each Middleware home.</p> <p>This guide refers to each Oracle home by product name or specifies the directory. For example, the Oracle home for Oracle SOA Suite is referred to as the SOA Oracle home or the <i>soa</i> Oracle home directory and its variable is <i>SOA_ORACLE_HOME</i>, while the Oracle home for Oracle Fusion Applications is referred to as the Oracle Fusion Applications Oracle home or the <i>applications</i> Oracle home and its variable is <i>FA_ORACLE_HOME</i>.</p>	<i>ATGPF_ORACLE_HOME</i> <i>FA_ORACLE_HOME</i> <i>JROCKIT_ORACLE_HOME</i> <i>SOA_ORACLE_HOME</i> <i>WC_ORACLE_HOME</i> <i>WG_ORACLE_HOME</i> <i>WT_ORACLE_HOME</i>
Oracle Web Tier configuration home	This directory contains the configuration files for the Oracle Web Tier installation.	<i>WT_CONFIG_HOME</i>
Oracle WebLogic Server domain home	This directory contains the configuration files for the Administration Server and the Managed Servers for the product family.	<i>DOMAIN_HOME</i>

Figure 1–6 shows the relationship of the home directories using the Oracle Fusion Financials product family on a UNIX environment as an example. The Oracle Database and the Oracle Identity Management is not included, as they are installed separately from provisioning.

Figure 1-6 Relationship of Home Directories



For more information about the detailed directory structure, see the *Oracle Fusion Applications Installation Guide*.

1.3 Roadmap for Administering Oracle Fusion Applications

To gain a better understanding of where administration fits into the overall process for installing, administering, and implementing Oracle Fusion Applications, see the *Oracle Fusion Applications Administrator and Implementor Roadmap*.

Prior to using the administrative procedures in this guide, ensure the Oracle Fusion Applications environment is provisioned, as described in the *Oracle Fusion Applications Installation Guide*.

The following provides a summary of the steps to configure, manage, and perform ongoing maintenance of an Oracle Fusion Applications environment after you have installed the software:

1. Start components. See [Section 3.3](#).
2. Configure ports. See [Section 3.4](#).
3. Configure passwords. See [Section 3.5](#).
4. Manage the Oracle database. See [Section 3.6](#).
5. Configure security. See [Chapter 4](#).
6. Deploy customizations. See the *Oracle Fusion Applications Extensibility Guide*.
7. Monitor your environment. See [Chapter 9](#), [Chapter 10](#), [Chapter 11](#), and [Chapter 12](#).
8. Manage log files and diagnostic information. See [Chapter 13](#).
9. Backup your environment. See [Chapter 15](#).
10. Configure load balancing between different components or applications. See [Chapter 16](#).
11. Configure Oracle Transactional Business Intelligence to perform real-time analysis of the Oracle Fusion applications. See the *Oracle Fusion Administrator's Guide for Transactional Business Intelligence*.
12. After the Oracle Fusion applications are running, configure the Oracle Business Intelligence Applications applications for analyzing historical data. See the *Oracle Fusion Middleware Installation and Configuration Guide for Oracle Business Intelligence Applications*.

In addition to this guide, refer to the following guides for tasks that you may need to perform, depending on your business needs:

- *Oracle Fusion Applications Common Implementation Guide* for common setup and maintenance tasks for Oracle Fusion applications
- *Oracle Fusion Applications Post-Installation Guide* for configuration tasks to perform after installation
- *Oracle Fusion Applications Administrator's Troubleshooting Guide* for common issues in the runtime environment

Part II

Basic Administration

This part describes basic administration tasks.

Part II contains the following chapters:

- Chapter 2, "Getting Started with Administering Oracle Fusion Applications"
- Chapter 3, "Performing Routine Administrative Tasks"
- Chapter 4, "Securing Oracle Fusion Applications"
- Chapter 5, "Managing Oracle Enterprise Scheduler Service and Jobs"
- Chapter 6, "Managing Report Delivery Servers"
- Chapter 7, "Managing Search with Oracle Enterprise Crawl and Search Framework"
- Chapter 8, "Managing Oracle WebLogic Communication Services for Click-to-Dial Functionality"

Getting Started with Administering Oracle Fusion Applications

This chapter provides information you need to get started administering your Oracle Fusion Applications environment, including information about the tools you use.

Prior to using the administrative procedures in this chapter and the rest of this guide, ensure the Oracle Fusion Applications environment is provisioned, as described in the *Oracle Fusion Applications Installation Guide*.

This chapter contains the following topics:

- [Introduction to the Tools Used to Administer Oracle Fusion Applications](#)
- [Key Differences Between Fusion Applications Control and Cloud Control](#)
- [When to Use Fusion Applications Control Tool, When to Use Oracle WebLogic Server Administration Console](#)
- [Finding the Administration Server Port and Other Administrative URLs for Domains](#)
- [Using Fusion Applications Control](#)
- [Using Oracle WebLogic Server Administration Console](#)
- [Using Cloud Control](#)
- [Using Oracle Fusion Functional Setup Manager](#)
- [Using Database Tools to Modify Oracle Fusion Applications Data](#)

2.1 Introduction to the Tools Used to Administer Oracle Fusion Applications

Oracle provides several tools to administer the Oracle Fusion Applications, the Oracle Fusion Middleware layer, and the Oracle Database within your Oracle Fusion Applications installations. Use these tools, rather than directly editing configuration files, to perform all administrative tasks unless a specific procedure requires you to edit a file. Editing a file may cause the settings to be inconsistent and generate problems.

The primary administration tools for Oracle Fusion Applications are described in [Table 2-1](#).

Table 2–1 Administration Tools for Oracle Fusion Applications

Tool	Description
Oracle Enterprise Manager Fusion Applications Control (Fusion Applications Control)	Enables you to monitor and administer a product family within the Oracle Fusion Applications environment. See Section 2.5 .
Oracle Enterprise Manager Cloud Control (Cloud Control)	Enables you to monitor and administer the entire Oracle Fusion Applications environment, including product families, Oracle Fusion Middleware components across all Oracle WebLogic Servers, and the Oracle Database. See Section 2.2 to understand the key difference between Cloud Control and Fusion Applications Control, and see Section 2.7 to get started with using Cloud Control.
Oracle Fusion Functional Setup Manager	Guides you through the tasks of setting up the functional data and reference data. See Section 2.5 .
Oracle Fusion Applications Diagnostic Dashboard	Enables you to register, execute, and monitor diagnostic tests for Oracle Fusion Applications, and also work with diagnostic test registration tags and purge diagnostic test results. See Chapter 13 .

The primary administration tools for Oracle Fusion Middleware are described in [Table 2–2](#).

Table 2–2 Administration Tools for Oracle Fusion Middleware

Tool	Description
Fusion Applications Control	Enables you to monitor and administer Oracle Fusion Middleware components within an Oracle WebLogic Server domain for a specific product family within the Oracle Fusion Applications environment. See Section 2.5 . Fusion Applications Control also includes a System MBean Browser that enables you browse the MBeans for an Oracle WebLogic Server or for a selected application. See "Getting Started Using the Fusion Middleware Control MBean Browsers" in the <i>Oracle Fusion Middleware Administrator's Guide</i> .
Oracle Enterprise Manager Cloud Control (Cloud Control)	Enables you to monitor and administer the entire Oracle Fusion Applications environment, including product families, Oracle Fusion Middleware components across all Oracle WebLogic Servers, and the Oracle Database. See Section 2.2 to understand the key difference between Cloud Control and Fusion Applications Control, and see Section 2.7 to get started with using Cloud Control.
Oracle WebLogic Server Administration Console	Enables you to manage the Oracle WebLogic Server in each domain. See Section 2.3 to understand the when to use Administration Console, and see Section 2.7 to get started with using the Administration Console.
Oracle Fusion Middleware command-line tools	Enables you to manage Oracle Fusion Middleware components. See "Getting Started Using Command Line Tools" in the <i>Oracle Fusion Middleware Administrator's Guide</i> to learn about the Oracle WebLogic Scripting Tool (WLST) and the Oracle Process Manager and Notification Server (OPMN).

Table 2–2 (Cont.) Administration Tools for Oracle Fusion Middleware

Tool	Description
Middleware support for application security	<p>The Oracle Fusion Applications environment includes the following tools for managing application security:</p> <ul style="list-style-type: none"> ■ Oracle WebLogic Server Administration Console enables you to manage policies of enterprise applications. See the <i>Oracle Fusion Middleware Securing Resources Using Roles and Policies for Oracle WebLogic Server</i>. ■ Oracle Authorization Policy Manager enables you to manage application policies. See the <i>Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)</i>. ■ Oracle Access Manager Administration Console enables you to manage system configurations, including Oracle Access Manager server instances and data sources, access policies, agent profiles, and other security elements. See the <i>Oracle Fusion Middleware Administrator's Guide for Oracle Access Manager with Oracle Security Token Service</i>. ■ Oracle Identity Manager Administrative and User Console enables you to access the organization's resources (through self-service) and enable administrators to manage the access (by creating users and roles, and defining authorization policies). See the <i>Oracle Fusion Middleware Administrator's Guide for Oracle Directory Integration Platform</i>. ■ Oracle Web Service Policy Manager enables you manage and secure web services for SOA composite applications. See the <i>Oracle Fusion Middleware Security and Administrator's Guide for Web Services</i>. <p>For an overview and information about Fusion security as it is relevant specifically to Oracle Fusion Applications, including details about how to interact with the Fusion Security infrastructure to implement and administer security for your enterprise, see the <i>Oracle Fusion Applications Security Guide</i>.</p>

The primary administration tools for the Oracle Database are described in [Table 2–3](#).

Table 2–3 Administration Tools for the Oracle Database

Tool	Description
Oracle Enterprise Manager Database Control	Enables you to monitor and manage the Oracle Database for the Oracle Fusion Applications. See the <i>Oracle Database 2 Day DBA</i> .
Cloud Control	Enables you to monitor historical data for the Oracle Database. See Section 2.7 .

2.2 Key Differences Between Fusion Applications Control and Cloud Control

You can manage your Oracle Fusion Applications environment using Fusion Applications Control or Cloud Control.

Fusion Applications Control organizes a wide variety of performance data and administrative functions into distinct, Web-based home pages for the farm, cluster, domain, servers, components, and applications. The Fusion Applications Control home pages make it easy to locate the most important monitoring data and the most commonly used administrative functions all from your Web browser. Fusion Applications Control is automatically installed with each Oracle WebLogic Server

domain for each product family during an Oracle Fusion Applications installation. For more information about this control, see [Section 2.5](#).

Managing several farms from different, distinct Fusion Applications Controls becomes very difficult in an Oracle Fusion Applications environment. For centralized management and additional management functionality (such as historical data collections for performance trending, email and page notifications for potential availability or performance problems, and service level management), you can use Cloud Control.

Cloud Control enables you to monitor all the product families, Oracle Fusion Middleware components, and the Oracle Database across the entire Oracle Fusion Applications environment from one console. To use Cloud Control, you need to separately install Oracle Enterprise Manager Cloud Control 12c (12.1.0.1) with the Application Management Pack for Oracle Fusion Applications. For more information about this control, see [Section 2.7](#). For more information about installing Cloud Control, see the following documentation:

- *Oracle Enterprise Manager Cloud Control Basic Installation Guide* for more information about the installation process
- *Oracle Enterprise Manager Licensing Information* for more information about the Application Management Pack

2.3 When to Use Fusion Applications Control Tool, When to Use Oracle WebLogic Server Administration Console

Fusion Applications Control and Oracle WebLogic Server Administration Console are graphical user interfaces that you can use to monitor and administer your environment. Use [Table 2-4](#) to locate the task you are performing.

Table 2-4 Comparing Fusion Applications Control and WebLogic Server Administration Console

Task	Tool to Use
Manage Oracle WebLogic Server	Use:
Create additional Managed Servers	WebLogic Server Administration Console
Clone Managed Servers	WebLogic Server Administration Console
Cluster Managed Servers	WebLogic Server Administration Console
Start and stop Oracle WebLogic Server	Fusion Applications Control or WebLogic Server Administration Console
Add users and groups	WebLogic Server Administration Console if using the default embedded LDAP; if using another LDAP server, use the LDAP server's tool
Manage Data Sources	Use:
Create data sources	WebLogic Server Administration Console
Create connection pools	WebLogic Server Administration Console
Manage JMS Resources	Use:
Create JMS queues	WebLogic Server Administration Console
Configure advanced queuing	WebLogic Server Administration Console
Manage SOA environment	Use:

Table 2–4 (Cont.) Comparing Fusion Applications Control and WebLogic Server Administration Console

Task	Tool to Use
Deploy and undeploy SOA Composite applications	Fusion Applications Control
Monitor SOA Composite applications	Fusion Applications Control
Modify Oracle BPEL Process Manager MBean properties	Fusion Applications Control
Debug applications such as Oracle BPEL Process Manager applications	Fusion Applications Control
ADF Applications	Use:
Deploy and undeploy ADF applications	Fusion Applications Control
Java EE applications	Use:
Deploy and undeploy Java EE applications	WebLogic Server Administration Console or Fusion Applications Control
Security	Use:
Configure and manage auditing	Fusion Applications Control
Configure SSL	WebLogic Server Administration Console for Oracle WebLogic Server Fusion Applications Control for Java components and system components. See Chapter 4 .
Change passwords	WebLogic Server Administration Console
Manage Components	Use:
View and manage log files	Fusion Applications Control for most log files. See Section 13.6 . WebLogic Server Administration Console for the following logs: <i>DOMAIN_HOME/servers/server_name/logs/access.log</i> <i>DOMAIN_HOME/servers/server_name/data/ldap/log/EmbeddedLDAP.log</i> <i>DOMAIN_HOME/servers/server_name/data/ldap/log/EmbeddedLDAPAccess.log</i>
Change ports	WebLogic Server Administration Console for Oracle WebLogic Server and Java components For some system components, Fusion Applications Control.
Manage Oracle HTTP Server	Fusion Applications Control
Start and stop components	Fusion Applications Control
Start and stop applications	Fusion Applications Control

2.4 Finding the Administration Server Port and Other Administrative URLs for Domains

To locate the port for the Administration Server and other administrative URLs for the Oracle WebLogic Server domains, use the summary file saved during provisioning.

During provisioning of Oracle Fusion Applications, on the Summary page of the Provisioning Wizard, when you clicked **Finish**, you saved the plan and the summary in a text document that summarizes the details of this plan. For more information about using the Provisioning Wizard, see "Creating a Provisioning Plan" section in the *Oracle Fusion Applications Installation Guide*.

In the summary file, under the Admin Server entry for the domain, locate the following subentries:

- Managed Server Port for the Administration Server port
- Enterprise Manager Welcome Page for the URL for Fusion Applications Control
- Admin Server for the URL for the Oracle WebLogic Server Administration Console

[Example 2-1](#) shows an excerpt of the summary file.

Example 2-1 Summary File Excerpt

```

...
Common Domain
  Admin Server
    Host: hostname.domain
    Managed Server Port: 7001
    Secure Managed Server Port: 7002
    Admin Console
      http://commoninternal.domain:7777/console
    Enterprise Manager Welcome Page
      http://commoninternal.domain:7777/em
  Functional Setup Manager
    Host: hostname.domain
    Managed Server Port: 7004
    Secure Managed Server Port: 7005
  Home Page
    Host: hostname.domain
    Managed Server Port: 7006
    Secure Managed Server Port: 7007
  Home Page
    Host: hostname.domain
    Managed Server Port: 7006
    Secure Managed Server Port: 7007
  Help Portal
    Host: hostname.domain
    Managed Server Port: 7008
    Secure Managed Server Port: 7009
  Business Process Server
    Host: hostname.domain
    Managed Server Port: 7010
    Secure Managed Server Port: 7011
  Content Server
    Host: hostname.domain
    Managed Server Port: 7012
    Secure Managed Server Port: 7013
  Secure Enterprise Search
    Host: hostname.domain
    Managed Server Port: 7016
    Secure Managed Server Port: 7017
  Enterprise Scheduler
    Host: hostname.domain
    Managed Server Port: 7020
    Secure Managed Server Port: 7021
  WebCenter Spaces

```

```

Host: hostname.domain
Managed Server Port: 7022
Secure Managed Server Port: 7023
WebCenter Collaboration
Host: hostname.domain
Managed Server Port: 7024
Secure Managed Server Port: 7025
OWLCS Server
Host: hostname.domain
Managed Server Port: 7030
Secure Managed Server Port: 7031
OWLCS SIP State Server
Host: hostname.domain
Managed Server Port: 7032
Secure Managed Server Port: 7033
Web Tier
Host: hostname.domain
CRM Domain
Admin Server
Host: hostname.domain
Managed Server Port: 9001
Secure Managed Server Port: 9002
Admin Console
  http://crminternal.domain:7777/console
Enterprise Manager Welcome Page
  http://crminternal.domain:7777/em
...
Financial Domain
Admin Server
Host: hostname.domain
Managed Server Port: 7401
Secure Managed Server Port: 7401
Admin Console
  http://fininternal.domain:7777/console
Enterprise Manager Welcome Page
  http://fininternal.domain:7777/em
...

```

2.5 Using Fusion Applications Control

Fusion Applications Control is a Web browser-based, graphical user interface that you can use to monitor a product family and its products. It also enables you to monitor and administer Oracle Fusion Applications and an Oracle Fusion Middleware farm.

A **farm** is a collection of components managed by Fusion Applications Control. It can contain an Oracle WebLogic Server domain, one Administration Server, one or more Managed Servers, clusters, and the Oracle Fusion Middleware components that are installed, configured, and running in the domain.

Fusion Applications Control organizes a wide variety of performance data and administrative functions into distinct, Web-based home pages for the product family, products, Oracle Fusion Applications, farm, domain, servers, and Oracle Fusion Middleware components. The Fusion Applications Control home pages make it easy to locate the most important monitoring data and the most commonly used administrative functions for various targets—all from your Web browser.

The section contains the following topics:

- [Understanding How Fusion Applications Control Compares with Fusion Middleware Control](#)

- [Starting Fusion Applications Control](#)
- [Using Fusion Applications Control Help](#)
- [Navigating within Fusion Applications Control](#)
- [Understanding the Targets in the Target Navigation Pane](#)
- [Navigating to the Product Family and Product Administration Pages](#)
- [Viewing the Performance of Targets](#)
- [Understanding Users and Roles for Fusion Applications Control](#)

For information about the System MBean Browser, see the "Getting Started Using the Fusion Middleware Control MBean Browsers" section in the *Oracle Fusion Middleware Administrator's Guide*.

2.5.1 Understanding How Fusion Applications Control Compares with Fusion Middleware Control

If you have worked with Oracle Fusion Middleware, then you may already be familiar with Fusion Middleware Control. Fusion Middleware Control enables you to monitor and administer a farm. Fusion Applications Control provides all the functionality available in Fusion Middleware Control plus functionality specific to Oracle Fusion Applications.

Many procedures throughout this guide reference content in Oracle Fusion Middleware guides, which describes using Fusion Middleware Control. Where these references occur, the same procedures also apply to Fusion Applications Control. Where the controls vary, the procedures for Fusion Applications Control are contained within this guide.

2.5.2 Starting Fusion Applications Control

To display Fusion Applications Control for a given Oracle WebLogic Server domain:

1. Enter the URL in your Web browser:

```
http://product_familyinternal.domain:port/em
```

where *product_familyinternal.domain*. *port* is the host and domain of the Oracle HTTP Server or a Load Balancer.

During provisioning of Oracle Fusion Applications, on the Summary page of the Provisioning Wizard, when you clicked **Finish**, you saved the plan and the summary in a text document that summarizes the details of this plan. For more information about using the Provisioning Wizard, see "Creating a Provisioning Plan" in the *Oracle Fusion Applications Installation Guide*.

In the summary file, for the domain, find the Enterprise Manager Welcome Page entry to locate the URL. For example:

```
....
CommonDomain
    Enterprise Manager Welcome Page
        http://commoninternal.domain:7777/em
....
CRM Domain
    Enterprise Manager Welcome Page
        http://crminternal.domain:7777/em
....
```

```

Financial Domain
  Enterprise Manager Welcome Page
    http://fininternal.domain:7777/em
  ...

```

The login page displays.



2. Enter Oracle Fusion Middleware administrative user name and password, and then click **Login**.

In an Oracle Fusion Applications installation, this user is the super user specified on the Identity and Policy Management page of the Provisioning Wizard. See "Identity and Policy Management Configuration" section in the *Oracle Fusion Applications Installation Guide* for information about using the Identity and Policy Management Configuration page and [Section 4.8.5.2](#) to change the password.

2.5.3 Using Fusion Applications Control Help

At any time while using the Fusion Applications Control, you can click **Help > Enterprise Manager** at the top of the page to get more information. In most cases, the Help window displays a help topic about the current page. Click **Contents** in the Help window to browse the list of help topics, or click **Search** to search for a particular word or phrase.

2.5.4 Navigating within Fusion Applications Control

Fusion Applications Control displays the target navigation pane on the left and the content pane on the right. For example, when you first log in to Fusion Applications Control, the product family home page is displayed on the right.

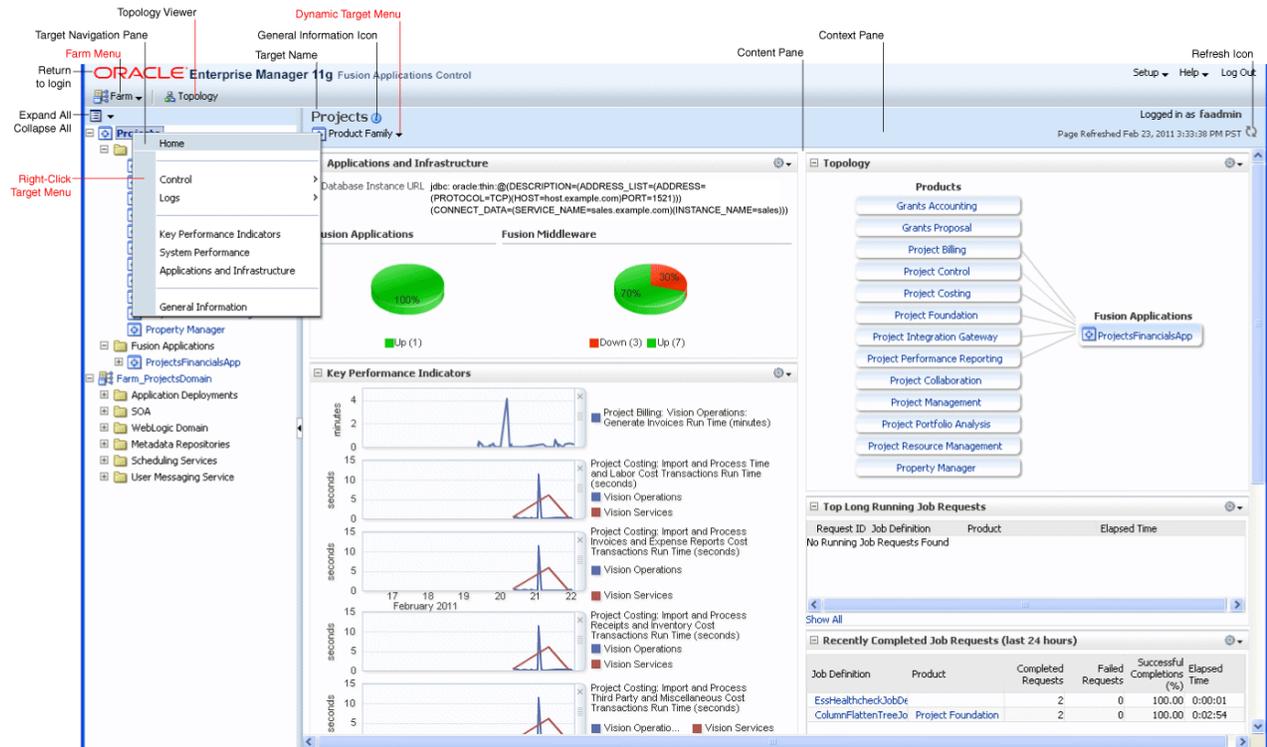
From the target navigation pane, you can expand the tree and select either targets for the product family or the farm. See [Section 2.5.5](#) for more information about the target navigation pane.

When you select a target, the target home page displays in the content pane and that target's menu is displayed at the top of the page, in the context pane. For example, if

you select a product family, the product family menu is displayed. You can also view the menu for a target by right-clicking the target in the navigation pane.

Figure 2–1 shows the target navigation pane and the home page of a product family. Because a product family was selected, the dynamic target menu listed in the context pane is the product family menu.

Figure 2–1 Overview of Fusion Applications Control



In the preceding figure, the following items are called out:

- **Target Navigation Pane** lists all of the targets for an Oracle Fusion Applications product family and the farm navigation tree.
- **Content Pane** shows the current page for the target. When you first select a target, that target's home page is displayed.
- **Dynamic Target Menu** provides a list of operations that you can perform on the currently selected target. The menu that is displayed depends on the target you select. The menu for a specific target contains the same operations as those in the **Right-Click Target Menu**.
- **Right-Click Target Menu** provides a list of operations that you can perform on the currently selected target. The menu is displayed when you right-click the target name in the target navigation pane. In the figure, the user has right-clicked the product family.

The menu for a specific target contains the same operations as those in the **Dynamic Target Menu**.

- **Topology Viewer** displays the topology of the product family, including the middleware relationships and the products and applications in the product family.

- **Target Name** is the name of the currently selected target.
- **General Information Icon** provides information about the target. For example, for a domain, it displays the target name, the version, and the domain home.
- **Context Pane** provides the name of the target, the name of the current user, the host name, and the time of the last page refresh, as well as the Refresh icon.
- **Expand All/Collapse All** lets you expand or collapse the navigation tree.
- **Refresh** indicates when the page is being refreshed. Click it to refresh a page with new data. (Refreshing the browser window refreshes the page but does not retrieve new data.)
- **Go To Home** takes you to the product family home page, when you click the Oracle logo.

Table 2–5 describes some common ways you can navigate within Fusion Applications Control.

Table 2–5 Navigating within Fusion Applications Control

To:	Take This Action:
View all of the targets in the farm	Click the Expand All icon at the top of the target navigation pane .
Navigate to a target	Select the target in the target navigation pane . Selecting the product family from the top of the displays the product family home page and selecting the farm displays the farm home page. You can expand the product family and the farm for additional targets.
Operate on a target	Right-click the target in the target navigation pane . The target menu display. Alternatively, you can select the target and use the dynamic target menu in the context pane.
Return to the target's home page	Click the target name at the top left-hand corner of the context pane .
Refresh a page with new data	Click the Refresh icon in the top right of the context pane .
Return to a previous page	Click the breadcrumbs, which appear below the context pane. The breadcrumbs appear when you drill down in a target. For example, from the WebLogic Server menu, choose Logs > View Log Messages . Select a log file and click View Log File . The breadcrumbs show: Log Messages > Log Files > View Log File: logfile_name
View the host on which the target is running	Select the target in the target navigation pane and view the host name in the target's context pane . You can also view the host name by clicking the General Information icon.
Return to the product family home page	Click the Oracle Enterprise Manager logo at the top left of the page.
View the topology	Click Topology .
View a server log file	Right-click the server name in the target navigation pane . Choose Logs > View Log Messages to see a summary of log messages and to search log files.

In addition, from Fusion Applications Control, from the home pages of targets such as the Administration Server or Managed Servers, you can access the WebLogic Server

Administration Console. For information about configuring single sign-on between Fusion Middleware Control and the WebLogic Server Administration Console, see "Configuring Single Sign-On for Administration Consoles" in the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Identity Management (Oracle Fusion Applications Edition)*.

2.5.5 Understanding the Targets in the Target Navigation Pane

Figure 2–2 shows the target navigation pane. The product family tree enables you to view an Oracle Fusion Applications product or application target. The farm tree enables you to select an Oracle WebLogic Server domain, an Oracle WebLogic Server Managed Server, an Oracle Fusion Middleware component, an application, or an Oracle Metadata Repository.

Figure 2–2 Target Navigation Pane



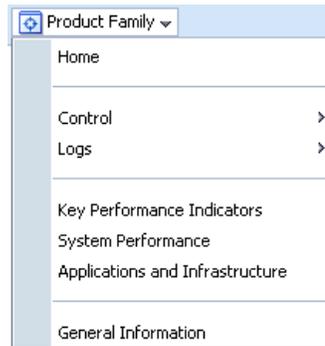
2.5.6 Navigating to the Product Family and Product Administration Pages

To navigate to the product family and product administration pages in Fusion Applications Control:

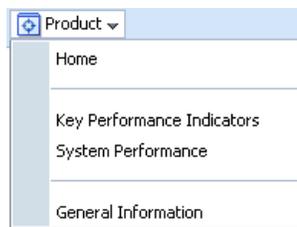
1. Navigate to the product family and product home page:

- To access the product family home page, from the navigation pane, select the name of the product family.
 - To access the product home page, from the navigation pane, expand the product family, then **Products**, and then select the product.
2. Select the **Product Family** or **Product** menu.

The following figure shows the **Product Family** menu:



The following figure shows the **Product** menu:



The **Product Family** and **Product** menu displays the options described in [Table 2-6](#).

Table 2-6 Product Family and Product Target Menu Options

Menu Option	Description
Home	Displays the product family or product home page. For more information about this page, see Section 9.2 and Section 9.3 .
Control	This option provides options for starting and stopping all the cluster applications in the product family. These options do not start and stop the Managed Servers, which is recommended when starting and stopping all applications in a product family. Instead of these options, use one of the start and stop use cases described in Section 3.3 .
Logs	Displays the options to either view log messages or configure log messages for all the applications in the product family. For more information, see Section 13.6 .
Key Performance Indicators	Displays the Key Performance Indicators page for viewing the middle-tier metrics for the product family or product. see Section 2.5.7 , Section 9.9 , and Section 9.10 .
System Performance	Displays the System Performance page for viewing the middle-tier metrics for the product family, see Section 2.5.7 and Section 10.3.1 .

Table 2–6 (Cont.) Product Family and Product Target Menu Options

Menu Option	Description
Applications and Infrastructure	Displays the Applications and Infrastructure page to view the status of Oracle Fusion Applications and Oracle Fusion Middleware components for the product family. See Section 9.8 for more information about using this page for monitoring applications, and see Section 10.2.1 for more information about using this page for monitoring Oracle Fusion Middleware.
General Information	Displays general details about the target.

2.5.7 Viewing the Performance of Targets

Fusion Applications Control automatically and continuously measures run-time performance. The performance metrics are automatically enabled; you do not need to set options or perform any extra configuration to collect them.

If you encounter a problem, such as an application that is running slowly or is hanging, you can view more detailed performance information, including performance metrics for a particular target, to find out more information about the problem. In many cases, these metrics are shown in interactive charts; other times they are presented in tabular format. The best way to use and correlate the performance metrics is from the Performance Summary page for the component or application you are monitoring.

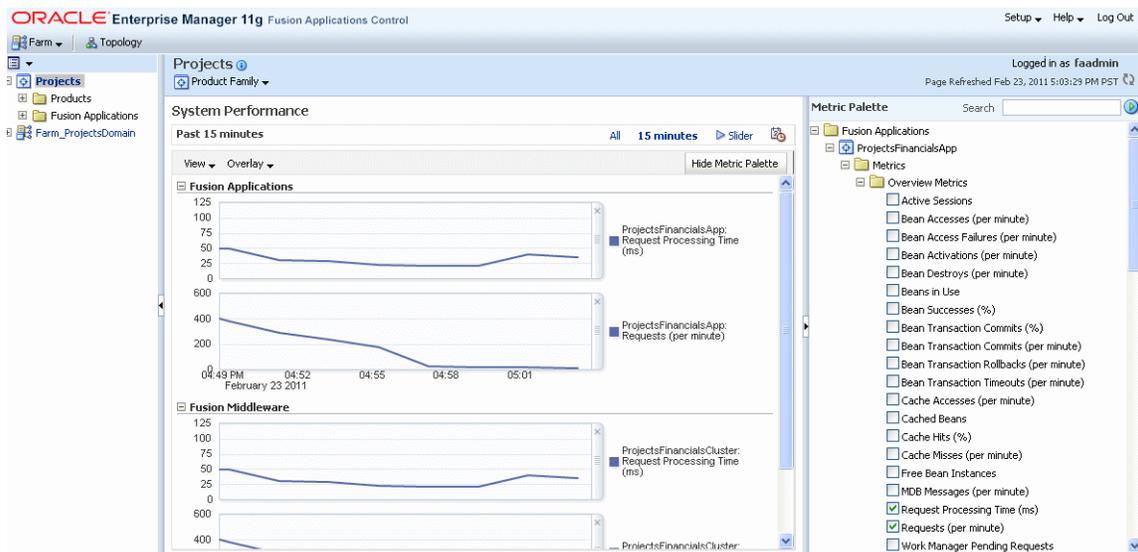
To view the performance of a target with Fusion Applications Control:

1. From the menu for a specific target, choose **Performance Summary**. If the target is a product family or product, choose **Key Performance Indicators** to view business metrics or **System Performance** to view middle tier metrics.

The Performance Summary page displays.

2. To see additional metrics, click **Show Metric Palette** and expand the metric categories.

The following figure shows the System Performance page for a product family with the Metric Palette displayed:



3. Select additional metrics to add them to the Performance Summary.

4. To re-organize the display of performance charts in the System Performance page, select **View > Reorder Chart Sets** and move charts up or down, or remove a particular chart from the page.
5. To overlay another target, click **Overlay**, and select the target. The target is added to the charts, so that you can view the performance of more than one target at a time, comparing their performance.

Alternatively, you can display historical business metrics of the same product instance. Select **Overlay > Today with Yesterday, This Week with Last Week or This Month with Last Month**.

6. To customize the time frame shown by the charts, you can:
 - Click **Slider** to display a slider tool that lets you specify that more or less time is shown in the charts. For example, to show the past 10 minutes, instead of the past 15 minutes, slide the left slider control to the right until it displays the last 10 minutes.
 - Click the **Enter Time** clock icon. Then, enter the **Start Time** and **End Time**, and click **OK**. If there is no data available from the earlier time point, a confirmation message displays, explaining the timeline will be automatically adjusted to the time period for which the data is available.
7. To save the chart set so as to preserve a specific set of metrics selected for display, as opposed to the default metrics, select **View > Save Chart Set** and enter a name for the chart set to be saved.

Use the Fusion Applications Control online help to obtain a definition of a specific performance metric. There are two ways to access this information:

- Browse or search for the metric in the Fusion Applications Control online help.
- Navigate to the Performance Summary page for your component or application and perform the following:
 - a. Click **Show Metric Palette**.
 - b. Browse the list of metrics available for the component or application to locate a specific metric.
 - c. Right-click the name of the metric and select **Help** from the context menu.

2.5.8 Understanding Users and Roles for Fusion Applications Control

To access Fusion Applications Control and perform tasks, you must have the appropriate role. Fusion Applications Control uses the Oracle WebLogic Server security realm and the roles defined in that realm. If a user is not granted one of these roles, the user cannot access Fusion Applications Control.

Each role defines the type of access a user has. For example, a user with the role Admin has full privileges. A user with the role Operator has privileges to perform essential day-to-day operations. A user with the role Monitor has privileges only to view the configuration.

For more information, see "Users, Groups, and Security Roles" in the *Oracle Fusion Middleware Securing Resources Using Roles and Policies for Oracle WebLogic Server*.

2.6 Using Oracle WebLogic Server Administration Console

Oracle WebLogic Server Administration Console is a Web browser-based, graphical user interface that you use to manage an Oracle WebLogic Server domain.

While Fusion Applications Control is the primary administrative interface, use Administration Console as specified by [Table 2-4](#) and wherever recommended in procedures in this guide.

The section contains the following topics:

- [Starting Oracle WebLogic Server Administration Console](#)
- [Locking the WebLogic Server Configuration](#)

2.6.1 Starting Oracle WebLogic Server Administration Console

To display Oracle WebLogic Server Administration Console for a given Oracle WebLogic Server domain:

1. Enter the URL in your Web browser:

```
http://product_familyinternal.domain:port/em
```

where *product_familyinternal.domain*. *port* is the host and domain of the Oracle HTTP Server or a Load Balancer.

During provisioning of Oracle Fusion Applications, on the Summary page of the Provisioning Wizard, when you clicked **Finish**, you saved the plan and the summary in a text document that summarizes the details of this plan. For more information about using the Provisioning Wizard, see "Creating a Provisioning Plan" in the *Oracle Fusion Applications Installation Guide*.

In the summary file, for the domain, find the Enterprise Manager Welcome Page entry to locate the URL. For example:

```
....
CommonDomain
  Admin Console
    http://commoninternal.domain:7777/console
....
CRM Domain
  Admin Console
    http://crminternal.domain:7777/console
....
Financial Domain
  Admin Console
    http://fininternal.domain:7777/console
...
```

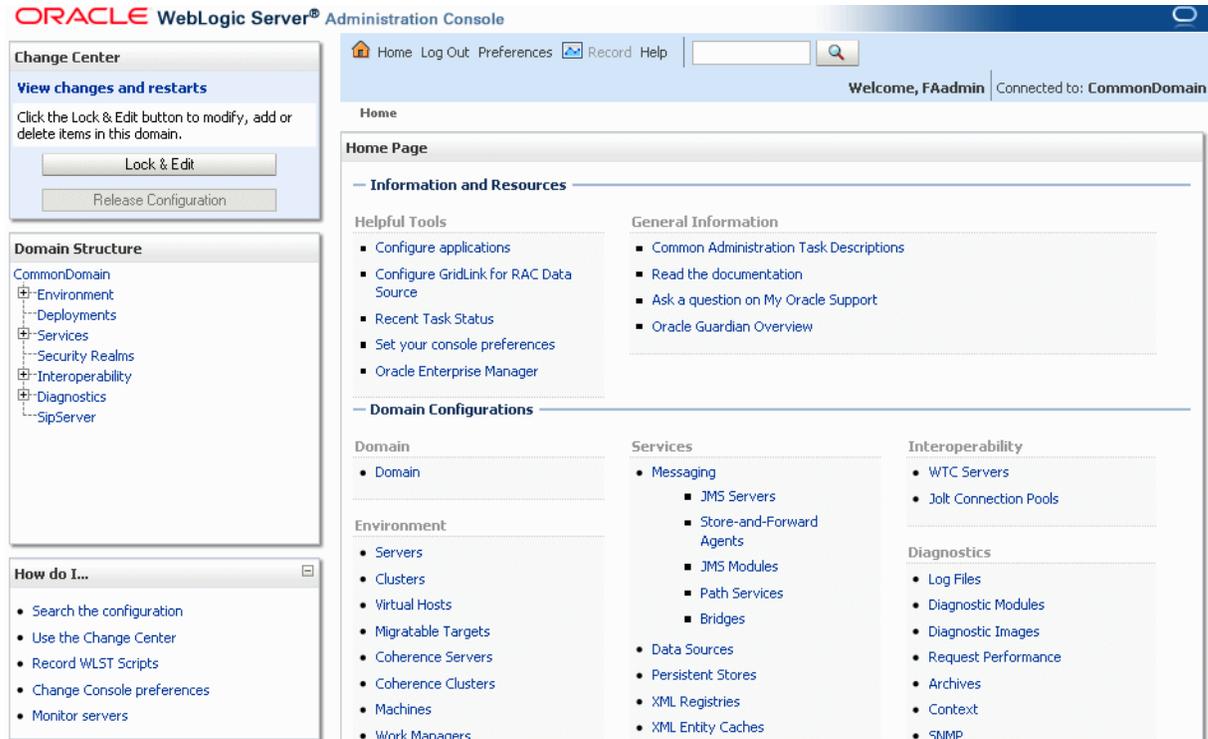
The login page displays.



2. Enter Oracle Fusion Middleware administrative user name and password, and then click **Login**.

In an Oracle Fusion Applications installation, this user is the super user specified on the Identity and Policy Management page of the Provisioning Wizard. See "Identity and Policy Management Configuration" section in the *Oracle Fusion Applications Installation Guide* for information about using the Identity and Policy Management Configuration page and [Section 4.8.5.2](#) to change the password.

The home page displays, as shown in the following figure:



Alternatively, you can access the Administration Console from Fusion Applications Control, from the home pages of targets such as the Administration Server or Managed Servers.

2.6.2 Locking the WebLogic Server Configuration

Before you make configuration changes, lock the domain configuration, so you can make changes to the configuration while preventing other accounts from making changes during your edit session. To lock the domain configuration:

1. Locate the Change Center in the upper left of the Administration Console screen.
2. Click Lock & Edit to lock the configuration edit hierarchy for the domain.

As you make configuration changes using the Administration Console, you click **Save** (or in some cases **Finish**) on the appropriate pages. This does not cause the changes to take effect immediately. The changes take effect when you click **Activate Changes** in the Change Center. At that point, the configuration changes are distributed to each of the servers in the domain. If the changes are acceptable to each of the servers, then they take effect. If any server cannot accept a change, then all of the changes are rolled back from all of the servers in the domain. The changes are left in a pending state; you can then either edit the pending changes to resolve the problem or revert to the previous configuration.

2.7 Using Cloud Control

Cloud Control enables you to monitor runtime and historical data for the Oracle Fusion applications, Oracle Fusion Middleware components, and Oracle Database across the entire environment.

For an overview of Cloud Control, see the "Overview of Enterprise Manager Cloud Control 12c" chapter in *Oracle Enterprise Manager Concepts*.

Application Management Pack for Oracle Fusion Application is required to be licensed with Oracle Fusion Applications, and must match the number of licenses for the associated Oracle WebLogic Suite and Oracle WebLogic Suite for Oracle Fusion Applications and Oracle SOA Suite for Oracle Middleware or Oracle SOA Suite for Oracle Middleware for Oracle Fusion Applications. For more information about other Oracle Enterprise Manager products that can be licensed with Oracle Fusion Applications, see the *Oracle Enterprise Manager Licensing Information*.

The section contains the following topics:

- [Starting Cloud Control and Discovering Oracle Fusion Applications Targets](#)
- [Rediscovering Oracle Fusion Applications Targets](#)
- [Removing Oracle Fusion Applications Targets](#)
- [Navigating the Fusion Applications Target Home Page in Cloud Control](#)
- [Using Cloud Control Help](#)

For specific uses of Cloud Control with an Oracle Fusion Applications environment, see the following sections:

- [Section 10.5, "Performing Proactive Performance Monitoring with Cloud Control"](#)
- [Chapter 11, "Diagnosing the Oracle Fusion Applications Middle Tier Performance"](#)

2.7.1 Starting Cloud Control and Discovering Oracle Fusion Applications Targets

To use Cloud Control, you need to separately install Oracle Enterprise Manager Cloud Control 12c (12.1.0.2) with the Application Management Pack for Oracle Fusion Applications. For more information about installing Cloud Control, see the *Oracle Enterprise Manager Cloud Control Basic Installation Guide*.

To display Cloud Control:

1. Enter the URL in your Web browser:

`http://oms_hostname.domain:port/em`

where `oms_hostname` and `oms_hostname` are the host and port number of the Oracle Management Service.

The login page displays.



2. Enter the default super administrator account, `SYSMAN`, that is created with the password you specified during the installation, and then click **Login**.

The home page displays. This page provides a high-level view of your entire enterprise. For more information about this page, see [Section 2.7.4](#) and click the **Help** link at the top of the page.

3. Enable licenses according to your licensing agreement with Oracle.

For more information, access the Enterprise Manager documentation library at the following URL and view the *Oracle Enterprise Manager Licensing Information Guide*: <http://www.oracle.com/technetwork/indexes/documentation/index.html>

4. To enable business metrics, perform the following prior to performing discovery:

- a. Install the agent on the host where the specified Administration Server is running.
- b. Give the agent read access to all the files in the directory that contains the `jps-config-jse.xml` file. The `jps-config-jse.xml` file is located in one of the following directories:

(UNIX) `DOMAIN_HOME/config/fmwconfig`
 (Windows) `DOMAIN_HOME\config\fmwconfig`

The agent must be installed a host that has NFS access to the domain home directory using the path as would be used on the Administration Server host. In addition, the agent user must have read access to this file.

5. Discover the Oracle Fusion Applications environment:

- a. From the **Targets** menu, choose **Fusion Applications**.
- b. Click **Add Fusion Instance**.

The Add Fusion Instance: Functional Setup wizard displays.

- c. Enter the host name for the Administration Server for the `CommonDomain` is in the **Administration Server Host** field.
- d. Enter the Administration Server port in the **Port** field.

- e. In the **Username** and **Password** fields, enter the username and password for the Administration Server.
- f. In the **Unique Fusion Instance Identifier**, enter an intuitive name for the Fusion Instance. This name will be used to identify the Fusion Instance throughout Cloud Control. It must be unique and can include alphanumeric as well as dash (-), underscore (_) or period (.) characters.
- g. If the Fusion Instance is deployed within a secured Functional Setup Domain, click **Secured**.
- h. Enter the host name and port for the Management Agent that will enable monitoring of the Fusion Instance targets. To search for an agent, click the icon next to the **Agent** field. Local Agents deployed on the WebLogic Admin Server host will be displayed first.
- i. Click **Next** to complete the configuration.

2.7.2 Rediscovering Oracle Fusion Applications Targets

As you work with your Oracle Fusion Applications environment, you may add additional targets through provisioning and scalability. When these types of changes occur, you can refresh the instance to add the new components as targets within a Fusion Instance already being managed by Cloud Control. You will typically do this because you need to:

- Add a new product family that is part of an already managed Fusion Instance as an additional target
- Add new instances within a product Family or an Oracle WebLogic Server domain that is already being managed by Cloud Control as new targets

To re-discover Oracle Fusion Applications environment:

1. From the **Targets** menu, choose **Fusion Applications**.
The Fusion Applications target home page displays.
2. From the table, choose the **Fusion Instance** you want to refresh.
3. From the **Fusion Instance** menu, choose **Refresh Fusion Instance**.
The Refresh Fusion Instance wizard opens, pre-populated with existing data for the Fusion Instance.
4. Use the wizard to complete the refresh.

2.7.3 Removing Oracle Fusion Applications Targets

To re-discover the Oracle Fusion Applications environment:

1. From the **Targets** menu, choose **Fusion Applications**.
The Fusion Applications target home page displays.
2. From the table, choose the **Fusion Instance**, **Product Family**, or **Product** you want to remove.
3. From the **Fusion Instance**, **Product Family**, or **Product** menu, choose **Target Setup > Remove Target**.
4. In the Remove Target page, choose **Yes** to remove the target.

2.7.4 Navigating the Fusion Applications Target Home Page in Cloud Control

Figure 2–3 shows the Fusion Applications target home page. This page displays the instance as a whole, product families, and products, giving an aggregate status of incidents and recent configuration changes.

Figure 2–3 Fusion Applications Target Home Page of Cloud Control

Name	Type	Database	Incidents				Configuration Changes (last 7 days)	Patches Applied
			Fatal	Critical	Warning	Escalation		
▼ Fusion1	Fusion Instance	Fusion1	0	21	88	0	0	0
▶ Oracle Fusion Financials	Fusion Applications Financials Product F...		0	21	14	0	0	
▶ Oracle Fusion Financials and Supply Chain Management	Fusion Applications Financials and Suppl...		0	0	2	0	0	
▼ Fusion2	Fusion Instance	Fusion2	0	9	18	0	0	0
▶ Oracle Fusion Applications Technology	Fusion Applications ATF Product Family		0	3	0	0	0	
▶ Oracle Fusion Common	Fusion Applications Common Product Fa...		0	1	0	0	0	
▶ Oracle Fusion Governance, Risk, and Compliance	Fusion Applications GRC Product Family		0	1	0	0	0	
▶ Oracle Fusion Incentive Compensation	Fusion Applications IC Product Family		0	1	2	0	0	
▶ Oracle Fusion Projects	Fusion Applications Projects Product Fa...		0	1	0	0	0	
▶ Oracle Fusion Setup	Fusion Applications SETUP Product Family		0	7	0	0	0	
▼ Fusion3	Fusion Instance	Fusion3	69	406	9	0	148	0
▶ Oracle Fusion Customer Relationship Management	Fusion Applications CRM Product Family		0	362	3	0	18	
▶ Oracle Fusion Financials	Fusion Applications Financials Product F...		2	21	0	0	53	
▶ Oracle Fusion Financials and Supply Chain Management	Fusion Applications Financials and Suppl...		0	0	0	0	8	
▶ Oracle Fusion Incentive Compensation	Fusion Applications IC Product Family		0	0	0	0	12	
▶ Oracle Fusion Procurement	Fusion Applications Procurement Produ...		0	0	0	0	0	

To access the Fusion Applications target home page in Cloud Control:

1. From the **Targets** menu, choose **Fusion Applications**.

The Fusion Applications target home page displays. The table on the page lists the Fusion Instances, product families and products for the environment.

The **Targets** menu enables you to select the type of target you want to monitor. For more information about the **Middleware** option, see [Section 10.5.1](#).

2. View the current data reported for each target.

The Fusion Applications overview page displays a table of Fusion instances. You can expand a given Fusion instance to display the Fusion Applications product families deployed to that instance. Each Fusion Applications product family can be expanded to display the applications within each product family. The following information displays for each item in the table.

- **Name:** This column displays the name of the Fusion instance or Fusion Applications product family.
- **Type:** This column indicates whether the item in the row is a Fusion instance, a Fusion Applications product family or a product.
- **Database:** This column displays the name of the database to which the Fusion instance is written.
- **Incidents:** This column displays incidents that have occurred within the scope of each item in the table. Incidents are categorized in sub-columns from left to right as follows: fatal, critical, warning and escalation level.

- **Configuration Changes (last 7 days):** This column displays any configuration changes to the given item in the row during the last seven days.
- **Patches Applied:** This column displays the number of software patches applied to the Fusion instance.

The Fusion Applications overview page allows you to create and view Fusion instances. Use the **View** menu to expand and collapse instances in the table. Click the **Add Instance** button to add a Fusion instance as a managed target in Cloud Control.

You can drill down to a Fusion instance, product family or product by clicking the relevant item in the table.

From this page, you can also access the global operations for the enterprise from the **Enterprise** menu:

- **Summary** enables to view the aggregate status of all managed targets, giving a big picture of system availability, alerts, and job status. Use this page to manage the grid by exception. Even with hundreds of managed targets, this page enables you to quickly detect and begin troubleshooting of problem areas without being distracted by noncritical issues.
- **Monitoring** provides options for global monitoring.
- **Job** provides information about all scheduled, currently running, and past jobs. It also enables you to display and manage jobs in the **job library**. The job library allows users to share and re-use library jobs that have been created.
- **Reports** enables you to manage all report definitions to which you have been granted access.
- **Configuration** enables you to manage configuration changes. For example, you can compare configuration to track down updates.
- **Compliance** enables you to view the policies for specific targets and obtain a summary of all the policy violations in your enterprise.
- **Quality Management** provide a complete testing solution for the Oracle Fusion applications.
- **My Oracle Support** provides access to software and patches.

2.7.5 Using Cloud Control Help

If you need information about the page or the use cases supported by a particular functionality exposed by the page, then click the **Help** link on that page at the top of the page. In most cases, the Help window displays a help topic about the current page. Click **Contents** in the Help window to browse the list of help topics, or click **Search** to search for a particular word or phrase

2.8 Using Oracle Fusion Functional Setup Manager

Oracle Fusion Functional Setup Manager guides you through the tasks of setting up the functional data and reference data after installation and helps you move the setup data from one environment to another (such as from the development environment to the production environment).

To access the administrative area of Oracle Fusion Functional Setup Manager:

1. Sign in to Oracle Fusion Applications with a user account that is provisioned with the necessary role. Contact your security administrator for details.

2. From the **Administration** menu in the work area of Oracle Fusion Applications, choose **Setup and Maintenance**.

2.9 Using Database Tools to Modify Oracle Fusion Applications Data

Do not use SQL*Plus, Oracle Data Browser, database triggers, or any other tool to modify Oracle Fusion Applications data unless otherwise instructed.

Oracle provides powerful tools you can use to create, store, change, retrieve, and maintain information in an Oracle database. But if you use Oracle tools such as SQL*Plus to modify Oracle Fusion Applications data, you risk destroying the integrity of your data and you lose the ability to audit changes to your data.

Because Oracle Fusion Applications tables are interrelated, any change you make using an Oracle Fusion Applications form can update many tables at once. But when you modify Oracle Fusion Applications data using anything other than Oracle Fusion Applications, you may change a row in one table without making corresponding changes in related tables.

If your tables get out of synchronization with each other, you risk retrieving erroneous information and you risk unpredictable results throughout Oracle Fusion Applications.

When you use Oracle Fusion Applications to modify your data, Oracle Fusion Applications automatically checks that your changes are valid. Oracle Fusion Applications also keeps track of who changes information. If you enter information into database tables using database tools, you may store invalid information. You also lose the ability to track who has changed your information because SQL*Plus and other database tools do not keep a record of changes.

Performing Routine Administrative Tasks

This chapter describes the common, essential tasks for administering your Oracle Fusion Applications environment.

This chapter contains the following topics:

- [Introduction to Performing Routine Administrative Tasks](#)
- [Running Administration Servers and Managed Servers from Shared and Local Disks](#)
- [Starting and Stopping](#)
- [Viewing and Changing Ports for Components](#)
- [Changing Passwords](#)
- [Managing the Oracle Database](#)
- [Patching](#)
- [Managing Oracle Fusion Applications-Specific Labels in the Oracle Metadata Repository](#)
- [Modifying Oracle Application Development Framework Connections in Oracle Fusion Applications](#)

3.1 Introduction to Performing Routine Administrative Tasks

An important aspect of system administration of the Oracle Fusion Applications environment is performing a range of maintenance life-cycle and management operations. [Section 1.3](#) provides a roadmap of when to perform these essential administrative tasks. In addition to following this roadmap, if you run into performance or security issues, then use the tasks in the chapter to maintain the environment.

In addition to the tasks in this guide, refer to the following documentation for tasks that you may need to perform, depending on your business needs:

- *Oracle Fusion Applications Common Implementation Guide* for common setup and maintenance tasks for Oracle Fusion applications
- *Oracle Fusion Applications Post-Installation Guide* for configuration tasks to perform after installation
- *Oracle Fusion Applications Administrator's Troubleshooting Guide* for common issues in the runtime environment

3.2 Running Administration Servers and Managed Servers from Shared and Local Disks

Configuration for domains is stored centrally on a shared disk. This location is accessible to all hosts, so the Administration Servers can run from this shared location. Managed Servers can run from either the shared location or from a local disk, that is, a non-network disk, visible only to the processes running on that host. During provisioning of Oracle Fusion Applications, in the Installation Location page of the Provisioning Wizard, you make this decision with the **Enable Local Application Configuration** option. When you select this option, provisioning copies the domain configuration from the shared server location to the specified local disk, and you run the Managed Servers from the local disk. If you chose not to select this option, then all the Managed Servers are configured to run from the shared disk. See the "Installation Location" section in the *Oracle Fusion Applications Installation Guide* for information about using the Installation Location page.

3.3 Starting and Stopping

You start and stop the Oracle Fusion Applications environment and its components to perform a range of maintenance operations that require process downtime. Understanding the state (that is, up, down, starting, and stopping) of each component in the Oracle Fusion Applications environment is an essential activity when diagnosing and resolving availability and performance issues, and when performing life-cycle and management operations.

A typical Oracle Fusion Applications environment contains the applications, the Oracle Database, and the Oracle Fusion Middleware components. Depending on whether you need to start and stop an individual component or the entire environment, there are many several supported scenarios for performing these operations correctly.

This sections contains the following topics:

- [Starting and Stopping Components in the Oracle Fusion Applications Environment](#)
- [Starting and Stopping the Entire Oracle Fusion Applications Environment](#)
- [Starting and Stopping the Oracle Fusion Applications Middle Tier Using the fastartstop Utility](#)
- [Starting and Stopping a Product Family Oracle WebLogic Server Domain](#)
- [Starting and Stopping an Oracle WebLogic Server Cluster for a Configuration Change](#)
- [Starting and Stopping Specific Applications](#)

3.3.1 Starting and Stopping Components in the Oracle Fusion Applications Environment

For many of the Oracle Fusion Middleware components, you can start and stop in different ways, depending on your requirements.

[Table 3–1](#) describes the Oracle Fusion Applications, Oracle Fusion Middleware, and Oracle Database components you can start and stop, and provides information on where to find related documentation for each component.

The procedures referenced in the Oracle Fusion Middleware guides describe using Fusion Middleware Control. These procedure also apply to Fusion Applications Control.

Table 3–1 Starting and Stopping

Component	Description	Documentation
Oracle Fusion Applications		
Oracle Fusion applications	You can start and stop applications using the WLST command line, Fusion Applications Control, Cloud Control, or WebLogic Server Administration Console.	Section 3.3.6 for WLST and Fusion Applications Control "Start applications and modules" and "Stop applications and modules" in the <i>Oracle Fusion Middleware Oracle WebLogic Server Administration Console Online Help</i> for the WebLogic Server Administration Console
Oracle Fusion Middleware		
Administration Server	The Administration server operates as the central control entity for the configuration of the entire Oracle WebLogic Server domain. It maintains the domain's configuration documents and distributes changes in the configuration documents to Managed Servers. The Administration Server serves as a central location from which to manage and monitor all resources in a domain. You can use the WLST command line to start an Administration Server or the <code>fastartstop</code> utility to start all the Administration Server for all the product families.	Section 3.3.2.1, Task 3, "Start the Administration Servers"
Node Manager	Node Manager is an Oracle WebLogic Server utility that enables you to start, shut down, and restart the Administration Servers and the Managed Server instances. On Windows, Node Manager is configured to automatically start by default.	"Starting Node Manager" section in the <i>Oracle Fusion Middleware Node Manager Administrator's Guide for Oracle WebLogic Server</i> Task 3, "Start Node Manager"
Managed Server for an application	Managed Servers host business applications, application components, Web services, and their associated resources. You can use Fusion Applications Control or WebLogic Server Administration Console to start the Managed Servers for a domain, or the <code>fastartstop</code> utility to start all the Managed Servers for all the product families.	Section 3.3.2.1, Task 4, "Start the Managed Servers" for Fusion Applications Control and the <code>fastartstop</code> script "Start Managed Servers from the Administration Console", "Start Managed Servers in a cluster" and "Shutdown servers in a cluster" in the <i>Oracle Fusion Middleware Oracle WebLogic Server Administration Console Online Help</i> for the WebLogic Server Administration Console
Oracle Business Intelligence Suite	You can use Fusion Applications Control or the <code>opmnctl</code> command to start and stop Oracle Business Intelligence system components.	"Starting and Stopping Oracle Business Intelligence" chapter in the <i>Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition</i>
Oracle Enterprise Scheduler	You can use Fusion Applications Control to start and stop the Oracle Enterprise Scheduler.	Section 5.7

Table 3–1 (Cont.) Starting and Stopping

Component	Description	Documentation
Oracle HTTP Server	You can use Fusion Applications Control or the <code>opmnctl</code> command to start and stop Oracle HTTP Server.	"Starting, Stopping, and Restarting Oracle HTTP Server" section in the <i>Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server</i>
Oracle WebCenter Portal: Spaces	Always ensure Content Server is running when starting Oracle WebCenter Portal: Spaces. If Content Server is not running when Spaces is started up Spaces will consider the Document service as unavailable. After Content Server is started, a restart of Spaces is needed for it to recheck if the Document service is configured.	"Starting Content Server" section in the <i>Oracle WebCenter Content System Administrator's Guide for Content Server</i> "Starting and Stopping the Spaces Application" section in the <i>Oracle WebCenter Content System Administrator's Guide for Content Server</i>
Oracle WebLogic Server domain for a product family	When applying a patch that only applies to a product family, you can stop the Oracle WebLogic Server domain for the product family and restart it after you apply the patch without stopping the entire Oracle Fusion Applications environment. You can use Fusion Applications Control to start and stop an Oracle WebLogic Server domain.	Section 3.3.4
Oracle WebLogic Server cluster for a product family	If you modify the Oracle WebLogic Server configuration, you must stop and restart the cluster for the configuration changes to take effect. You can start and stop Oracle WebLogic Server clusters with either Fusion Applications Control or Node Manager.	Section 3.3.5.2 for Fusion Applications Control "Using Node Manager to Start Managed Servers in a WebLogic Domain or Cluster" section in <i>Oracle Fusion Middleware Oracle WebLogic Scripting Tool for Node Manager</i>
Oracle Database		
Oracle Database	You can start and stop the Oracle database with the Oracle Enterprise Manager Database Control.	"Accessing the Database Home Page" section in the <i>Oracle Database 2 Day DBA</i>

3.3.2 Starting and Stopping the Entire Oracle Fusion Applications Environment

Some components in the Oracle Fusion Applications environment dependent on one another. Therefore, it is important to start and stop components in the proper order. In the course of normal IT operations, common operations include shutting down computers and starting them back up. Therefore, it is crucial to start and stop Oracle Fusion Applications in a sequential manner.

This section contains the following topics:

- [Starting an Oracle Fusion Applications Environment](#)
- [Stopping an Oracle Fusion Applications Environment](#)

3.3.2.1 Starting an Oracle Fusion Applications Environment

To perform a complete start of the Oracle Fusion Applications environment, you must start the Oracle WebLogic Server domain for the Oracle Fusion Setup product family before the domains for the other product families.

Applications may not work correctly until all the Managed Servers in all domains have been started.

The section contains the following topics for performing a complete start of the Oracle Fusion Applications environment:

- [Task 1, "Start the Oracle Database"](#)
- [Task 2, "Start the Oracle Identity Management Suite"](#)
- [Task 3, "Start the Administration Servers"](#)
- [Task 4, "Start the Managed Servers"](#)
- [Task 5, "Start Oracle Business Intelligence"](#)
- [Task 6, "Start the Oracle HTTP Server"](#)

The `faststartstop` utility provides an alternative way to starting the Administration Servers and Managed Servers for all the product families in one step. See [Section 3.3.3](#).

Task 1 Start the Oracle Database

To start the Oracle database instance using Oracle Database Control:

1. Go to the Database Home page.

For information about how to access the Database Home page in Oracle Enterprise Manager, see "Accessing the Database Home Page" section in the *Oracle Database 2 Day DBA*.

2. In the Database Home page, click **Startup**.

The Startup/Shutdown Credentials page appears.

3. Enter credentials as follows:

- a. Enter the host computer credentials of the user who installed Oracle Database, or of any user who is authorized to use SQL*Plus.
- b. Enter the database credentials consisting of the user name `SYS` and the password that you assigned to `SYS` during the installation.
- c. From the **Connect As** list, select **SYSOPER**.

4. Click **OK**.

A confirmation page appears.

5. Click **Yes**.

The Startup page appears, indicating that the database is being started up. When startup is complete, the Login page appears.

6. Log in to the database.

The Database Home page appears indicating that the database instance status is **Up**.

For information about shutting down and starting up the Oracle database instance, see the "Shutting Down and Starting Up the Oracle Instance" section in the *Oracle Database 2 Day DBA*.

Task 2 Start the Oracle Identity Management Suite

To start the Oracle Identity Management system components:

1. Set the `ORACLE_HOME` environment variable to the Oracle home for the Oracle Identity Management components.
2. Start Oracle Process Manager and Notification Server (OPMN) and all system components:

```
opmnctl startall
```

Task 3 Start the Administration Servers

You must start the Oracle WebLogic Server Administration Server for the `CommonDomain` domain in the Oracle Fusion Setup product family from a shared disk. The applications and other product families depend on topology information in the Administration Server in the Oracle Fusion Setup product family. For more information about running the Administration Server from a shared location, see [Section 3.2](#).

Optionally, if you need to perform administration tasks for other product families, start the Administration Server for those product families.

When you start the Administration Server, you also start the applications that run on the Administration Server, including the WebLogic Server Administration Console and Fusion Applications Control.

To start an individual Administration Server:

1. Start the host-specific Node Manager for the domain for which you want to start the Administration Servers:

- a. Set the `JAVA_OPTIONS` environment variable:

```
setenv JAVA_OPTIONS "${JAVA_OPTIONS} -Dserver.group=AdminServer"
```

- b. Set the `WLST_PROPERTIES` environment variable:

```
setenv WLST_PROPERTIES "-Dweblogic.security.SSL.trustedCAKeyStore=WL_HOME/server/lib/fusion_trust.jks"
```

where `WL_HOME` is the Oracle WebLogic Server home directory under the `fusionapps` Middleware subdirectory.

- c. Start Node Manager with the `startNodeManagerWrapper` utility from the Oracle WebLogic Server home directory under the `fusionapps` Middleware subdirectory:

```
(UNIX) WL_HOME/common/nodemanager/host_name/startNodeManagerWrapper.sh
(Windows) WL_HOME\common\nodemanager\host_name\startNodeManagerWrapper.cmd
```

- d. For all the hosts, verify that Node Manager is running. For example, on UNIX platforms, enter the following command:

```
netstat -a|grep 5556
```

The 5556 port should be in a listening state. If your environment uses a different port, then use the port number in the System Port Allocation page during the creation of a provisioning plan. See the "System Port Allocation" section in the *Oracle Fusion Applications Administrator's Guide*.

Note: Do not start Node Manager with WLST or the following Oracle WebLogic Server scripts:

```
(UNIX) WL_HOME/server/bin/startNodeManager.sh
(Windows) WL_HOME\server\bin\startNodeManager.cmd
```

2. Start the Administration Server:

- a. At an operating system command prompt, navigate to the location of the WLST script:

```
(UNIX) FA_MW_HOME/oracle_common/common/bin
(Windows) FA_MW_HOME\oracle_common\common\bin
```

FA_MW_HOME is named fusionapps and located under the *APPLICATIONS_BASE*.

b. Run the wlst script:

```
(UNIX) wlst.sh
(Windows) wlst.cmd
```

c. Connect to Oracle WebLogic Server with the following WLST command:

```
nmConnect('node_manager_user_name','node_manager_password',
'node_manager_host_name','node_manager_host_port','domain_name',
'path_to_domain_home_of_the_domain')
```

The following shows an example of a connection to the CommonDomain domain:

```
nmConnect('bootstrap_admin','welcome','weblogic1.company.com','5556',
'CommonDomain','/oracle/fusion/top/instance/domains/weblogic1.company.com/C
ommonDomain')
```

See the "nmConnect" section in the *Oracle Fusion Middleware WebLogic Scripting Tool Command Reference*.

d. Start the Administration Server:

```
nmStart('WLS_server_name')
```

For example a connection to the CommonDomain domain might be:

```
nmStart('AdminServer')
```

See the "nmStart" section in the *Oracle Fusion Middleware WebLogic Scripting Tool Command Reference*.

To start all the Administration Servers for all the product families:

1. Start the host-specific Node Managers for all of the product family domains:

Each host has its own copy of the Node Manager and all the Node Managers should be running before attempting to start any Administration or Managed Servers.

a. Set the JAVA_OPTIONS environment variable:

```
setenv JAVA_OPTIONS "${JAVA_OPTIONS} -Dserver.group=AdminServer"
```

b. Set the WLST_PROPERTIES environment variable:

```
setenv WLST_PROPERTIES "-Dweblogic.security.SSL.trustedCAKeyStore=WL_
HOME/server/lib/fusion_trust.jks"
```

where *WL_HOME* is the Oracle WebLogic Server domain home directory under the fusionapps Middleware subdirectory.

c. Start Node Manager with the startNodeManagerWrapper utility from the Oracle WebLogic Server home directory under the fusionapps Middleware subdirectory:

```
(UNIX) WL_HOME/common/nodemanager/host_name/startNodeManagerWrapper.sh
(Windows) WL_HOME\common\nodemanager\host_name\startNodeManagerWrapper.cmd
```

d. For all the hosts, verify that Node Manager is running. For example, on UNIX platforms, enter the following command:

```
netstat -a|grep 5556
```

The 5556 port should be in a listening state. If your environment uses a different port, then use the port number in the System Port Allocation page during the creation of a provisioning plan. See the "System Port Allocation" section in the *Oracle Fusion Applications Administrator's Guide*.

Note: Do not start Node Manager with WLST or the following Oracle WebLogic Server scripts:

```
(UNIX) WL_HOME/server/bin/startNodeManager.sh
(Windows) WL_HOME\server\bin\startNodeManager.cmd
```

2. Start the Administration Servers for all the product families:

- a. If the Oracle Fusion Applications patching framework was installed, the `fastartstop` utility is available from the following directories:

```
(UNIX) FA_ORACLE_HOME/applications/lcm/ad/bin
(Windows) FA_ORACLE_HOME\applications\lcm\ad\bin
```

`FA_ORACLE_HOME` is a directory named `applications`, located under the `fusionapps` Oracle Fusion Applications Middleware home.

- b. From the shared host location, run `fastartstop` to start all the domains or specific domains:

```
(UNIX) fastartstop.sh -Start -all
      -username user_name
      -appbase APPLICATIONS_BASE
      -adminServersOnly
      [-setup setup_details.xml_file_location]
      [-loglevel log_level]
      [-timeout timeout_period]
```

```
(Windows) fastartstop.cmd -Start -all
      -username user_name
      -appbase APPLICATIONS_BASE
      -adminServersOnly
      [-setup setup_details.xml_file_location]
      [-loglevel log_level]
      [-timeout timeout_period]
```

For more information about the syntax, see [Table 3–2](#). For more information about starting the Administration Servers and Managed Servers for all the product families is the `fastartstop` utility. See [Section 3.3.3](#).

Note: Do not use the following Oracle WebLogic Server scripts to start the Administration Servers:

```
(UNIX) DOMAIN_HOME/bin/startWebLogic.sh
(Windows) DOMAIN_HOME\bin\startWebLogic.cmd
```

Task 4 Start the Managed Servers

Depending on the configuration established during provisioning, you can start Oracle WebLogic Server Managed Servers from either a shared or local disk. See [Section 3.2](#).

To start the Managed Servers for a domain using Fusion Applications Control:

1. From the navigation pane, expand the farm, **WebLogic Domain**.
2. From the **WebLogic Domain** menu, choose **Control**, then **Start Up**.

If you prefer to use the WebLogic Server Administration Console, "Start Managed Servers from the Administration Console" and "Start Managed Servers in a cluster" in the *Oracle Fusion Middleware Oracle WebLogic Server Administration Console Online Help*

To start all the Administration Servers and Managed Servers for all the product families in one step, use the `fastartstop` utility. See [Section 3.3.3](#).

Note: Do not use the following Oracle WebLogic Server scripts to start the Managed Servers:

```
(UNIX) DOMAIN_HOME/bin/startManagedWebLogic.sh
(Windows) DOMAIN_HOME\bin\startManagedWebLogic.cmd
```

Task 5 Start Oracle Business Intelligence

See "Starting and Stopping Oracle Business Intelligence" in the *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

Task 6 Start the Oracle HTTP Server

To start the Oracle HTTP Server using `opmnctl`:

```
(UNIX) WT_CONFIG_HOME/bin/opmnctl startproc process-type=OHS
(Windows) WT_CONFIG_HOME\bin\opmnctl startproc process-type=OHS
```

To start Oracle HTTP Server using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Web Tier** installation type.
2. Select the Oracle HTTP Server.
3. From the **Oracle HTTP Server** menu, choose **Control**, then **Start Up**.

3.3.2.2 Stopping an Oracle Fusion Applications Environment

You can follow these procedures when you need to completely shut down the Oracle Fusion Applications environment. For example, when preparing to perform a complete backup of your environment, or apply a patch.

The section contains the following topics for performing a complete stop of the Oracle Fusion Applications environment:

- [Task 1, "Stop the Oracle HTTP Server"](#)
- [Task 2, "Stop Oracle Business Intelligence"](#)
- [Task 3, "Stop the Oracle Identity Management Suite"](#)
- [Task 4, "Stop the Administration Servers"](#)
- [Task 5, "Stop the Managed Servers"](#)
- [Task 6, "Stop the Oracle Database"](#)

The `fastartstop` utility provides an alternative way to stopping the Administration Servers and Managed Servers for all the product families in one step. See [Section 3.3.3](#).

Task 1 Stop the Oracle HTTP Server

To stop the Oracle HTTP Server using `opmnctl`:

```
(UNIX) WT_CONFIG_HOME/bin/opmnctl stopproc process-type=OHS
(Windows) WT_CONFIG_HOME\bin\opmnctl stopproc process-type=OHS
```

To stop the Oracle HTTP Server using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Web Tier** installation type.
2. Select the Oracle HTTP Server.
3. From the **Oracle HTTP Server** menu, choose **Control**, then **Shut Down**.

Task 2 Stop Oracle Business Intelligence

See "Starting and Stopping Oracle Business Intelligence" in the *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

Task 3 Stop the Oracle Identity Management Suite

To stop the Oracle Identity Management system components:

1. Set the `ORACLE_HOME` environment variable to the Oracle home for the Identity Management components.
2. Start OPMN and all system components:

```
opmnctl stopall
```

Task 4 Stop the Administration Servers

When stopping the Oracle Fusion Applications environment, stop the Administration Server for the product families other than the Oracle Fusion Setup product family first, and then stop the Administration Server for the Oracle Fusion Setup product family. The applications and other product families depend on topology information in the Administration Server in the Oracle Fusion Setup product family. Therefore, you need to stop the Administration Server in the Oracle Fusion Setup product family after the other product families. You stop the Administration Servers from a shared disk.

To stop an individual Administration Server:

1. At an operating system command prompt, navigate to the location of the WLST script:

```
(UNIX) FA_MW_HOME/oracle_common/common/bin
(Windows) FA_MW_HOME\oracle_common\common\bin
```

`FA_MW_HOME` is named `fusionapps` and located under the `APPLICATIONS_BASE`.

2. Run the `wlst` script:

```
(UNIX) wlst.sh
(Windows) wlst.cmd
```

3. Connect to Oracle WebLogic Server with the following WLST command:

```
nmConnect('node_manager_user_name','node_manager_password',
'node_manager_host_name','node_manager_host_port', 'domain_name',
'path_to_domain_home_of_the_domain')
```

The following shows an example of a connection to the `CommonDomain` domain:

```
nmConnect('bootstrap_admin','welcome','weblogic1.company.com','5556',
'CommonDomain','/oracle/fusion/top/instance/domains/weblogic1.company.com/Comm
nDomain')
```

See the "nmConnect" section in the *Oracle Fusion Middleware WebLogic Scripting Tool Command Reference*.

4. Stop the Administration Server:

```
nmStop('WLS_server_name')
```

For example, to stop the Administration Server:

```
nmStop('AdminServer')
```

To stop all the Administration Servers for all the product families:

1. Stop the Administration Servers for all the product families:

- a. If the Oracle Fusion Applications patching framework was installed, the `fastartstop` utility is available from the following directories:

```
(UNIX) FA_ORACLE_HOME/applications/lcm/ad/bin
(Windows) FA_ORACLE_HOME\applications\lcm\ad\bin
```

`FA_ORACLE_HOME` is a directory named `applications`, located under the `fusionapps` Oracle Fusion Applications Middleware home.

- b. From the shared host location, run `fastartstop` to start all the domains or specific domains:

```
(UNIX) fastartstop.sh -Stop -all
-username user_name
-appbase APPLICATIONS_BASE
-adminServersOnly
[-setup setup_details.xml_file_location]
[-loglevel log_level]
[-timeout timeout_period]
```

```
(Windows) fastartstop.cmd -Stop -all
-username user_name
-appbase APPLICATIONS_BASE
-adminServersOnly
[-setup setup_details.xml_file_location]
[-loglevel log_level]
[-timeout timeout_period]
```

For more information about the syntax, see [Table 3–2](#). For more information about stopping the Administration Servers and Managed Servers for all the product families is the `fastartstop` utility. See [Section 3.3.3](#).

Note: Do not use the following Oracle WebLogic Server scripts to stop the Administration Servers:

```
(UNIX) DOMAIN_HOME/bin/stopWebLogic.sh
(Windows) DOMAIN_HOME\bin\stopWebLogic.cmd
```

Task 5 Stop the Managed Servers

Depending on the configuration established during provisioning, you can stop Oracle WebLogic Server Managed Servers from either a shared or local disk. See [Section 3.2](#).

To stop the Managed Servers for a domain using Fusion Applications Control:

1. From the navigation pane, expand the farm, and then **WebLogic Domain**.
2. From the **WebLogic Domain** menu, choose **Control**, then **Shut Down**.

If you prefer to use the WebLogic Server Administration Console, see "Start Managed Servers from the Administration Console" and "Shutdown servers in a cluster" in the *Oracle Fusion Middleware Oracle WebLogic Server Administration Console Online Help*

To stop all the Administration Servers and Managed Servers for all the product families in one step, use the `faststartstop` utility. See [Section 3.3.3](#).

Note: Do not use the following Oracle WebLogic Server scripts to stop the Managed Servers:

```
(UNIX) DOMAIN_HOME/bin/stopManagedWebLogic.sh  
(Windows) DOMAIN_HOME\bin\stopManagedWebLogic.cmd
```

Task 6 Stop the Oracle Database

To stop the Oracle database instance using Oracle Database Control:

1. Go to the Database Home page.

For information about how to access the Database Home page in Oracle Enterprise Manager, see "Accessing the Database Home Page" section in the *Oracle Database 2 Day DBA*.

2. In the Database Home page, click **Shutdown**.

The Startup/Shutdown Credentials page appears.

3. Enter credentials as follows:

- a. Enter the host computer credentials of the user who installed Oracle Database, or of any user who is authorized to use SQL*Plus.
- b. Enter the database credentials consisting of the user name `SYS` and the password that you assigned to `SYS` during the installation.
- c. From the **Connect As** list, select **SYSOPER**.

4. (Optional) Select the **Save as Preferred Credential** option if you want these credentials to be automatically filled in for you the next time that this page appears.

5. Click **OK**.

A confirmation page appears.

6. Click **Yes** to start the `SHUTDOWN IMMEDIATE` operation on the database.

The Shutdown: Activity Information page appears, informing you that the database is being shut down.

7. After a short period (approximately 2 minutes), click **Refresh** to be returned to the Database Home page.

The Database Home page indicates that the database instance status is Down.

For information about shutting down and starting up the Oracle database instance, see the "Shutting Down and Starting Up the Oracle Instance" section in the *Oracle Database 2 Day DBA*.

3.3.3 Starting and Stopping the Oracle Fusion Applications Middle Tier Using the fastartstop Utility

Another alternative to starting and stopping the Oracle Fusion Applications Administration Servers and Managed Servers is the `fastartstop` utility.

The `fastartstop` script installed with the Oracle Fusion Applications patching framework. After installation, the registry file for each product family is populated with Oracle WebLogic Server domain entries for all product families, including the Oracle Fusion Applications Common Domain in the Oracle Fusion Setup product family. You can use this utility to start all the domains or a select domain. When the `fastartstop` utility is used to start a domain, all the Managed Servers will be started and the Administration Server of the domain will be stopped after all the Managed Servers have been started. The `fastartstop` utility requires Node Manager to run.

While the utility starts and stops the domains, it does not start and stop Oracle HTTP Server, Oracle Identity Management, and the Oracle Database. Therefore, you must start these components separately.

This section contains the following topics:

- [fastartstop Syntax](#)
- [Registry Entries](#)
- [Starting an Oracle Fusion Applications Environment with the fastartstop Utility for the Middle Tier](#)
- [Stopping an Oracle Fusion Applications Environment with the fastartstop Utility](#)

3.3.3.1 fastartstop Syntax

If the Oracle Fusion Applications patching framework was installed, the `fastartstop` utility is available from the following directories:

```
(UNIX) FA_ORACLE_HOME/applications/lcm/ad/bin
(Windows) FA_ORACLE_HOME\applications\lcm\ad\bin
```

`FA_ORACLE_HOME` is a directory named `applications`, located under the `fusionapps` Oracle Fusion Applications Middleware home.

```
(UNIX) fastartstop.sh -Start|-Stop|-Bounce|-StartBIPS|-StopBIPS
-all|-domains domain_name, ...domain_nameN|-BIPServerComponents
[-clusterType soa]
-username user_name
-appbase APPLICATIONS_BASE
[-adminServersOnly]
[-startAdminServer true|false]
[-setup setup_details.xml_file_location]
[-loglevel log_level]
[-timeout timeout_period]
[-help]
```

```
(Windows) fastartstop.cmd -Start|-Stop|-Bounce|-StartBIPS|-StopBIPS
-all|-domains domain_name, ...domain_nameN|-BIPServerComponents
[-clusterType soa]
-username user_name
-appbase APPLICATIONS_BASE
[-adminServersOnly]
[-startAdminServer true|false]
[-setup setup_details.xml_file_location]
[-loglevel log_level]
```

```
[-timeout timeout_period]
[-help]
```

The following example shows starting Managed Servers in SOA cluster for the HCMDomain domain and the Administration Server of that domain.

```
fastartstop.sh -Start -domains HCMDomain -username username
appbase /scratch/aimel/appltop/fusionapps/applications
-clusterType soa -startAdminServer true
```

Table 3–2 describes the options for the fastartstop script.

Table 3–2 Options for the fastartstop Script

Options	Description
-appbase	Specify the path for the Oracle Fusion Applications installation.
-adminServersOnly	Use to start, stop, or restart the Administration Servers only, leaving Managed Servers untouched. This option is mutually exclusive to the <code>-startAdminServer</code> option. You cannot include both the <code>-adminServersOnly</code> and the <code>-startAdminServer</code> options in the same <code>fastartstop</code> command. When used with <code>Bounce</code> , <code>fastartstop</code> supports retry for Administration Servers. If some Administration Servers fail during the <code>Bounce</code> command, in next run, failed servers will be acted upon. The servers bounced successfully in the previous runs will be skipped in the next run.
-all -domains -BIPServer Components	Specify what to start, stop, or bounce: <ul style="list-style-type: none"> Use <code>-all</code> for to start or stop all domains. Use <code>-domains</code> for specific domains. If using the <code>-domains</code> option, delineate the domains names with a comma and space (<i>domain_name, domain_name</i>). Use <code>-BIPServerComponents</code> for specific BI Presentation Server components. This option should be used only in conjunction with the <code>-StartBIPS</code> or <code>-StopBIPS</code> option.
-clusterType	Specify to start or stop Managed Servers within a cluster. This parameter is currently only supported for Oracle SOA Suite clusters (<code>soa</code>).
-help	Specify to display the syntax.
-loglevel	Specify the log level for the <code>fastartstop.log</code> file: <ul style="list-style-type: none"> INFO: Provides informational messages. FINER: Provides detailed log messages. WARNING: Provides messages for potentially problems. INFO is the default.
-Start -Stop -Bounce -StartBIPS -StopBIPS	Specify whether to start, stop, or restart: <ul style="list-style-type: none"> Use <code>-Start</code> to start Managed Servers. Use <code>-Stop</code> to stop Managed Servers. Use <code>-Bounce</code> to restart Managed Servers. Use <code>-StartBIPS</code> to start only BI Presentation Server components. Use <code>-StopBIPS</code> to stop only BI Presentation Server components.

Table 3–2 (Cont.) Options for the fastartstop Script

Options	Description
-setup	Specify the location of the SetupDetails.xml file. If this parameter is not specified, then the following defaults are used: (UNIX) <code>FA_ORACLE_HOME/admin</code> (Windows) <code>FA_ORACLE_HOME\admin</code> <code>FA_ORACLE_HOME</code> is a directory named applications, located under the fusionapps Oracle Fusion Applications Middleware home.
-startAdminServer	Specify true to start the Administration Server; specify false to stop the Administration Server. The default is false. This option is mutually exclusive to the -adminServersOnly option. You cannot include both the -adminServersOnly and the -startAdminServer options in the same fastartstop command.
-timeout	Specify the timeout in seconds for the utility to run.
-username	Specify the Oracle Fusion Middleware administration user.

The utility prompts you for the Oracle Fusion Middleware administration user. See [Section 4.8.5.2](#) for more information about changing this password.

The fastartstop utility generates fastartstop.log in the following subdirectories:

(UNIX) `FA_ORACLE_HOME/lcm/ad/bin`
(Windows) `FA_ORACLE_HOME\lcm\ad\bin`

`FA_ORACLE_HOME` is a directory named applications, located under the fusionapps Oracle Fusion Applications Middleware home.

3.3.3.2 Registry Entries

[Example 3–1](#) shows the format of a typical entry in the SetupDetails.xml. It is populated with the CommonDomain domains details. The other domains are not represented, because they are automatically read from the API. The file contains Oracle WebLogic Server domain information, including the Administration Server, the Managed Servers, and Node Manager.

Example 3–1 Registry File Format

```
<ServerInfo DomainName="CommonDomain" DomainType="CommonDomain"
ServerHostName="host_name AdminServerName="AdminServer" AdminPort="port"
DomainRootDir="DOMAIN_HOME" NodeManagerHost="host_name"
NodeManagerPort="port" />
```

3.3.3.3 Starting an Oracle Fusion Applications Environment with the fastartstop Utility for the Middle Tier

You can reduce the number of steps required to start the Oracle Fusion Applications environment with the fastartstop utility.

The section contains the following topics for starting the Oracle Fusion Applications environment with the fastartstop utility:

- [Task 1, "Start the Oracle Database"](#)
- [Task 2, "Start the Oracle Identity Management Suite"](#)
- [Task 3, "Start Node Manager"](#)

- [Task 4, "Run the faststartstop Utility to Start"](#)
- [Task 5, "Start the Oracle HTTP Server"](#)

Task 1 Start the Oracle Database

To start the Oracle database instance using Oracle Database Control:

1. Go to the Database Home page.
For information about how to access the Database Home page in Oracle Enterprise Manager, see "Accessing the Database Home Page" section in the *Oracle Database 2 Day DBA*.
2. In the Database Home page, click **Startup**.
The Startup/Shutdown Credentials page appears.
3. Enter credentials as follows:
 - a. Enter the host computer credentials of the user who installed Oracle Database, or of any user who is authorized to use SQL*Plus.
 - b. Enter the database credentials consisting of the user name `SYS` and the password that you assigned to `SYS` during the installation.
 - c. From the **Connect As** list, select **SYSOPER**.
4. Click **OK**.
A confirmation page appears.
5. Click **Yes**.
The Startup page appears, indicating that the database is being started up. When startup is complete, the Login page appears.
6. Log in to the database.
The Database Home page appears indicating that the database instance status is **Up**.

For information about shutting down and starting up the Oracle database instance, see the "Shutting Down and Starting Up the Oracle Instance" section in the *Oracle Database 2 Day DBA*.

Task 2 Start the Oracle Identity Management Suite

To start the Oracle Identity Management system components:

1. Set the `ORACLE_HOME` environment variable to the Oracle home for the Identity Management components.
2. Start OPMN and all system components:

```
opmnctl startall
```

Task 3 Start Node Manager

Start Node Manager for each of the hosts for which you want to start the Administration Servers:

1. Set the `JAVA_OPTIONS` environment variable:

```
setenv JAVA_OPTIONS "${JAVA_OPTIONS} -Dserver.group=AdminServer"
```
2. Set the `WLST_PROPERTIES` environment variable:

```
setenv WLST_PROPERTIES "-Dweblogic.security.SSL.trustedCAKeyStore=WL_
```

```
HOME/server/lib/fusion_trust.jks"
```

where `WL_HOME` is the Oracle WebLogic Server home directory under the `fusionapps` Middleware subdirectory.

3. Start Node Manager with the `startNodeManagerWrapper` utility from the Oracle WebLogic Server home directory under the `fusionapps` Middleware subdirectory:

```
(UNIX) WL_HOME/nodemanager/host_name/startNodeManagerWrapper.sh
(Windows) WL_HOME\nodemanager\host_name\startNodeManagerWrapper.cmd
```

4. For all the hosts, verify that Node Manager is running. For example, on UNIX platforms, enter the following command:

```
netstat -a|grep 5556
```

The 5556 port should be in a listening state. If your environment uses a different port, then use the port number in the System Port Allocation page during the creation of a provisioning plan. See the "System Port Allocation" section in the *Oracle Fusion Applications Administrator's Guide*.

Task 4 Run the `faststartstop` Utility to Start

From the shared host location, run `faststartstop` to start all the domains or specific domains:

```
(UNIX) faststartstop.sh -Start
-all|-domains domain_name, ...domain_nameN
[-clusterType soa]
-username user_name
-appbase APPLICATIONS_BASE
[-startAdminServer true]
[-setup setup_details.xml_file_location]
[-loglevel log_level]
[-timeout timeout_period]
```

```
(Windows) faststartstop.cmd -Start
-all|-domains domain_name, ...domain_nameN
[-clusterType soa]
-username user_name
-appbase APPLICATIONS_BASE
[-startAdminServer true]
[-setup setup_details.xml_file_location]
[-loglevel log_level]
[-timeout timeout_period]
```

For more information about the syntax, see [Table 3-2](#).

Task 5 Start the Oracle HTTP Server

To start the Oracle HTTP Server using `opmnctl`:

```
(UNIX) WT_CONFIG_HOME/bin/opmnctl startproc process-type=OHS
(Windows) WT_CONFIG_HOME\bin\opmnctl startproc process-type=OHS
```

To start Oracle HTTP Server using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Web Tier** installation type.
2. Select the Oracle HTTP Server.
3. From the **Oracle HTTP Server** menu, choose **Control**, then **Start Up**.

3.3.3.4 Stopping an Oracle Fusion Applications Environment with the fastartstop Utility

You can reduce the number of steps required to stop the Oracle Fusion Applications environment with the `fastartstop` utility.

The section contains the following topics for stopping the Oracle Fusion Applications environment with the `fastartstop` utility:

- [Task 1, "Stop the Oracle HTTP Server"](#)
- [Task 2, "Start Node Manager"](#)
- [Task 3, "Run the fastartstop Utility to Stop"](#)
- [Task 4, "Stop the Oracle Identity Management Suite"](#)
- [Task 5, "Stop the Oracle Database"](#)

Task 1 Stop the Oracle HTTP Server

To stop the Oracle HTTP Server using `opmnctl`:

```
(UNIX) WT_CONFIG_HOME/bin/opmnctl stopproc process-type=OHS
(Windows) WT_CONFIG_HOME\bin\opmnctl stopproc process-type=OHS
```

To stop the Oracle HTTP Server using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Web Tier** installation type.
2. Select the Oracle HTTP Server.
3. From the **Oracle HTTP Server** menu, choose **Control**, then **Shut Down**.

Task 2 Start Node Manager

The `fastartstop` utility requires the Node Manager be running.

To start Node Manager, use the `startNodeManagerWrapper` utility from the Oracle WebLogic Server home directory under the `fusionapps` Middleware subdirectory:

```
(UNIX) WL_HOME/nodemanager/host_name/startNodeManagerWrapper.sh
(Windows) WL_HOME\nodemanager\host_name\startNodeManagerWrapper.cmd
```

Task 3 Run the fastartstop Utility to Stop

From the shared host location, run `fastartstop` to stop all the domains or specific domains:

```
(UNIX) fastartstop.sh -Stop
-all|-domains domain_name, ...domain_nameN
[-clusterType soa]
-username user_name
-appbase APPLICATIONS_BASE
[-setup setup_details.xml_file_location]
[-loglevel log_level]
[-timeout timeout_period]

(Windows) fastartstop.cmd -Stop
-all|-domains domain_name, ...domain_nameN
[-clusterType soa]
-username user_name
-appbase APPLICATIONS_BASE
[-setup setup_details.xml_file_location]
[-loglevel log_level]
```

```
[-timeout timeout_period]
```

For more information about the syntax, see [Table 3–2](#).

Task 4 Stop the Oracle Identity Management Suite

To stop the Oracle Identity Management system components:

1. Set the `ORACLE_HOME` environment variable to the Oracle home for the Identity Management components.
2. Stop OPMN and all system components:

```
opmnctl stopall
```

Task 5 Stop the Oracle Database

To stop the Oracle database instance using Oracle Database Control:

1. Go to the Database Home page.

For information about how to access the Database Home page in Oracle Enterprise Manager, see "Accessing the Database Home Page" section in the *Oracle Database 2 Day DBA*.

2. In the Database Home page, click **Shutdown**.

The Startup/Shutdown Credentials page appears.

3. Enter credentials as follows:

- a. Enter the host computer credentials of the user who installed Oracle Database, or of any user who is authorized to use SQL*Plus.
- b. Enter the database credentials consisting of the user name `SYS` and the password that you assigned to `SYS` during the installation.
- c. From the **Connect As** list, select **SYSOPER**.

4. (Optional) Select the **Save as Preferred Credential** option if you want these credentials to be automatically filled in for you the next time that this page appears.

5. Click **OK**.

A confirmation page appears.

6. Click **Yes** to start the `SHUTDOWN IMMEDIATE` operation on the database.

The Shutdown: Activity Information page appears, informing you that the database is being shut down.

7. After a short period (approximately 2 minutes), click **Refresh** to be returned to the Database Home page.

The Database Home page indicates that the database instance status is Down.

For information about shutting down and starting up the Oracle database instance, see the "Shutting Down and Starting Up the Oracle Instance" section in the *Oracle Database 2 Day DBA*.

3.3.4 Starting and Stopping a Product Family Oracle WebLogic Server Domain

When applying a patch that only applies to a product family, you can stop the Oracle WebLogic Server domain for the product family and restart it after you apply the

patch. Your environment may require a restart of a domain for reasons other than a patch.

This section contains the following topics:

- [Stopping an Oracle WebLogic Server Domain for a Product Family](#)
- [Starting an Oracle WebLogic Server Domain for a Product Family](#)

3.3.4.1 Stopping an Oracle WebLogic Server Domain for a Product Family

You stop an Oracle WebLogic Server domain for a product family before applying a patch.

This section contains the following topics for stopping an Oracle WebLogic Server domain for a product family:

- [Task 1, "Stop the Oracle HTTP Server"](#)
- [Task 2, "Stop the Product Family Domain"](#)

Task 1 Stop the Oracle HTTP Server

Oracle HTTP Server is typically installed on a computer separate from the product family. By stopping Oracle HTTP Server, you stop requests to the entire Oracle Fusion Applications environment. Stopping Oracle HTTP Server before the product family ensures a cleaner shutdown.

To stop the Oracle HTTP Server using `opmnctl`:

```
(UNIX) WT_CONFIG_HOME/bin/opmnctl stopproc process-type=OHS  
(Windows) WT_CONFIG_HOME\bin\opmnctl stopproc process-type=OHS
```

To stop Oracle HTTP Server using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Web Tier** installation type.
2. Select an Oracle HTTP Server.
3. From the **Oracle HTTP Server** menu, choose **Control**, then **Shut Down**.

Task 2 Stop the Product Family Domain

When you stop an Oracle WebLogic Server domain for a product family, you stop all the applications on the cluster of Managed Servers in the domain. For example, the Oracle Enterprise Scheduler `ESSAPP` application and the SOA Infrastructure `soa-infra` application stop running.

To stop a product family domain using Fusion Applications Control:

1. From the navigation pane, expand the farm, then **WebLogic Domain**.
2. Select the Oracle WebLogic Server domain.
3. From the **WebLogic Domain** menu, choose **Control**, then **Shut Down**.

3.3.4.2 Starting an Oracle WebLogic Server Domain for a Product Family

You start an Oracle WebLogic Server domain for a product family after applying a patch.

This section contains the following topics for starting an Oracle WebLogic Server domain for a product family:

- [Task 1, "Start the Oracle HTTP Server"](#)

- [Task 2, "Start the Product Family Domain"](#)

Task 1 Start the Oracle HTTP Server

To start the Oracle HTTP Server using `opmnctl`:

```
(UNIX) WT_CONFIG_HOME/bin/opmnctl startproc process-type=OHS
(Windows) WT_CONFIG_HOME\bin\opmnctl startproc process-type=OHS
```

To stop Oracle HTTP Server using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Web Tier** installation type.
2. Select the Oracle HTTP Server.
3. From the **Oracle HTTP Server** menu, choose **Control**, then **Start Up**.

Task 2 Start the Product Family Domain

When you start an Oracle WebLogic Server domain for a product family, you start the cluster of Managed Servers in the domain and all the applications.

To start a product family domain using Fusion Applications Control:

1. From the navigation pane, expand the farm, then **WebLogic Domain**.
2. Select the Oracle WebLogic Server domain.
3. From the **WebLogic Domain** menu, choose **Control**, then **Start Up**.

3.3.5 Starting and Stopping an Oracle WebLogic Server Cluster for a Configuration Change

If you modify the Oracle WebLogic Server configuration, you must stop and restart the cluster for the configuration changes to take effect.

This section contains the following topics:

- [Starting and Stopping an Oracle WebLogic Server Cluster Using WLST and Node Manager](#)
- [Starting and Stopping an Oracle WebLogic Server Cluster Using Fusion Applications Control](#)

3.3.5.1 Starting and Stopping an Oracle WebLogic Server Cluster Using WLST and Node Manager

See "Using Node Manager to Start Managed Servers in a WebLogic Domain or Cluster" section in *Oracle Fusion Middleware Oracle WebLogic Scripting Tool*.

3.3.5.2 Starting and Stopping an Oracle WebLogic Server Cluster Using Fusion Applications Control

To stop or restart a cluster using Fusion Applications Control:

1. Modify the Oracle WebLogic Server configuration as required.
2. From the navigation pane, expand the farm, **WebLogic Domain**, and then domain name.
3. Select the Oracle WebLogic Server cluster.
4. Expand the Oracle WebLogic Server cluster to show each target Managed Server.

- From the **WebLogic Cluster** menu, choose **Control**, then **Shut Down** or **Start Up**.

If you prefer to use the WebLogic Server Administration Console, see "Start Managed Servers in a cluster" and "Shutdown servers in a cluster" in the *Oracle Fusion Middleware Oracle WebLogic Server Administration Console Online Help*.

3.3.6 Starting and Stopping Specific Applications

Applications may not work correctly until all Managed Servers in all domains have been started.

This section contains the following topics:

- Starting and Stopping Java EE Applications Using WLST
- Starting and Stopping an Individual Application Instance for a Cluster Using Fusion Applications Control or Cloud Control
- Starting and Stopping All Application Instances for a Cluster Using Fusion Applications Control or Cloud Control

3.3.6.1 Starting and Stopping Java EE Applications Using WLST

To start or stop applications with the WebLogic Scripting Tool (WLST):

- At an operating system command prompt, navigate to the location of the WLST script:

```
(UNIX) FA_MW_HOME/oracle_common/common/bin
(Windows) FA_MW_HOME\oracle_common\common\bin
```

FA_MW_HOME is named *fusionapps* and located under the *APPLICATIONS_BASE*.

- Run the `wlst` script:

```
(UNIX) wlst.sh
(Windows) wlst.cmd
```

- Connect to Oracle WebLogic Server with the following WLST command:

```
connect([options])
```

See the "connect" section in the *Oracle Fusion Middleware WebLogic Scripting Tool Command Reference*.

- Use the following WLST commands:

```
startApplication(appName, [options])
stopApplication(appName, [options])
```

The application must be fully configured and available in the domain. The `startApplication` command returns a `WLSTProgress` object that you can access to check the status of the command. In the event of an error, the command returns a `WLSTException`. For more information about the `WLSTProgress` object, see "WLSTProgress Object" in the *Oracle Fusion Middleware Oracle WebLogic Scripting Tool*.

3.3.6.2 Starting and Stopping an Individual Application Instance for a Cluster Using Fusion Applications Control or Cloud Control

To start or stop an application using Fusion Applications Control or Cloud Control:

- Navigate to a product family home page:

- From Fusion Applications Control, from the navigation pane, select the product family.
- From Cloud Control:
 - a. From the **Targets** menu, choose **Fusion Applications**.
The Fusion Applications target home page displays.
 - b. In the table on the Fusion Applications target home page, click the appropriate **Product Family** target.
- 2. From the navigation pane, expand the product family, then **Fusion Applications**, and then the cluster application.
- 3. Select the application instance you want to stop.
- 4. From the **Fusion J2EE Application** menu, choose **Control**, then **Start Up** or **Shut Down**.

If you prefer to use the WebLogic Server Administration Console, see "Start applications and modules" and "Stop applications and modules" in the *Oracle Fusion Middleware Oracle WebLogic Server Administration Console Online Help*.

3.3.6.3 Starting and Stopping All Application Instances for a Cluster Using Fusion Applications Control or Cloud Control

To start or stop all the application instances in a cluster using Fusion Applications Control or Cloud Control

1. Navigate to a product family home page:
 - From Fusion Applications Control, from the navigation pane, select the product family.
 - From Cloud Control:
 - a. From the **Targets** menu, choose **Fusion Applications**.
The Fusion Applications target home page displays.
 - b. In the table on the Fusion Applications target home page, click the appropriate **Product Family** target.
2. From the navigation pane, expand the product family, then **Fusion Applications**, and then the cluster application.
3. From the **Fusion Cluster Application** menu, choose **Control**, then **Start Up** or **Shut Down**.

3.4 Viewing and Changing Ports for Components

Oracle Database and Oracle Fusion Middleware components and services use ports. Most numbers are assigned during installation. As an administrator, it is important to know the numbers used by these services, and to ensure that the same number is not used by two services on your host.

For some ports, you can specify a number assignment during installation.

[Table 3–3](#) lists the port types and provides information on where to find documentation related to changing ports. The procedures referenced in the Oracle Fusion Middleware guides describe using Fusion Middleware Control. These procedure also apply to Fusion Applications Control.

Table 3–3 Changing Ports

Port Type	Documentation for Changing Ports
Oracle Database	"Changing the Oracle Database Net Listener" section in the <i>Oracle Fusion Middleware Administrator's Guide</i>
Oracle Business Intelligence Suite	"Using Fusion Middleware Control to Scale System Components" section in the <i>Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition</i>
Oracle Essbase	<p>"Using Fusion Middleware Control to Scale System Components" section in the <i>Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition</i></p> <p>You may also need to modify the port range used by Essbase Applications. These are managed separately from the rest of the ports.</p> <p>To modify the port range:</p> <ol style="list-style-type: none"> At an operating system command prompt, navigate to the location of the <code>opmn.xml</code> file: <ul style="list-style-type: none"> (UNIX) <code>APPLICATIONS_ BASE/instance/BIInstance/config/OPMN/opmn</code> (Windows) <code>APPLICATIONS_ BASE\Instance\BIInstance\config\OPMN\opmn</code> <p>This is the port range used by Essbase Applications. These are allocated by the Essbase server as required. The Essbase server will use the ports, even if they are explicitly allocated elsewhere to another component. Therefore, it is important to not use the same port range as the rest of the Oracle Business Intelligence domain port range.</p> Modify the port range: <pre data-bbox="691 1075 1295 1102"><port id="essbase-port-range" range="9500-9599" /></pre> Ensure that the OPMN server re-reads its configuration after changing the contents of <code>opmn.xml</code>: <ul style="list-style-type: none"> (UNIX) <code>APPLICATIONS_ BASE/instance/BIInstance/bin/opmnctl reload</code> (Windows) <code>APPLICATIONS_ BASE\instance\BIInstance\bin\opmnctl reload</code> Restart the Essbase server: <ul style="list-style-type: none"> (UNIX) <code>APPLICATIONS_ BASE/instance/BIInstance/bin/opmnctl restartproc process-type=Essbase</code> (Windows) <code>APPLICATIONS_ BASE\instance\BIInstance\bin\opmnctl restartproc process-type=Essbase</code>
Oracle HTTP Server	"Manage Ports" section in the <i>Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server</i>
Oracle Enterprise Content Management Suite	"Modifying Server Configuration Parameters for Oracle UCM Content Server" section in the <i>Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory</i> for information about changing the for Oracle WebCenter Content Content Server

Table 3–3 (Cont.) Changing Ports

Port Type	Documentation for Changing Ports
Oracle Internet Directory	"Configuring Server Properties" section or the "Setting System Configuration Attributes by Using ldapmodify" section in the <i>Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory</i>
Oracle Virtual Directory	"Configuring Oracle Virtual Directory to Listen on Privileged Ports" section in the <i>Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory</i>
Node Manager	"Overview of Node Manager Configuration" section in the <i>Oracle Fusion Middleware Node Manager Administrator's Guide for Oracle WebLogic Server</i>

3.5 Changing Passwords

Before you begin configuration, change the passwords for the various accounts to secure passwords. [Table 3–4](#) describes the accounts and provides information on where to find related documentation for each account. The procedures referenced in the Oracle Fusion Middleware guides describe using Fusion Middleware Control. These procedure also apply to Fusion Applications Control.

Table 3–4 Changing Passwords

Password Account	Description	Documentation for Changing the Passwords
Oracle Fusion Middleware administration for the Oracle WebLogic Server domain	During the Oracle Fusion Applications installation, you must provide a password for the Oracle Fusion Middleware administration user. By default, this administrator takes the Super User value you specified on the Identity Management page when creating the provisioning plan. The super user is described in the next row of this table.	Section 4.8.5.2
Super User	<p>The super user has the following privileges:</p> <ul style="list-style-type: none"> ■ Administrative privileges for all Oracle WebLogic Server domains and all middleware. ■ Functional setup privileges for all Oracle Fusion applications. ■ Administrative privileges to Oracle Fusion applications. These do not include transactional privileges. You use this username and password to access the Fusion Applications Control and the Oracle WebLogic Server Administration Console. <p>You can change the credentials using Fusion Applications Control or Oracle Platform Security Services (OPSS) scripts.</p> <p>This user is specified on the Identity and Policy Management page of the Provisioning Wizard during installation. See "Identity and Policy Management Configuration" section in the <i>Oracle Fusion Applications Installation Guide</i> for information about this page. Installation establishes the same username and password for all the domains.</p>	<p>Section 4.8.5.2 to change the administrative password</p> <p>"Managing the Credential Store" section in the <i>Oracle Fusion Middleware Application Security Guide</i></p>
Oracle Metadata Repository schema	Oracle Metadata Services (MDS) repository contains metadata for the Oracle Fusion Applications and some Oracle Fusion Middleware component applications. The schema passwords are stored in the Oracle database.	"Changing Metadata Repository Schema Passwords" section in the <i>Oracle Fusion Middleware Application Security Guide</i>

Table 3–4 (Cont.) Changing Passwords

Password Account	Description	Documentation for Changing the Passwords
App IDs	Oracle Fusion Applications must rely on a type of credential known as the App ID. Each application has its own App ID which is initially provisioned for the application.	Section 4.8.5.1 <i>Oracle Fusion Applications Security Hardening Guide</i>
Node Manager	The Node Manager account authenticates the connection between a client (for example, the Administration Server) and Node Manager. In an Oracle Fusion Applications installation, this user is specified on the Installation Location page of the Provisioning Wizard. See the "Installation Location" section in the <i>Oracle Fusion Applications Installation Guide</i> for information about using this page.	"Specify Node Manager Username and Password" section in the <i>Oracle Fusion Middleware Node Manager Administrator's Guide for Oracle WebLogic Server</i>
BISystemUser	The BISystemUser account provides access to the Oracle Business Intelligence system components.	"Default Users and Passwords" section in the <i>Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition</i>

3.6 Managing the Oracle Database

To manage your Oracle database:

1. Start the Oracle database instance, if not already started. See [Section 3.3.2](#).
2. Review database initialization parameters. Modify initialization parameters as needed. See [Section 3.6.1](#).
3. Review your database storage structures: tablespaces and data files, online redo log files, and control files. Create or modify storage structures as needed. See [Section 3.6.2](#).
4. Review memory allocation and adjust as needed. See [Section 3.6.3](#).
5. Review, unlock, and reset passwords for predefined database users as needed. Create new users, and assign privileges and roles to them as needed. See [Section 3.6.4](#).
6. Create or review the backup strategy for the database and back up the database. See [Chapter 15](#).

3.6.1 Viewing and Modifying Initialization Parameters

Managing an Oracle instance includes configuring parameters that affect the basic operation of the Oracle instance. These parameters are called initialization parameters. The Oracle instance reads initialization parameters from a file at startup.

After being read from a file, initialization parameters are retained in memory, where the values for many of them can be changed dynamically. There are two types of parameter files. The type of file used to start the instance determines if dynamic initialization parameter changes persist across database shutdown and startup. The parameter file types are:

- Server parameter file

The **server parameter file** is a binary file that can be written to and read by the database. It *must not* be edited manually. It is stored on the host system on which Oracle Database is running. Changes are made when you use Database Control to modify one or more initialization parameters, or when Oracle Database itself makes changes for self-tuning purposes. Any changes to it persist across database shutdown and startup operations.

- Text initialization parameter file

A **text initialization parameter file** is a text file that can be read by the Oracle instance, but it is not written to by the instance. You can change a text initialization parameter file with a text editor, but changes do not take effect until you restart the Oracle instance. When you start the instance with this type of file, you can still change many initialization parameters dynamically with Database Control, but only for the current instance. Unless you also edit the text initialization parameter file and make the same change, the change is lost when you restart the database instance.

As the number of database users increases and the workload increases, you might have to alter some initialization parameters. You can make these changes using the Initialization Parameter page in Oracle Enterprise Manager Database Control, accessible from the **Database Configuration** section of the **Server** tab.

Table 3–5 lists specific initialization parameters that you should set and their recommended values.

Table 3–5 Initialization Parameters

Parameter	Recommended Values
CONTROL_MANAGEMENT_PACK_ACCESS	Set this parameter to <code>DIAGNOSTIC+TUNING</code> (default) or <code>DIAGNOSTIC</code> to enable Automatic Database Diagnostic Monitor (ADDM).
STATISTICS_LEVEL	Set this parameter to <code>TYPICAL</code> (default) to enable the automatic performance tuning features of Oracle Database, including Automatic Workload Repository (AWR) and ADDM.

For information about viewing and modifying initialization parameters, see the "Viewing and Modifying Initialization Parameters" section in the *Oracle Database 2 Day DBA*.

3.6.2 Managing Database Storage Structures

Oracle Database is made up of physical and logical structures. Physical structures can be seen and operated on from the operating system, such as the physical files that store data on a disk.

Logical structures are created and recognized by Oracle Database and are not known to the operating system. The primary logical structure in a database, a tablespace, contains physical files. The applications developer or administrator may be aware of the logical structure, but may not be aware of the physical structure. The database administrator (DBA), on the other hand, must understand the relationship between the physical and logical structures of a database.

Oracle Database can automate much of the management of its structure. To view a database storage structure using Oracle Enterprise Manager Database Control, go to the **Storage** section of the **Server** tab, where you can access the following storage options:

- Control files

- Tablespaces
- Temporary tablespace groups
- Datafiles
- Rollback segments
- Redo log groups
- Archive logs
- Disk groups
- Other storage structures

For more information about managing database storage structures, see the "Managing Database Storage Structures" chapter in the *Oracle Database 2 Day DBA*.

3.6.3 Managing Memory

Memory management involves maintaining optimal sizes for the Oracle instance memory structures as demands on the database change. The memory that must be managed are the System Global Area (SGA) memory and the instance Program Global Area (PGA) memory. The instance PGA memory is the collection of memory allocations for all individual PGAs.

Oracle Database can manage the SGA memory and instance PGA memory automatically. You designate only the total memory size to be used by the instance, and Oracle Database dynamically exchanges memory between the SGA and the instance PGA as needed to meet processing demands. This capability is referred to as **automatic memory management**. In this memory management mode, the database also dynamically tunes the sizes of the individual SGA components and the sizes of the individual PGAs.

To have more direct control over the sizes of the SGA and instance PGA, use the Memory Advisors page of Oracle Enterprise Manager Database Control to disable automatic memory management and enable automatic shared memory management. With **automatic shared memory management**, you set target and maximum sizes for the SGA. Oracle Database then tunes the total size of the SGA to your designated target, and dynamically tunes the sizes of all SGA components. In this memory management mode, you also implicitly enable automatic PGA memory management. With **automatic PGA memory management**, you set a target size for the instance PGA. The database then tunes the size of the instance PGA to your target, and dynamically tunes the sizes of individual PGAs.

If you want complete control of individual SGA component sizes, then use the Memory Advisors page of Oracle Enterprise Manager Database Control to disable both automatic memory management and automatic shared memory management. This is called **manual shared memory management**. In this mode, you set the sizes of several individual SGA components, thereby determining the overall SGA size. You then manually tune these individual SGA components on an ongoing basis. Manual shared memory management mode is intended for experienced DBAs only. Note that in this mode, automatic PGA memory management remains enabled.

To manage memory, use the Memory Advisor page in Database Control, accessible from the **Database Configuration** section of the **Server** tab.

For more information about memory management, see the "Managing Memory" section in the *Oracle Database 2 Day DBA*.

3.6.4 Administering User Accounts

For users to access your database, you must create user accounts and grant appropriate database access privileges to those accounts. A user account is identified by a user name and defines the attributes of the user, including the following:

- Authentication method
- Password for database authentication
- Default tablespaces for permanent and temporary data storage
- Tablespace quotas
- Account status (locked or unlocked)
- Password status (expired or not)

When you create a user account, you must not only assign a user name, a password, and default tablespaces for the account, but you must also do the following:

- Grant the appropriate system privileges, object privileges, and roles to the account.
- If the user will be creating database objects, then give the user account a space usage quota on each tablespace in which the objects will be created.

In addition, you may want to create user accounts that are used by applications only, such as Fusion Applications. Users do not log in with these accounts; instead, applications use these accounts to connect to the database, and users log in to the applications. This type of user account avoids giving application users the ability to log in to the database directly, where they could unintentionally cause damage.

To administer user accounts using Oracle Enterprise Manager Database Control, go to the **Security** section of the **Server** tab, where you can access users and roles.

For more information about administering user accounts, see the "Administering User Accounts and Security" chapter in the *Oracle Database 2 Day DBA*.

3.7 Patching

It is necessary to perform various maintenance actions on your applications, their middleware dependencies, and their database components. Maintenance actions include fixing issues that affect the way the applications perform, adding new functionality and features, updating to a higher maintenance level, or providing interoperability to new technology stacks. Patches may be required for maintenance of middleware artifacts, database artifacts, or both. [Table 3–6](#) describes the types of patching and provides information on where to find related documentation.

Table 3–6 Patching

Type of Patching	Description	Documentation for Patching
Oracle Fusion Applications	The Oracle Fusion Applications Patching Framework provides the tools needed to sup updates to Oracle Fusion Applications software between major or patch set releases. These tools manage the processes for applying individual patches, patch sets, and release update packs.	<i>Oracle Fusion Applications Patching Guide</i>
Oracle Fusion Middleware	OPatch patches any Oracle Fusion Middleware component, except Oracle WebLogic Server. For Oracle WebLogic Server, use Smart Update.	<i>Oracle Fusion Middleware Patching Guide</i> for information about patching middleware artifacts
Oracle Database	Oracle Universal Installer and OPatch manage the patching for Oracle database and third-party software.	<i>Oracle Universal Installer and OPatch User's Guide for Windows and UNIX</i>

3.8 Managing Oracle Fusion Applications-Specific Labels in the Oracle Metadata Repository

The Oracle Metadata Repository (MDS Repository) contains metadata for Oracle Fusion Middleware components. It can also contain metadata about the configuration of Oracle Fusion Middleware and metadata for your applications. For general instructions on managing the metadata for Oracle Fusion Middleware components in the Oracle Metadata Repository, see the "Managing the Metadata Repository" chapter in the *Oracle Fusion Middleware Administrator's Guide*.

As a part of MDS Repository management, you work with metadata labels. You select a version for the repository to select a particular version of each object from a metadata repository partition. Also, if you take actions that affect the MDS Repository, such as deploying, patching, or doing some customizations, you can potentially put the repository into a broken state. When this happens, you can rollback to a previous label. For example, when you roll back to a label, all objects within that repository move to that state. To manage labels, note the following prefixes for Oracle Fusion Applications in the MDS Repository:

- `postDeployLabel_` for labels created by MDS Repository at deployment time
- `Creation_` for labels created by Oracle Fusion Middleware Extensions for Applications (Applications Core) in the MDS Repository
- `PreMerge_` for labels created by Applications Core in the MDS Repository during sandbox creation.
- `composer_PostMerge_` for labels created by Applications Core in the MDS Repository during a sandbox publish operation
- `mds_` for labels created during a Flexfield changes applied during a patch
- `soa_` for labels created during SOA deployments
- `composer_` for labels created by Oracle Composer when using the **Save and Label** button. For more information about **Save and Label** button, see "Create Labels On Saving Application Customizations" in the *Oracle Fusion Middleware Developer's Guide for Oracle WebCenter Portal*.

For more information about managing labels, see the "Managing Metadata Labels in the MDS Repository" section in the *Oracle Fusion Middleware Administrator's Guide*.

3.9 Modifying Oracle Application Development Framework Connections in Oracle Fusion Applications

A connection configuration (`connections.xml`) contains information that a client application uses to identify the Oracle Application Development Framework (Oracle ADF) application module's deployment scenario. For more information about modifying the configuration of a single application, see the following sections in the *Oracle Fusion Middleware Administrator's Guide for Oracle Application Development Framework*:

- How to Modify Connection Configurations
- Configuring Application Properties Using the MBean Browser

In an Oracle Fusion Applications environment, you may need to change the internal and external settings for several applications. For example, if there was an outage of the Oracle Fusion Customer Relationship Management (CRM) applications due to a lost Oracle HTTP Server host, then you would need update all the applications that depend on the CRM-provided web services. If all the applications in an Oracle WebLogic Server domain use the same `connections.xml`, which is the default behavior after provisioning, you can make this edit in one place using Fusion Applications Control.

To modify the connections for all the applications in an Oracle WebLogic Server domain using Fusion Applications Control:

1. From the navigation pane, expand the farm and then **WebLogic Domain**.
2. Click the domain.

The Oracle WebLogic Server Domain home page displays.

3. From the **WebLogic Domain** menu, choose **ADF Domain Configuration**.

The ADF Common Properties page displays.

4. In the **Property Sets** table, click the property set to modify, and in the **Details for Property Set** table, modify the property values.

These property sets and properties reside in the `adf-domain-config.xml` file, and are used in the ELs in the `connections.xml` file.

5. Click **Save**.

Securing Oracle Fusion Applications

This chapter explains the security features available to all Oracle Fusion applications. It explains the enterprise identity store, identity provisioning, authorization policies, roles, audit trail, SSL configuration, data masking, securing Web services, and customizing security.

This chapter explains the security features available to all Oracle Fusion applications, and it contains the following topics:

- [Introduction to Security](#)
- [About the Enterprise Identity Store](#)
- [Provisioning Identities](#)
- [Managing Authorization Policies](#)
- [Configuring Roles](#)
- [Configuring Audit Trail](#)
- [Configuring SSL for Oracle Fusion Applications](#)
- [Managing Wallets, Keystores, Credentials, and Certificates](#)
- [Data Masking](#)
- [Securing Web Services](#)
- [Securing Oracle Fusion Middleware Products](#)
- [Extracting Data from an LDAP- Based Store to a File](#)
- [Customizing Security from Installation to Deployment](#)

The high-level information presented in this chapter includes links to other documents where the topic is explained in detail.

For additional information about application security, see the following documents:

- *Oracle Fusion Applications Security Guide*
- *Oracle Fusion Applications Security Hardening Guide*

For a detailed list of administrative tasks and pointers to further documentation, see *Oracle Fusion Applications Administrator and Implementor Roadmap*.

4.1 Introduction to Security

Oracle Fusion Applications use the services of the Oracle Platform Security Services (OPSS) to secure applications.

OPSS is a security platform that provides enterprise product development teams, systems integrators, and independent software vendors with a standards-based enterprise-grade security framework for Java SE and Java EE applications. Using OPSS, Oracle Fusion Applications benefit from the same, uniform security, identity management, and audit services across the enterprise.

The intended audience for this chapter are application security administrators and system security administrators.

Oracle Fusion Applications provisioning sets up the security infrastructure, including:

- the identity store
- authorization policies
- enterprise roles
- SSL wiring and its support structure including keystores and certificates
- data masking
- setting protected URIs for the infrastructure to work as expected

For instructions about protected URIs and configuring Oracle ADF applications with Oracle Access Manager SSO, see the "Integration with Oracle ADF Applications" section in the *Oracle Fusion Middleware Administrator's Guide for Oracle Access Manager with Oracle Security Token Service*.

For details about security tasks, including creating and managing users and roles, function and data security, audit, and compliance tasks, see the *Oracle Fusion Applications Security Guide*.

4.2 About the Enterprise Identity Store

Oracle Fusion applications run within a container in the Oracle WebLogic Server. This container handles authentication automatically for the application running in it by intercepting all requests to the application and ensuring that users are properly authenticated and the security context is propagated, as appropriate, before the request can proceed forward.

Note: The Subject creation is automatic, but the security context propagation requires application configuration.

Fusion Applications use LDAP-based authenticators; Fusion Application identity provisioning sets up and wires WebLogic domains with the appropriate authenticators during the Fusion Application installation.

Important: Any LDAP-based authenticator, other than the DefaultAuthenticator, requires that the flag `UseRetrievedUserNameAsPrincipal` be set. During installation, this flag is automatically set in the DefaultAuthenticator.

For details about bootstrap identity provisioning, such as super administrators for Fusion pillars, see [Section 4.3](#).

4.2.1 Supported LDAP Identity Providers

Fusion Applications support the following LDAP identity store types:

- Oracle Internet Directory 11g
- Active Directory 2008

4.2.2 Configuring the Identity Store

Multiple LDAP authenticators can be configured in a given context. For the algorithm that selects the identity store to initialize from a stack of authenticators, see the "Configuring the Identity Store Service" section in the *Oracle Fusion Middleware Application Security Guide*.

Note: The Oracle WebLogic Server Administration Console is the recommended tool to configure authenticators, but this configuration can also be alternatively carried out with WLST commands. For the list of all available WLST commands, see *Oracle Fusion Middleware WebLogic Scripting Tool Command Reference*.

The specification and configuration of LDAP authenticators is carried out with the Oracle WebLogic Administration Console. For details, see the "Configuring Authentication Providers" chapter in the *Oracle Fusion Middleware Securing Oracle WebLogic Server*.

It is important to keep the username attribute on the authenticator synchronized with the corresponding identity store property. For details, see note at the end of the table of identity store properties in the "LDAP Identity Store Properties" section in the *Oracle Fusion Middleware Application Security Guide*.

It is also important that the following two time intervals be equal:

- The number of seconds that cached entries stay in the cache. This value is controlled by the WebLogic authenticator parameter `Group Hierarchy Cache TTL`, which by default is 60 seconds.
- The number of seconds after which group membership changes are in effect. This value is controlled by the system property `jps.subject.cache.ttl`, which by default is 60 seconds.

If the `Group Hierarchy Cache TTL` value is changed, then that new value must also be set with the system property `jps.subject.cache.ttl`. For example, if the value of `Group Hierarchy Cache TTL` is changed to 55,000 (milliseconds), then `jps.subject.cache.ttl` must be reset as follows:

```
-Djps.subject.cache.ttl 55000
```

4.3 Provisioning Identities

Provisioning as a whole encompasses all the operations required to install, configure, and deploy applications product offerings. Identity provisioning is a subset of this process which populates the users and groups needed for deployment and ongoing administration.

This section contains the following topics:

- [Identity Provisioning Concepts](#)
- [WebLogic Authenticators and the Primary Identity Store](#)

- [Provisioning Steps](#)
- [Best Practices for the Administrator Groups](#)
- [Managing Identities after Deployment](#)

4.3.1 Identity Provisioning Concepts

During the identity provisioning stage of installation, Oracle Fusion Applications require the existence of certain users with specific privileges. These administrative users reside in a secure, central repository called the identity store. This section explains the phases of identity provisioning, what users or groups are needed for provisioning, and the users and groups that exist at the end of the process.

This section contains the following topics:

- [Administrators For Fusion Applications](#)
- [Two Types of Users During Provisioning](#)

4.3.1.1 Administrators For Fusion Applications

The application provisioning process bootstraps the provisioned environment with two administrator groups for each application family.

These two administrator groups are:

- A system administrator
A directory group representing the WebLogic Server domain administrators for all the domains.
- An application administrator
A directory group with an assigned enterprise role reflecting all the application roles and delegation privileges for all the applications in a given family.

The purpose of creating these "Super Administrators" during provisioning is to enable ongoing administration and/or delegation privileges.

The above process facilitates separation of duties between system administration and application administration responsibilities, but you are free to assign the same user to both hierarchies ("system admin" and "application admin").

[Table 4–1](#) shows the groups that are created for each application family:

Table 4–1 Provisioned Administrator Groups

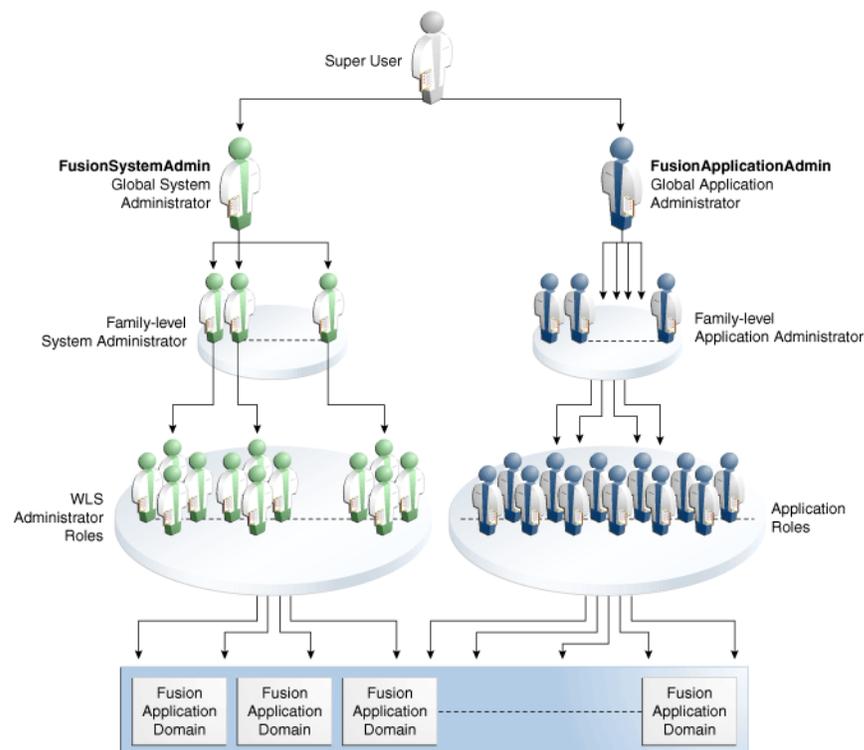
Product Family/Product	System Administrator Group	Application Administrator Group
Oracle Fusion Supply Chain Management	FSCMSysAdmin	FSCMAppAdmin
Oracle Fusion Customer Relationship Management	CRMSysAdmin	CRMAppAdmin
Oracle Fusion Human Capital Management	HCMSysAdmin	HCMAppAdmin
Oracle Fusion Financials	FINSysAdmin	FINAppAdmin

Table 4–1 (Cont.) Provisioned Administrator Groups

Product Family/Product	System Administrator Group	Application Administrator Group
Oracle Fusion Procurement	PRCSysAdmin	PRCAppAdmin
Oracle Fusion Project	PRJSysAdmin	PRJAppAdmin
Oracle Fusion Incentive Compensation	OICSysAdmin	OICAppAdmin

In addition a single user, known as the super user, is set up to belong to all the administrator groups. That user becomes the administrator for all middleware and the application administrator for all product families.

Figure 4–1 shows the relationship between these groups.

Figure 4–1 Super-User and Administrators

4.3.1.2 Two Types of Users During Provisioning

It is important to distinguish between the two types of super-administrators that exist in the provisioning process.

- Pre-seeded bootstrap user
- Designated super-user

In the context of the pre-seeded user, provisioning employs an identity known as the App ID that is required to bootstrap the WebLogic domains. The pre-configuration

phase of provisioning automatically generates the credential needed for this App ID user.

In the context of the designated super user, during the interview phase of provisioning, you are asked to specify the user ID of the designated "real" user who will be set up as the Middleware Administrator and Functional Setup Manager.

For example, if you want a "real" user such as "cn=john.doe,cn=users,cn=acme,cn=com" to be the super user, provide "john.doe" as the user ID during provisioning. This user will be set up as the super user in the identity store.

Note: It should be emphasized that the identity seed data that is used in the LDIF file to configure the WebLogic domains does not use real user DNs.

4.3.2 WebLogic Authenticators and the Primary Identity Store

When Oracle WebLogic Server is installed, the default authenticator is based on an embedded LDAP store.

As part of Oracle Fusion Applications provisioning, the default authenticator based on the embedded LDAP is deleted. Upon completion of Oracle Fusion Applications provisioning, the primary and only identity store will be your external LDAP store. The bootstrap identity used to configure the domains during the provisioning process will be pre-seeded in the external LDAP through the LDIF file, as explained in [Section 4.3.1.2](#).

4.3.3 Provisioning Steps

The identity provisioning process consists of distinct phases.

In the interview phase, Provisioning Wizard collects the following information:

- The DN of the user designated as the super user. This user must already exist in the identity store.
- Whether the system administrators group exists or must be created.
- If the group exists, the DN of the group.
- The LDAP authenticator, either Oracle Internet Directory (OIDAuthenticator) or Oracle Virtual Directory (OVDAAuthenticator) that will serve as the LDAP identity store.

The next step of the process verifies that the designated super-user exists in the identity store.

The system administrator group is created if needed, and the super user is made a member of the group.

Next, the application domains are created, the LDAP authenticator is enabled, and the WebLogic domain is started up.

Following configuration, the system administrator groups are assigned the appropriate family-level enterprise roles.

At the end of this process, the super user has:

- Administrator privileges for all WebLogic domains and all middleware.
- Function setup privileges for all Oracle Fusion applications.

- Administration privileges to Oracle Fusion Applications. These do not include transactional privileges.

For more information about identity provisioning and using the interview wizard to create a provisioning plan, see the *Oracle Fusion Applications Installation Guide*.

4.3.4 Best Practices for the Administrator Groups

While there are logical sets of "super" administrative groups (a set of two per application family, consisting of the super-user administrator and the application administrator) you can choose to distribute these functions among fewer individuals. The recommended best practice is to carefully plan the separation of duties, taking into account the real-world operational needs of your site.

4.3.5 Managing Identities after Deployment

Oracle Identity Manager is the best-in-class user provisioning and administration component in Oracle Fusion Middleware.

Oracle Identity Manager automates the process of adding, updating, and deleting user accounts from applications and directories.

Identities can be created and managed when user records are created through activities such as employee hiring and creation of contacts through CRM.

In addition, identities can also be created and managed administratively through the Oracle Identity Manager Administrative Console.

For more information about provisioning and managing identities, see:

- *Oracle Fusion Middleware User's Guide for Oracle Identity Manager*
- *Oracle Fusion Middleware Administrator's Guide for Oracle Identity Manager*

4.4 Managing Authorization Policies

Authorization is the most sensitive and application-specific security concept. At its very core, authorization protects access to application resources through the enforcement of policies, which are stored in the domain policy store. Authorization determines what types of actions, tasks, or services a user can access.

In most cases, the definition of application policies begins during the design of the application. This definition includes identifying application privileges and application roles, the hierarchical relationships between application roles, and categorizing them into products and Java EE applications.

The policy model is based on a number of logical entities, such as resource types, resource instances, entitlements (also known as permission sets), application roles, and enterprise roles. For details about these entities and the logical model of a policy, see the "Terminology" section in the *Oracle Fusion Middleware Application Security Guide*.

This section includes the following topics:

- [Managing Oracle Fusion Application Policies](#)
- [Managing System Policies](#)
- [Reconciling GUIDs](#)
- [Managing Data Security](#)

4.4.1 Managing Oracle Fusion Application Policies

An Oracle Fusion application policy is either a functional policy or a data security policy. Both these policies define who can do what on a resource.

A data security policy includes a condition, while a functional policy does not. The condition identifies a row or a set of rows in a business object, and the privileges the data security grants are for *only* the data that meets the condition. A functional policy, instead, assigns permissions to resources or code artifacts (such as task flows, pages, Java methods, or UI components) and grants a specific set of actions on each resource.

Data security policies are stored in the transactional database; functional policies are stored in the domain policy store.

Oracle Authorization Policy Manager is the recommended tool to administer application policies once the application has been deployed. This graphical interface tool allows application security administrators to provision, search, and modify application functional and data security policies. Using this tool they can, for example, remove a resource from an entitlement or change the actions granted to a resource in an entitlement.

For details about the most frequently uses of Oracle Authorization Policy Manager, see the "Managing Policies and Policy Objects" section in the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)* where the following typical administrative tasks are described:

- Managing application roles
- Managing application resource types
- Managing application resources
- Managing application entitlements
- Creating and modifying an application policy
- Viewing the enterprise role hierarchy
- Managing the application role hierarchy
- Mapping application roles to an enterprise role
- Mapping enterprise roles to an application role

For details about configuring application roles, see [Section 4.5.1](#).

4.4.2 Managing System Policies

Oracle Enterprise Manager Fusion Middleware Control is the recommended tool to administer system policies. These are policies that pertain to the whole domain, as opposed to application policies, which pertain a particular application.

A principal policy is a system policy that grants permissions to a list of users or enterprise groups. A codebase policy is a system policy that grants permissions to a piece of code or a URL (typically represented by an EAR or a JAR file); for example, an application using the Credential Store Framework requires an appropriate codebase policy.

Fusion Middleware Control allows the creation and modification of both these types of system policies. For details about the procedure to follow, see the "Managing System Policies" section in the *Oracle Fusion Middleware Application Security Guide*.

An alternative way to administer policies (both system and application policies), although not as convenient but occasionally useful, is using WLST commands. For a

complete list of security-related commands, see the "OPSS Scripts" appendix in the *Oracle Fusion Middleware Application Security Guide*.

4.4.3 Reconciling GUIDs

The recipient of a grant can be either an application role or an enterprise role. In Oracle Fusion Data Security and policy grants this recipient is identified by a GUID, and it is crucial for the security system to work as expected that these GUIDs be consistent. Since GUIDs are *not* preserved by migration, the GUIDs in the Oracle Fusion Data Security policies and in the policy store policies must be reconciled when, for instance, migrating to a staging or a production environment.

The java utility program `DSDataMigrator` reconciles GUIDs by modifying GUIDs in Oracle Fusion Data Security so that the GUIDs of the role entries in the identity and policy stores are consistent with those in Oracle Fusion Data Security.

`DSDataMigrator` needs not be run when:

- The user interface is used to create new policies or to modify existing ones.
- The system is provisioned for the first time (GUID reconciliation happens automatically in this case).
- A patch containing Grants Seed Data is applied to the environment (GUID reconciliation happens automatically in this case).

`DSDataMigrator` must be run when:

- A policy stripe was dropped or recreated.
- A role was dropped or recreated.
- An SQL script was used to upload grants data.
- Patching the infrastructure fails, such as, when the LDAP server went down.
- The LDAP server was swapped out.
- `FndGrantsSD.xml` data was loaded to the database using SDF programs directly bypassing the patching infrastructure.
- The file `jps-config-jse.xml` contains invalid or incorrect configurations.
- Security data was migrated to a staging or to a production environment.
- Enterprise roles in the identity store were dropped and imported again using a migration tool.

This section includes the following topics:

- [Prerequisites to Running DSDataMigrator](#)
- [DSDataMigrator Syntax](#)

4.4.3.1 Prerequisites to Running DSDataMigrator

Before an administrator runs this command, it is assumed that:

- The Fusion Application has been installed (so that Oracle Fusion Data Security has been loaded).
- The XML policies (`jazn-data.xml`) have been migrated to an Oracle Internet Directory server.
- The `FND_GRANTS` table has been backed up, as illustrated in the following invocation:

```
>sqlplus sys as sysdba
create table FUSION.FND_GRANTS_OLD as select * from FUSION.FND_GRANTS;
```

- The classpath in the shell where the command is to be run contains the following JAR files:
 - MW_HOME/atgpf/atgpf/modules/oracle.applcore.model_11.1.1/Common-Model.jar.
 - MW_HOME/atgpf/atgpf/modules/oracle.applcore.model_11.1.1/DataSecurity-Model.jar.
 - MW_HOME/oracle_common/modules/oracle.adf.model_11.1.1/adfm.jar.
 - MW_HOME/oracle_common/modules/oracle.adf.share_11.1.1/adf-share-support.jar.
 - MW_HOME/oracle_common/modules/oracle.adf.share.ca_11.1.1/adf-share-ca.jar.
 - MW_HOME/oracle_common/modules/oracle.adf.share.ca_11.1.1/adf-share-base.jar.
 - MW_HOME/oracle_common/modules/oracle.adf.share_11.1.1/jsp-el-api.jar.
 - MW_HOME/oracle_common/modules/oracle.adf.businesseditor_11.1.1/adf-businesseditor.jar.
 - MW_HOME/oracle_common/modules/oracle.adf.share_11.1.1/adfflogginghandler.jar.
 - MW_HOME/oracle_common/modules/oracle.jps_11.1.1/jps-manifest.jar.
 - MW_HOME/modules/javax.jsp_1.2.0.0_2-1.jar.
 - MW_HOME/oracle_common/modules/oracle.mds_11.1.1/mdsrt.jar.
 - MW_HOME/oracle_common/modules/oracle.javatools_11.1.1/resourcebundle.jar.
 - MW_HOME/oracle_common/modules/oracle.javatools_11.1.1/javatools-nodeps.jar.
 - MW_HOME/wlserver_10.3/server/ext/jdbc/oracle/11g/ojdbc5.jar.
- Optionally, the identity store has been seeded.

Notes: The files `Common-Model.jar` and `DataSecurity-Model.jar` are expected to be in the paths indicated above, but, depending on the environment, they could be installed in some other location. To find out the location of those files in your environment, invoke the following commands at the top of the `MW_HOME` directory:

```
>find . -name "Common-Model.jar"
>find . -name "DataSecurity-Model.jar"
```

`DSDDataMigrator` is executed automatically when (and only when) a patch containing the data security grants seed data file `FndGrantsSD.xml` is applied to the environment. In particular, the script is not executed automatically if the seed data file is applied manually.

4.4.3.2 DSDataMigrator Syntax

DSDataMigrator is located in the directory `oracle.apps.fnd.applcore.dataSecurity.util`, and it has the following syntax (arguments are written in separate lines for the sake of clarity only):

```
java -classpath $CLASSPATH
-Doracle.security.jps.config=<path to the jps-config-jse.xml file>
-DFND_DS_GUID_RECON_LOG_DIR=<path to the log output directory>
DSDataMigrator -dsdburl <dsdbURL>
                -dsdbuser <dsdbUser>
                -silentMode <true_or_false>
                -forceProcessAllRows <true_or_false>
                -policyStripe <FA policy stripe name>
                -idStoreOnly <true_or_false>
                -validationMode <true_or_false>
                -multiTenantAware <true_or_false>
                -enterprise_id <enterpriseId>
```

When run and before processing security data, the command prompts the administrator for the database password.

The meaning of the arguments is as follows:

- `oracle.security.jps.config` specifies the location of the configuration file `jps-config-jse.xml` where the policy store and identity store are configured. This file must include the appropriate bootstrap credentials to access the policy store. The following fragment of a configuration file illustrates the specification of credentials for a policy store instance:

```
<serviceInstance provider="ldap.policystore.provider" name="policystore.ldap">
  <property value="OID" name="policystore.type" />
  <property value="bootstrap_123456" name="bootstrap.security.principal.key" />
  <property value="cn=st2_d8b3" name="oracle.security.jps.farm.name" />
  <property value="cn=FusionAppsPolicies"
name="oracle.security.jps.ldap.root.name" />
  <property value="ldap://adc2110301.us.oracle.com:33060" name="ldap.url" />
</serviceInstance>

<serviceInstance location="./bootstrap" provider="credstoressp"
name="bootstrap.cred">
  <property value="./bootstrap" name="location" />
</serviceInstance>
```

Typically, the file `jps-config-se.xml` is located in the directory `<WLS Domains>/<server name>/CommonDomain/config/fmwconfig`. If this file is moved to a different location, the bootstrap directory (which contains the file `cwallet.sso`) must also be moved to that new location.

- `dsdburl` specifies the URL of the data base where Oracle Fusion Data Security is stored.
- `dsdbuser` specifies the name of the user that can access the data base.
- `silentMode` specifies whether the command should raise exceptions when an entry is not found in the OID server. Set to `TRUE` to prevent raising these kind of exceptions; otherwise, set to `FALSE`. Default value: `FALSE`.
- `forceProcessAllRows` specifies whether all rows in the `FND_GRANTS` table should be processed. Set to `TRUE` to process all rows in that table; otherwise, set to `FALSE`. The default behavior is `FALSE` and processes just those rows with the `compile_` flag is set to `Y`.

- `policyStripe` specifies the name of the Oracle Fusion application stripe in the policy store. Typical values are `fscm`, `crm`, and `hcm`.
- `idStoreOnly` specifies whether only data security grants granting to enterprise roles should be processed. Set to `TRUE` to process only grants granting to enterprise roles; otherwise, set to `FALSE` to process all grants. When set to `TRUE`, the value of the argument `policyStripe` is ignored. Default value: `FALSE`.
- `validationMode` specifies to run the script in read only mode (validation mode). Set to `TRUE` to run the script *without* updating the database data. Set to `FALSE` to run the script and update database data. Default value: `FALSE`. A validation run of the script is useful to find out if any reconciliation is needed between the database and the LDAP store.
- `multiTenantAware` specifies whether the environment is a multi-tenant environment; applies only to multi-tenant environments; set to `TRUE` to indicate a multi-tenant environment. Default value: `FALSE`. When set to `TRUE`, the argument `enterprise_id` must be passed.
- `enterprise_id` specifies the tenant id for which the reconciliation will be performed; if specified, the value passed for `multiTenantAware` must be `TRUE`. This value is the attribute `orclMTTenantGUID` in the LDAP repository that is stored in the `ENTERPRISE_ID` column of the `AppCore/Application` tables.

4.4.3.3 Examples of Use

Running the program without any arguments prints out the usage instructions; a sample invocation (in a non-multi-tenant environment) is the following:

```
java -classpath $CLASSPATH
-Doracle.security.jps.config=/home/sayarram/work/jps/jps-config-jse.xml
oracle.apps.fnd.applcore.dataSecurity.util.DSDataMigrator
-dsdburl myURL -dsdbuser fusion -silentMode true
-forceProcessAllRows true -policyStripe fscm
```

When run, the program generates a log file with the details of records processed, warnings, errors, and a summary of what has been processed.

A sample invocation in a multi-tenant environment is the following:

```
java -classpath $CLASSPATH
-Doracle.security.jps.config=/home/sayarram/work/jps/jps-config-jse.xml
oracle.apps.fnd.applcore.dataSecurity.util.DSDataMigrator
-dsdburl myURL -dsdbuser fusion -silentMode true -forceProcessAllRows true
-policyStripe fscm -multiTenantAware true -enterprise_id 202
```

4.4.4 Managing Data Security

For details on topics related to data security management, see the following documents:

- *Oracle Fusion Applications Security Guide*
- *Oracle Fusion Applications Security Hardening Guide*

For details about creating data role templates, see the "Oracle Fusion Applications Data Role Templates" chapter in the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)*.

4.5 Configuring Roles

An enterprise role or enterprise group is a collection of users and other enterprise roles, and it is stored in the domain identity store. An application role is a collection of users, enterprise roles, and application roles, and it is stored in the domain policy store.

In the Oracle Authorization Policy Manager environment, enterprise roles are referred to as external roles. In the Oracle WebLogic Administration Console, enterprise roles are referred to as enterprise groups.

Roles can be structured in a hierarchy by the relation "inherits." If a parent role inherits a child role, then the parent role can do anything that the child role can do (in addition to what the parent role can do).

Mapping an enterprise role to an application role establishes that the enterprise role inherits the application role, thus, the privileges of the enterprise role become the union of its privileges and those of the application roles to which it is mapped.

For details about managing the role mapping after the application has been deployed, see the following topics in *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)*:

- [Managing Role Mapping Policies](#)
- [Mapping External Roles to an Application Role](#)

For details about managing data security policies in Oracle Fusion applications, see the "Managing Oracle Fusion Applications Data Security Policies" section in the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)*.

For details about generating data roles with data role templates, see the "Oracle Fusion Applications Data Role Templates" section in the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)*.

This section includes the following topics:

- [Configuring Oracle Fusion Application Roles](#)
- [Configuring Enterprise Roles](#)

4.5.1 Configuring Oracle Fusion Application Roles

Oracle Fusion applications use the following enterprise and application roles in their application policies and data security policies:

- Data role, an enterprise role used exclusively in data security policies. It can inherit Job, Duty, and Abstract roles. A number of data roles are provisioned with each Oracle Fusion application.
- Job role, or Business role, is an enterprise role that corresponds with a job or business occupation. A Job role *must* inherit at least a Duty role.
- Duty role, or Task role, is an application role that corresponds with the duties of a job.
- Abstract role is an enterprise role that can be associated with any user, irrespective of his job or duties. Typical examples of this role are Employee, Manager, Customer, and Supplier. Several job roles are provisioned with each Oracle Fusion application. An Abstract role *must* inherit at least a Duty role.

Oracle Authorization Policy Manager is the recommended tool to manage application roles once the application has been deployed. Using this tool an administrator can, for example, create, remove, or modify an application role; or modify the hierarchy of application roles; or create, remove, or modify the application role category. For details about the tool, see the "Managing Applications" section in the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)*.

[Table 4–2](#) lists the equivalent terms used in the physical and reference implementations. The terminology in the physical implementation follows the one used in Oracle Fusion applications; the terminology in the reference implementation follows the one used in the Oracle Authorization Policy Manager graphic interface.

Table 4–2 Equivalent Terminology

Physical Implementation	Reference Implementation
Data role	Enterprise role used only in data security policies. Typically, the name of a data role has the suffix <code>_DATA</code> .
Job	Enterprise role mapped to application role. Typically, the name of this role has the suffix <code>_JOB</code> .
Abstract role	Enterprise role, which are persisted as LDAP groups and can be managed with Oracle Authorization Policy Manager and Oracle Identity Management.
Duty	Application role used only in application policies.
Privilege	Entitlement (or permission set).
FND Grant (or Foundation Grant)	Data security policy, which ties a data role or job role to a specific set of data.

For definitions and details about the terms in the reference implementation, see the "Understanding the Policy Model" section in the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)*.

4.5.2 Configuring Enterprise Roles

A security administrator uses integrated Oracle Identity Management pages to create and manage enterprise (job) roles in Oracle Fusion applications. For details, see the "Organization and Role Management" section in the *Oracle Fusion Middleware User's Guide for Oracle Identity Manager*.

4.6 Configuring Audit Trail

Auditing features are provided through Audit Trail, a history of the changes that have been made to data in Oracle Fusion Applications. Audit Trail enables you to track who made changes to data, at what time, and how the value changed.

For more information and configuration details, see the "Implementing Audit Trail Reporting" chapter in *Oracle Fusion Applications Developer's Guide*.

Note: Oracle Fusion Middleware Audit Framework is a separate service that provides a centralized audit framework for the middleware family of products. For details about this feature, see "Configuring and Managing Auditing" in the *Oracle Fusion Middleware Application Security Guide*.

4.7 Configuring SSL for Oracle Fusion Applications

SSL provides secure communication between the paths that connect endpoints. For example, the path between Oracle WebLogic Server and an LDAP directory server is secured through SSL.

This section contains the following topics:

- [SSL Configuration in Oracle Fusion Middleware](#)
- [SSL Configuration for Oracle Fusion Applications](#)
- [Implementing SSL for Identity Management Configuration](#)
- [Additional SSL Configuration](#)
- [Enabling Secure Sockets Layer on ECSF](#)

4.7.1 SSL Configuration in Oracle Fusion Middleware

Oracle Fusion Middleware provides SSL configuration features across the three tiers of the enterprise stack (Web, Middle, and Data tiers). SSL configuration is consistent and uniform across all Oracle Fusion Middleware system components and applications.

This section contains the following topics:

- [SSL and Infrastructure Hardening](#)
- [Communication in the Three-Tier Model](#)

4.7.1.1 SSL and Infrastructure Hardening

SSL-enabling communication paths is one element in a hardening process whose aim is to ensure that all appropriate security features are activated and configured correctly in the various major systems and subsystems that comprise Oracle Fusion Middleware.

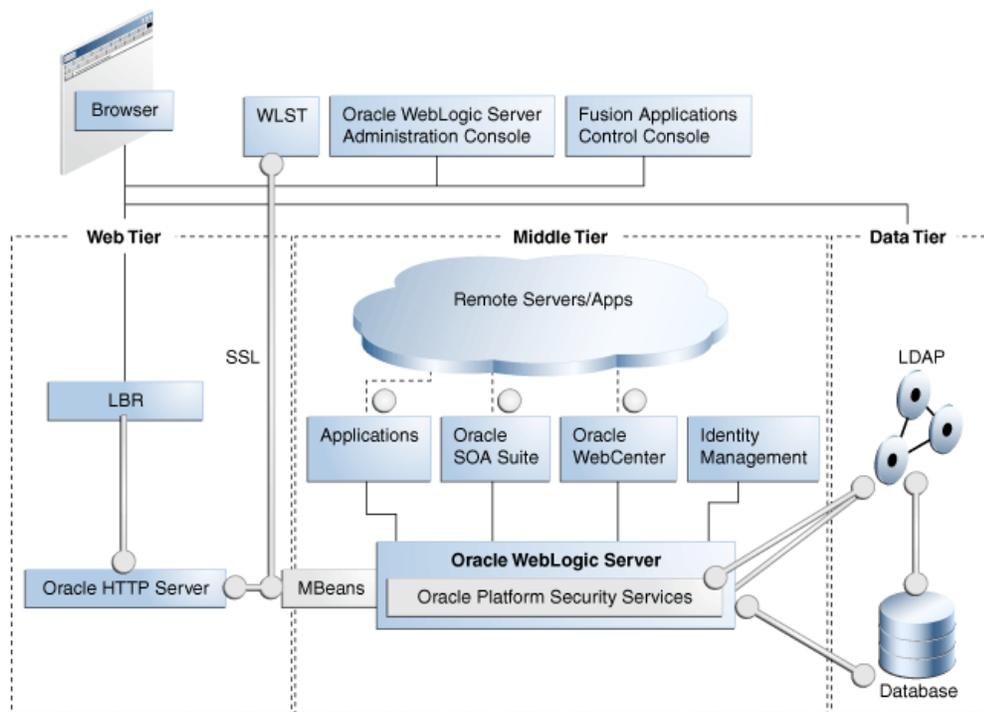
The discussion in this document is limited to SSL features available to Oracle Fusion applications. For information about other elements of hardening, see the *Oracle Fusion Applications Security Hardening Guide*.

4.7.1.2 Communication in the Three-Tier Model

Oracle Fusion Middleware supports a three-tier structure: the Web tier contains load balancers and other components outside the firewall, the middle tier hosts Oracle WebLogic Server and its applications, and the Data tier contains databases and directories. Different administration tools are shown at the top of the figure.

[Figure 4–2](#) shows the location of key elements in this architecture:

Figure 4–2 Oracle Fusion Middleware and the Three-Tier Model



In the figure, the vertical broken lines represent firewalls. The circles represents listeners that can be SSL-enabled for secure communication. As the figure shows, all critical communication paths can be protected with SSL regardless of the tier(s) involved.

For more information, see the "About SSL in Oracle Fusion Middleware" section in the *Oracle Fusion Middleware Administrator's Guide*.

4.7.2 SSL Configuration for Oracle Fusion Applications

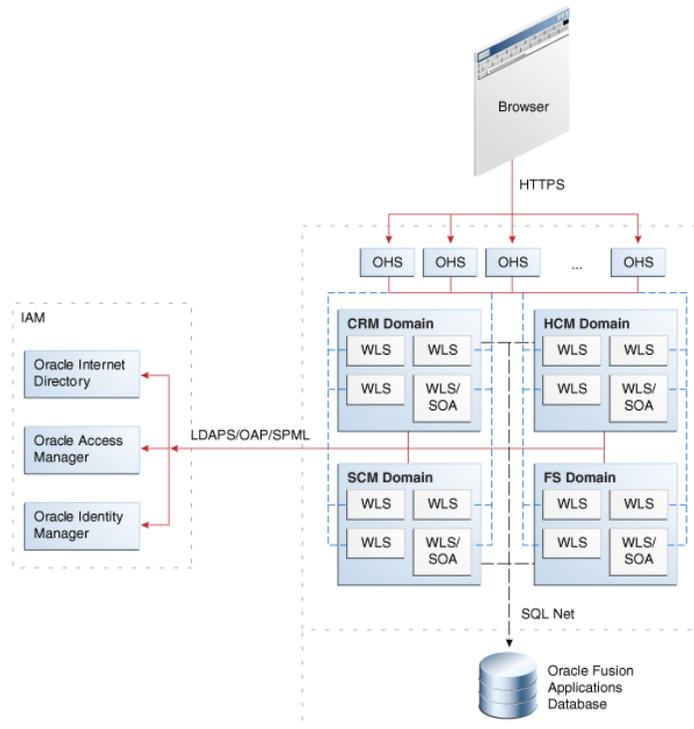
Key connections in Oracle Fusion Applications can be secured either during provisioning or post-provisioning.

This section contains the following topics:

- [Basic Network Topology](#)
- [Provisioned SSL Connections](#)

4.7.2.1 Basic Network Topology

Figure 4–3 shows a high-level representation of a three-tier network topology in the typical Oracle Fusion Applications environment:

Figure 4–3 SSL Connections for Basic Network Topology**Notes:**

- The diagram shows some representative domains.
- OAP is the Oracle Access Protocol used by Oracle Access Manager. LDAPS is the LDAP-over-SSL protocol.

Several approaches to configuring SSL are available in the Oracle Fusion Applications environment:

1. You can choose not to enable SSL connections during provisioning.
2. You can choose to SSL-enable connections to certain components during provisioning.

These connections are shown as solid lines (red) in the diagram and include IdM components like Oracle Access Manager and Oracle Internet Directory. [Table 4–3](#) lists these connections.

3. You can SSL-enable connections post-provisioning.

If you started with Option 2, you can now protect service-to-service connections like those shown in dotted lines (blue) in the diagram; this includes, for example, connections to Oracle Business Intelligence, ECSF, and external Web services. For details about this wiring, see [Section 4.7.4](#).

If you started with Option 1, you would SSL-enable Oracle Identity Management first (see [Section 4.7.3](#)), followed by the service-to-service connections (see [Section 4.7.4](#)).

Note the following assumptions about this topology:

- In most environments, the Fusion Applications middleware servers operate on an isolated network within the larger corporate network.
This means that such components as Oracle HTTP Server (OHS), Oracle WebLogic Server, Oracle Business Intelligence, and others required for the Oracle Fusion Applications instance all run within a single isolated network.
- Likewise, the Oracle Fusion Applications database runs in either the same isolated network or on its own isolated network that can only be reached by means of the application's private network.
- Because most business applications do not face the extranet, the HTTP server (OHS) tier is typically not segregated into its own DMZ.
- External dependencies for Oracle Identity Management components, represented in the figure by Oracle Access Manager and Oracle Internet Directory servers, are assumed to reside outside this isolated network.

4.7.2.2 Provisioned SSL Connections

As mentioned earlier, Oracle Fusion applications are configured with SSL for client traffic inbound to OHS and traffic to and from the identity management zone.

Table 4–3 shows the connections that are SSL-enabled at provisioning:

Table 4–3 Provisioned SSL Connections

Connection Path	Protocol	Default SSL Connection
Incoming HTTP Traffic (client to HTTP server)	HTTPS	One-way SSL (trust in server)
mod_webgate to Oracle Access Manager	OAP	One-way SSL (trust in server)
Oracle WebLogic Server to LDAP server (Oracle Internet Directory/Oracle Virtual Directory)	LDAPS	One-way SSL (trust in server)

SSL can also be enabled for these connections post-provisioning by following the instructions in Section 4.7.3.

4.7.3 Implementing SSL for Identity Management Configuration

As mentioned, you can SSL-enable the connections shown in Table 4–3 during Oracle Fusion Applications provisioning.

You can also SSL-enable these connections post-provisioning, by following the instructions provided in this section.

These instructions rely on the fact that, during the provisioning process, self-signed server certificates are created and configured for all Oracle Fusion Applications domain Weblogic Server instances. Thus:

- Each self-signed certificate is associated with an Oracle WebLogic Server instance through its identity store, located at:

`fusionapps/wlserver_10.3/server/lib/hostname_fusion_identity.jks`

- The self-signed certificates are imported as trusted certificates into a common trust store named `fusion_trust.jks`, located at:

`fusionapps/wlserver_10.3/server/lib`

This common trust store also contains the identity management trust CA certificates provided during provisioning's configuration steps.

- The common trust store is ready for all outbound SSL connections to external web services deployed with the same trusted certificates.

Using the common trust store `fusion_trust.jks` and the identity key store `hostname_fusion_identity.jks`, you can configure SSL between the Oracle HTTP Server instances and the related Oracle Fusion Applications domain Weblogic Server instances as follows:

1. Enable outbound SSL for the Web tier. For each OHS instance:
 - a. Locate the configuration file `FusionVirtualHost_fs.xml`, which resides in the following directory:


```
ohs_instance/CommonDomain_
webtier/config/OHS/ohs1/moduleconf/FusionVirtualHost_fs.xml
```
 - b. Check `Enable SSL` in the configuration file `FusionVirtualHost_fs.xml`.
 - c. Change the current (non-SSL HTTP) ports to point to the Oracle Fusion Applications Weblogic instance SSL ports.

All OHS virtual hosts are now configured for HTTPS.

2. Enable SSL for inbound connections to the middle tier. For each Oracle WebLogic Server instance in the Oracle Fusion Applications domain:
 - a. Enable SSL for the configured SSL port. You can use the Oracle WebLogic Server administrative console or a WLST command.
 - b. For Oracle WebCenter Portal, change all connections specified in the `connections.xml` file to use HTTPS instead of HTTP. Change all ports to SSL ports. As an alternative, the internal virtual host can be configured to route all incoming traffic to use HTTPS.

For an example of `connections.xml`, see *Oracle Fusion Middleware Developer's Guide for Oracle WebCenter Portal*.
 - c. For Oracle SOA Suite composite deployment connections, all composites must be rewired and redeployed.

Since all composites are already wired to the internal virtual host, a simpler approach is to redirect all HTTP traffic coming in to the internal virtual host to use HTTPS.
3. If the topology includes a load balancer (LBR), since LBR and OHS are both SSL endpoints, you must SSL-enable all internal virtual IP's, and import the CA certificates from the Web tier (Oracle HTTP Server or Oracle Web Cache) `cwallet.sso` to the LBR trust store:
 - a. Export all the trusted certificates to PEM (text format) using the `orapki` utility.
 - b. Import the PEM format certificates to the load balancer. Refer to the LBR documentation for SSL trust configuration.

4.7.4 Additional SSL Configuration

Besides the connections that are SSL-enabled during provisioning, there are additional connections in the Oracle Fusion Applications environment that can be SSL-enabled. For example, you can secure traffic to Oracle Database.

This section contains the following topics:

- [SSL-enable Oracle Identity Management](#)

- [SSL-enable Oracle Business Intelligence](#)
- [SSL-enable ECM](#)
- [SSL to External Web Services](#)

4.7.4.1 SSL-enable Oracle Identity Management

Oracle Identity Management components reside in either the middle tier or the data tier of Oracle Fusion Middleware and can be individually configured for SSL.

Table 4–4 lists the components and references for configuration details.

Table 4–4 SSL for Oracle Identity Management Components

Task	Reference
SSL-enable Oracle Internet Directory	"Enabling SSL on Oracle Internet Directory Listeners" section in the <i>Oracle Fusion Middleware Administrator's Guide</i>
SSL-enable Oracle Virtual Directory	"Enabling SSL on Oracle Virtual Directory Listeners" section in the <i>Oracle Fusion Middleware Administrator's Guide</i>
Enable SSL Between Oracle Internet Directory and the Database	"Enabling Outbound SSL from Oracle Internet Directory to Oracle Database" section in the <i>Oracle Fusion Middleware Administrator's Guide</i>
Enable SSL Between Oracle Internet Directory and Oracle Virtual Directory	"Creating LDAP Adapters" section in the <i>Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory</i>
Enable SSL for data source	"SSL-Enable a Data Source" section in the <i>Oracle Fusion Middleware Administrator's Guide</i>

4.7.4.2 SSL-enable Oracle Business Intelligence

You can configure the components of Oracle Business Intelligence to communicate over SSL.

For configuration details, see "SSL Configuration in Oracle Business Intelligence" chapter in the *Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition*.

4.7.4.3 SSL-enable ECM

You can configure SSL for Oracle ECM applications running in a production or development environment.

For configuration details, see "Configuring SSL for Oracle WebCenter Content Applications" section in the *Oracle WebCenter Content Installation Guide*.

4.7.4.4 SSL to External Web Services

For information about securing Oracle Fusion Web Services, see the following:

- "Locking Down Web Services: Points to Consider" section in the *Oracle Fusion Applications Security Hardening Guide*
- "Hardening Web Services" section in the *Oracle Fusion Applications Security Hardening Guide*

4.7.5 Enabling Secure Sockets Layer on ECSF

Enable Secure Sockets Layer (SSL) on ECSF to secure all connections that transmit passwords. These include connections to Oracle WebLogic Server, the ECSF servlet, and the Oracle SES server.

To enable SSL for ECSF, perform the following tasks:

- [Task 1, "Generate a Certificate for the Oracle WebLogic Server Instance"](#)
- [Task 2, "Enable SSL on the Oracle WebLogic Server Instance"](#)
- [Task 3, "Add the Oracle WebLogic Server's Certificate to Oracle SES's Trust Store"](#)
- [Task 4, "Enable SSL for Oracle SES Services"](#)
- [Task 5, "Configure Search Engine Instances to Use SSL"](#)

Task 1 Generate a Certificate for the Oracle WebLogic Server Instance

A certificate that contains the public key of the Oracle WebLogic Server instance hosting ECSF must be generated. This can be either a CA-signed certificate or a self-signed server certificate.

For information about how to generate the certificate, see *Oracle Fusion Middleware Securing Oracle WebLogic Server*.

As described later (Task 3), if this is a self-signed certificate, then it must be imported to Oracle SES's truststore. If it is a CA-signed certificate, it likely already exists in the truststore; if not, you must import the CA certificate into Oracle SES's truststore.

Task 2 Enable SSL on the Oracle WebLogic Server Instance

Enabling SSL connections on the Oracle WebLogic Server instance, on which the ECSF application is running, secures the ECSF servlet connections used to access the feeds and the security service. For more information, see the "Configuring SSL" chapter in *Oracle Fusion Middleware Securing Oracle WebLogic Server*.

To enable SSL on Oracle WebLogic Server:

1. Configure SSL.
 - Select the server on which the ECSF application is running.
 - Make sure that the **Allow Unencrypted Null Cipher** checkbox is deselected (default).
2. Configure the listen ports to enable the SSL listen port (check the **SSL Listen Port Enabled** checkbox).

Note: The SSL Listen Port is usually set as 7002. You can use this port number to access the ECSF servlet through SSL (for example <https://mywlserver.oracle.com:7002/approot/searchfeedservlet/ConfigFeed>). Also note that the protocol in the URL must be `https`.

3. Save your changes.

Task 3 Add the Oracle WebLogic Server's Certificate to Oracle SES's Trust Store

When Oracle SES initiates SSL connections to ECSF, the Oracle WebLogic Server instance on which the ECSF Servlet is running sends a digital certificate containing its public key to Oracle SES. If this certificate is a self-signed certificate, then the certificate must be extracted from Oracle WebLogic Server and added to Oracle SES's trust store.

(If a CA certificate, ensure that it exists in the SES trust store.) For more information about adding the certificate to the trust store, see the *Oracle Secure Enterprise Search Administrator's Guide*.

To add the Oracle WebLogic Server certificate to the Oracle SES trust store:

1. Use the `keytool` utility (located in `ORACLE_HOME/jdk/bin`) to export the Oracle WebLogic Server's certificate from the keystore, for example:

```
keytool -export -alias weblogic -keystore JAVA_
HOME/jre/lib/security/cacerts -file /temp/weblogic.cer
```

where Oracle WebLogic Server's certificate, located in the keystore at `JAVA_HOME/jre/lib/security/cacerts`, is exported to the `weblogic.cer` file in the `/temp` directory; this file now contains the server's certificate.

2. Navigate to the Oracle SES installation directory, and use `keytool` to import the Oracle WebLogic Server's certificate into the Oracle SES keystore, for example:

```
keytool -import -alias weblogic -file weblogic.cer -keystore ORACLE_
HOME/jdk/jre/lib/security/cacerts
```

where Oracle WebLogic Server's certificate in the `weblogic.cer` file is imported into the Oracle SES keystore at `ORACLE_HOME/jdk/jre/lib/security/cacerts`.

Task 4 Enable SSL for Oracle SES Services

You must also enable SSL on Oracle SES Query and Admin Services. For more information about how to enable SSL, see *Oracle Secure Enterprise Search Administrator's Guide*.

Since the Oracle SES server's certificate has not been signed by a reputable certificate authority (CA) but instead is a self-signed certificate, you must extract the certificate from Oracle SES and add it to both Oracle WebLogic Server's trust store and ECSF Command Line Administration Utility's trust store using these steps:

1. Use `keytool` to add the extracted Oracle SES server's certificate to the trust store that Oracle WebLogic Server is configured to use, for example:

```
keytool -import -alias oses -file oses.cer -keystore WLS_keystore
```

where Oracle SES server's certificate in the `oses.cer` file is imported into `WLS_keystore`, which is the keystore that Oracle WebLogic Server is configured to use.

2. Add the extracted Oracle SES server's certificate to ECSF Command Line Administration Utility's trust store by modifying the `runCmdLineAdmin.bat` and `runCmdLineAdmin.sh` scripts, located in `ORACLE_HOME/jdeveloper/ecsf`, so that `JAVA_HOME` points to the path of the keystore that Oracle WebLogic Server is using.

Task 5 Configure Search Engine Instances to Use SSL

After SSL is enabled, you must reconfigure the search engine instance parameters that contain URLs that point to the ECSF server. These parameters are `ECSF_DATA_SERVICE`, `ECSF_SECURITY_SERVICE`, and `ECSF_REDIRECT_SERVICE`.

The protocol must change from `http` to `https`, and the `http` port must change from 7101 to 7002, the SSL port. For example, you need to change `http://wlsserver.com:7101/approot/searchfeedservlet` to `https://wlsserver.com:7002/approot/searchfeedservlet`.

You must also reconfigure the search engine instance parameters that contain URLs that point to the Oracle SES server. These parameters are `SES_ADMIN_SERVICE` and `SES_QUERY_SERVICE`.

The protocol must change from `http` to `https`. For example, you need to change `http://sesserver.com:7777/search/api/admin/AdminService` to `https://sesserver.com:7777/search/api/admin/AdminService`. You do not need to change the port number because the instructions will switch the current port to be an SSL port instead of adding a separate SSL port.

You can use Fusion Applications Control to configure the engine instance parameters. For information, see [Section 7.5.3.2](#).

4.8 Managing Wallets, Keystores, Credentials, and Certificates

Enabling SSL requires the use of a number of security artifacts such as certificates and keys, as well as the containers (keystores) where they are stored. Oracle Fusion Middleware provides a number of tools to create and maintain these artifacts.

This section contains the following topics:

- [Wallets and Keystores](#)
- [Management Tools](#)
- [Managing Wallets and their Contents](#)
- [Managing Keystores and their Contents](#)
- [Managing Credentials](#)

4.8.1 Wallets and Keystores

Oracle Fusion Middleware provides two types of repositories (keystores) for keys and certificates:

- [JKS Keystore and Truststore](#)
- [Oracle Wallet](#)

For more information about wallets and keystores, see [Section 4.7.2](#).

4.8.1.1 JKS Keystore and Truststore

A JKS keystore is the default JDK implementation of Java keystores. In 11g Release 1 (11.1.1), all Java components and Java EE applications use the JKS-based keystore and truststore.

While creating a keystore, you can pre-populate it with a keypair wrapped in a self-signed certificate. Such a keystore is typically used in development and testing phases.

You can also generate a certificate signing request for a keypair and request a signed certificate back from a Certificate Authority (CA). Once the CA sends the certificate back, it is imported into the keystore. The keystore now contains a trusted certificate since it comes from a trusted third-party. Such a keystore is typically used in production.

For more information about creating and managing keystores, see the "JKS Keystore and Truststore" section in the *Oracle Fusion Middleware Administrator's Guide*.

4.8.1.2 Oracle Wallet

An Oracle wallet is a type of keystore or container that stores your credentials, such as certificates, trusted certificates, certificate requests, and private keys. You can store

Oracle wallets on the file system or in LDAP directories such as Oracle Internet Directory.

In the Oracle Fusion Applications environment, Oracle wallets are used by Oracle HTTP Server, LDAP clients, and SQL*net clients.

When creating a wallet, you can pre-populate it with a self-signed certificate. Such a wallet is called a test wallet and is typically used in development and testing phases.

You can also create a certificate request and request a signed certificate back from a Certificate Authority (CA). Once the CA sends the certificate back it is imported into the wallet. Such a wallet is called a third-party wallet.

For more information about creating and managing wallets, see the "Oracle Wallet" section in the *Oracle Fusion Middleware Administrator's Guide*.

4.8.1.3 Keystore Types Used by Products

Table 4–5 shows the type of keystore, either Oracle wallet or JKS keystore, used by various products:

Table 4–5 Keystore Types for Products

Product	Type of Keystore Used
Oracle HTTP Server	Oracle Wallet
Oracle Internet Directory	Oracle Wallet
Oracle Virtual Directory	JKS-based Keystore
Oracle WebLogic Server	JKS-based Keystore
Oracle SES Services	JKS-based Keystore
Oracle ECM	JKS-based Keystore
Oracle Business Intelligence	JKS-based Keystore

See Section 4.8.2 for a survey of common tools used to create and manage keystores. See the SSL references in Section 4.7 for the keystore management tools used by individual products.

4.8.2 Management Tools

Oracle Fusion Middleware provides these tools for keystore operations:

- WLST, a command-line interface for JKS keystores and wallets
- orapki, a command-line tool for wallets
- Fusion Middleware Control, a graphical user interface
- Oracle Wallet Manager, a stand-alone GUI tool for wallets, recommended for managing PKCS#11 wallets
- the `keytool` utility

For more information about these management tools, see the "Keystore Management Tools" section in the *Oracle Fusion Middleware Administrator's Guide*.

4.8.3 Managing Wallets and their Contents

When working with certificates and the wallets in which they are stored, you must be aware of the operations that can be performed on these objects in the course of routine administration:

- [Wallet Lifecycle](#)
- [Wallet Operations](#)
- [Certificate Lifecycle](#)
- [Certificate Operations](#)

4.8.3.1 Wallet Lifecycle

Typical life cycle events for an Oracle wallet are as follows:

- The wallet is created. Wallets can be created directly, or by importing a wallet file from the file system.
- The list of available wallets is viewed and specific wallets are selected for update.
- Wallets are updated or deleted. Update operations for password-protected wallets require that you enter the wallet password.
- The wallet password can be changed for password-protected wallets.
- The wallet can be deleted.
- Wallets can be exported and imported.

4.8.3.2 Wallet Operations

Typical operations for the Oracle Wallet include the following:

- Creating wallets, including auto-login, self-signed, and password-protected wallets
- Changing a self-signed wallet to a third-party wallet

Note: Third-party wallets contains certificates signed by a trusted Certificate Authority (CA).

- Exporting and importing wallets to and from the file system
- Deleting a wallet.

4.8.3.3 Certificate Lifecycle

The following provides a summary of the steps in a certificate's lifecycle:

1. Create an empty wallet (that is, a wallet that does not contain a certificate request).
2. Add a certificate request to the wallet.
3. Export the certificate request.
4. Use the certificate request to obtain the corresponding certificate.
5. Import trusted certificates.
6. Import the certificate.

4.8.3.4 Certificate Operations

Common certificate operations include:

- Creating a certificate request
- Exporting a Certificate, Certificate Request, or a Trusted Certificate
- Importing a Certificate or a Trusted Certificate
- Deleting a Certificate Request, a Certificate, or a Trusted Certificate
- Converting a Self-Signed Certificate into a Third-Party Certificate.

Note: Third-party certificates are signed by a trusted Certificate Authority (CA).

4.8.4 Managing Keystores and their Contents

When working with JKS certificates and the keystores in which they are stored, you must be aware of the operations that can be performed on these objects in the course of routine administration:

- [Keystore Lifecycle](#)
- [Keystore Operations](#)
- [Certificate Lifecycle](#)
- [Certificate Operations](#)

4.8.4.1 Keystore Lifecycle

Typical life cycle events for a JKS keystore are as follows:

- The keystore is created. Keystores can be created directly, or by importing a keystore file from the file system.
- The list of available keystores are viewed and specific keystores selected for update.
- Keystores are updated or deleted. Update operations require that the keystore password be entered.
- The keystore password can be changed.
- The keystore can be deleted.
- Keystores can be exported and imported.

4.8.4.2 Keystore Operations

Typical keystore operations include the following:

- Creating or updating a keystore.
- Exporting and importing keystores.
- Deleting a keystore.
- Changing the keystore password.

4.8.4.3 Certificate Lifecycle

Typical life cycle events for a certificate residing in a keystore are as follows:

- A self-signed certificate is automatically created for the keypair.
- A certificate signing request (CSR) is generated, and can then be exported to a file.
- Certificates are imported into the keystore. You can import both user certificates and trusted certificates (also known as CA certificates) in this way.
- Certificates or trusted certificates are exported from the keystore out to a file.
- Certificates or trusted certificates are deleted from the keystore.

4.8.4.4 Certificate Operations

Common operations on JKS certificates include the following:

- Generating a new key (that is, a new self-signed certificate) for a keystore.
- Generating a certificate signing request.
- Importing/exporting a certificate or trusted certificate into/from a keystore.
- Deleting a certificate or trusted certificate from a keystore.

4.8.5 Managing Credentials

In Oracle Fusion Applications, user and role information such as passwords are maintained in a domain credential store. The tools you use to update existing passwords (for routine administration or regulatory compliance), depend on the type of credentials:

- WebLogic data sources

These objects contain properties such as the URL or user name and password. Application components use data sources to obtain connections to a relational database.

You can create and manage JDBC data sources using the administration tools provided with Oracle WebLogic Server. For more information about managing data sources, see *Oracle Fusion Middleware Configuring and Managing JDBC Data Sources for Oracle WebLogic Server* and the "Configure JDBC data sources" topic in *Oracle Fusion Middleware Oracle WebLogic Server Administration Console Online Help*.

- Application-specific credentials

These credentials provide access to different components of your application. An example is the credential associated with the designated super-user.

Application credentials are maintained through Fusion Middleware Control or Oracle WebLogic Scripting Tool (WLST). For details, see the "Configuring the Credential Store" topic in the *Oracle Fusion Middleware Application Security Guide*.

This section contains the following topics:

- [Changing App ID Passwords](#)
- [Changing the Oracle Fusion Middleware Administrative User Password](#)

4.8.5.1 Changing App ID Passwords

When invoking Web services, Oracle Fusion Applications must rely on a type of credential known as the Application ID or App ID. Each application has its own App ID which is initially provisioned for the application.

For information about resetting App ID passwords (a task typically done during scheduled downtime), see the Application Identity Password Reset And Password Policy Management section in the *Oracle Fusion Applications Security Guide*.

4.8.5.2 Changing the Oracle Fusion Middleware Administrative User Password

During the Oracle Fusion Applications installation, you must provide a password for the Oracle Fusion Middleware administration user. By default, this administrator takes the Super User value you specified on the Identity Management page when creating the provisioning plan. The password is the one specified when adding the user to the identity store.

Then, you can use this account to log in to Fusion Applications Control and the Oracle WebLogic Server Administration Console for the first time. You can create additional administrative accounts using the WLST command line or the Oracle WebLogic Server Administration Console.

You can change the password of the administrative user using the Oracle WebLogic Server Administration Console or the WLST command line.

This section contains the following topics:

- [Changing the Oracle Fusion Middleware Administrative User Password Using the Command Line](#)
- [Changing the Oracle Fusion Middleware Administrative User Password Using the Administration Console](#)

4.8.5.2.1 Changing the Oracle Fusion Middleware Administrative User Password Using the Command Line

To change the Oracle Fusion Middleware administrative user password or other user passwords using the command line, you invoke the `UserPasswordEditorMBean.changeUserPassword` method, which is extended by the security realm's `AuthenticationProvider` MBean.

For more information, see the `changeUserPassword` method in the *Oracle Fusion Middleware Oracle WebLogic Server MBean Reference*.

4.8.5.2.2 Changing the Oracle Fusion Middleware Administrative User Password Using the Administration Console

To change the password of the Oracle Fusion Middleware administrative user using the Oracle WebLogic Server Administration Console:

1. Navigate to the Oracle WebLogic Server Administration Console. (For example, from the home page of the domain in Fusion Applications Control, select **To configure and managed this WebLogic Domain, use the Oracle WebLogic Server Administration Console**.)
2. From the target navigation pane, select **Security Realms**.
The Summary of Security Realms page is displayed.
3. Select a realm, such as **myrealm**.
The Settings for the realm page is displayed.
4. Select the **Users and Groups** tab, then the **Users** tab. Select the user.
The Settings for *user* page is displayed.
5. Select the **Passwords** tab.

6. Enter the new password, then enter it again to confirm it.
7. Click **Save**.

4.9 Data Masking

Data masking is the ability to replace sensitive data with realistic but false data on test and development databases. Features include:

- the ability to keep data properties (data type, width, and so on) intact and provide realistic data sets for analysis and testing.
- ensuring various constraints like (Primary Keys, Uniqueness, Foreign keys) are maintained, preserving relational integrity.
- using user specified formatting rules that ensure custom and packaged applications continue to work after the data is masked.

This section contains the following topics:

- [Introduction to Data Masking](#)
- [Data Masking in Oracle Fusion Applications](#)
- [Managing the Masking Definitions](#)
- [Data Masking in Oracle Enterprise Manager Cloud Control](#)
- [Best Practices when Masking Test Databases](#)
- [References](#)

4.9.1 Introduction to Data Masking

Data masking is an on-going activity for the Oracle Fusion Applications administrator. It requires an understanding of masking concepts, methodology, and the implementation tools:

- [Masking Terminology](#)
- [Types of Sensitive Data](#)
- [The FAST Methodology](#)
- [Administration Tools](#)

See Also: *Oracle Database Real Application Testing User's Guide.*

4.9.1.1 Masking Terminology

Key terms used in data masking are as follows:

- Pre Masking script - A SQL script that runs prior to the start of masking.
- Post masking script - A SQL script that executes after all masking completes. For example, such a script can be used to recompute aggregated columns after the detailed data is masked. This ensures that aggregated and masked columns are consistent and the totals match.
- User-Defined Function (UDF) - A user-defined function takes the original value, the row id, and column name to generate the mask value. A single column format can be a combination of one or more formats including UDF.
- Post-processing function - This is a special case of a user-defined function. A post processing function (PPF) is called after the mask value is generated using the

specified format. The function takes the generated mask value and further modifies it to produce the actual mask value.

For example if the format used Random Number (1000,10000) and Post Processing Function (checksum) a number between 1000 and 10000 is generated and this value is fed into the PPF. The PPF computes the checksum and appends it to the original number and returns the new mask value.

There can be only one PPF to a column. A PPF cannot be the only format for a column. There has to be some other format preceding the function.

4.9.1.2 Types of Sensitive Data

Examples of sensitive data include:

- Salary information
- Government IDs like drivers' license numbers and social security numbers
- Geographical data such as GPS coordinates

Each product family has its own set of sensitive data types.

For details about sensitive data types, see the *Oracle Fusion Applications Security Guide*.

4.9.1.3 The FAST Methodology

Data masking requires the user to implement the FAST (Find, Assess, Secure, Test) methodology:

- Find

The first step is to identify sensitive data that should be subject to data masking. This includes personally identifiable data or information that could be misused. Examples include:

- Salary Information
- Driver's License Number
- Military Service ID
- Biometrics Data

Another aspect of this effort is to determine where this data is located. The recommended procedure involves:

- Defining pattern-matching rules against sensitive tables and columns. Examples of such rules include a column name like `"*SSN*"` or column format `"###-##-####"`.

For database independence, applications typically do not store the primary key-foreign key relationships in the database itself; rather, the relationships are enforced in the application. To support this, the Data Masking Pack provides administrators with the ability to register these relationships so that columns in related tables, such as `EMPLOYEE_ID`, `MGR_ID`, are masked identically using the same masking rules.

- Searching Oracle databases to find these patterns.
- Importing the relevant database fields into a data privacy catalog.
- Maintaining the catalog as new fields are identified.

- Assess

In this phase you specify how each sensitive field is to be masked, that is, transformed into a non-sensitive representation while maintaining the field's structure. This is typically accomplished through a data masking definition, which associates tables and columns in a schema with appropriate masking formats. An example is a mask format that converts the name "John" in the Employee table to "Andy".

The Data Masking definition contains a list of sensitive columns in the application tables, such as employee social security numbers, and its corresponding association with data masking formats, such as a fictitious social security number generator.

Note: Oracle Fusion Applications provide a set of default mask templates for basic sensitive fields; administrators can update these definitions according to application requirements. For details about viewing and updating masking templates, see [Section 4.9.3](#)

- Secure

This phase enables the masks to be executed securely to generate the masked data. Notable aspects of this phase include:

- Production database cloned in restricted mode for the execution
- Privilege delegation allowing the mask to be executed with tools like sudo or PowerBroker
- Generation of test database

- Test

The test phase involves comparing before and after values for verification. Redo logs are available to restore data to pre-mask state.

4.9.1.4 Administration Tools

Data masks and masking definitions are created and managed with Oracle Enterprise Manager Grid Control.

4.9.2 Data Masking in Oracle Fusion Applications

The Oracle Fusion Applications administrator must take certain considerations into account when implementing data masking for an application. For example, free space requirements must be evaluated to ensure that adequate resources are available to the masking job. The following instructions provide guidelines:

- [Requirements for Data Masking](#)
- [Sensitive Data in Oracle Fusion Applications](#)
- [Masking Definitions](#)

4.9.2.1 Requirements for Data Masking

You should be aware of certain background information, prerequisites, and requirements before undertaking data masking operations in your environment.

This section contains the following topics:

- [Data Model Descriptions](#)

- [Required Versions](#)
- [Preliminary Steps](#)
- [Temporary Space Requirements](#)
- [Database Free Space Requirements](#)
- [Role Requirements](#)
- [Custom Field Masks](#)
- [Production-to-Test Requirement](#)

4.9.2.1.1 Data Model Descriptions You must understand the data model for your application data when deciding how and what to mask. For data model descriptions, see the product-specific documentation from Oracle Enterprise Repository for Oracle Fusion Applications.

4.9.2.1.2 Required Versions

Data masking requires the following:

- Oracle Database 10g Release 2 or Oracle Database 11g Release 1
- Oracle Enterprise Manager Grid Control 10g Release 4 or Oracle Enterprise Manager Cloud Control 12c
- The Oracle Data Masking Pack

4.9.2.1.3 Preliminary Steps

You must perform certain steps, such as installing a format library and configuring temp spaces, before using data masking in an Oracle Fusion Applications environment:

1. Install the FMTLIB package, which contains functions needed to perform data masking.
 - a. Locate these scripts in your Oracle Enterprise Manager installation:

```
$ORACLE_HOME/sysman/admin/emdrep/sql/db/latest/masking/dm_fmtlib_pkgdef.sql
$ORACLE_HOME/sysman/admin/emdrep/sql/db/latest/masking/dm_fmtlib_pkgbody.plb
```
 - b. Copy these scripts to a directory in your target database installation.
 - c. Execute the scripts using SQL*Plus, connected as a user that can create packages in the DBSNMP schema.
2. Apply any required patches. Data masking requires a specific Oracle Database version and patch set. See [Section 3.7](#).
3. Change the temp spaces used for data masking to "Auto Extend" from Grid Control:
 - a. Click the **Targets** tab.
 - b. Click the **Databases** secondary tab.
 - c. Select the database, and click the **Server** tab.

ORACLE Enterprise Manager
Grid Control 11g

Hosts | Databases | Middleware | Web Applications | Services | Systems | Groups | Virtual Servers | All Targets

Database Instance: May3_DB

Home Performance Availability **Server** Schema Data Movement Software and Support

General Shutdown Black Out


 Status Up
 Up Since Jun 30, 2010 12:00:49 PM PDT
 Instance Name g01ys580
 Version 11.1.0.7.0
 Host dbhost12.us.example.com
 Listener

[View All Properties](#)

Host CPU



Load 3.30 Paging 310.19

d. Click Datafiles.

ORACLE Enterprise Manager
Grid Control 11g

Hosts | Databases | Middleware | Web Applications | Services | Systems | Groups | Virtual Servers | All Targets

Database Instance: May3_DB

Home Performance Availability **Server** Schema Data Movement Software and Support

Storage

- Control Files
- Tablespaces
- Temporary Tablespace Groups
- Datafiles**
- Rollback Segments
- Redo Log Groups
- Archive Logs
- Migrate to ASM
- Make Tablespace Locally Managed

Database Configuration

- Memory Advisors
- Automatic Undo Management
- Initialization Parameters
- [View Database Feature Usage](#)

Statistics Management

- Automatic Workload Repository
- AWR Baselines

Resource Manager

- Getting Started
- Consumer Groups
- Consumer Group Mappings

e. On the Datafiles page, enter temp in the Object Name search box. Press Go. Two temp spaces are displayed:

ORACLE Enterprise Manager
Grid Control 11g

Hosts | Databases | Middleware | Web Applications | Services | Systems | Groups | Virtual Servers | All Targets

Database Instance: May3_DB >

Datafiles

Search

Select an object type and optionally enter an object name to filter the data that is displayed in your results set.

Object Name

By default Datafiles are case-sensitive searches. To run an exact match, double quote the search string. You can use the wildcard symbol (%) in a double quoted string.

Actions

Select	File Name	Tablespace	Status
<input checked="" type="radio"/>	/slot/ems4207/oracle/db/apps_st/data/temp1_03.dbf	TEMP1	ONLINE
<input type="radio"/>	/slot/ems4207/oracle/db/apps_st/data/temp1_04.dbf	TEMP2	ONLINE

f. Edit each temp space in turn by selecting its radio button and clicking Edit. Change the value for each space as shown here:

Database Instance: May3_DB > Datafiles >
 Edit Datafile: temp1_03.dbf

File Name **temp1_03.dbf**
 File Directory **/slot/ems4207/oracle/db/apps_st/data/**
 Tablespace **TEMP1**
 Status Online Offline
 File Size MB

Storage
 Automatically extend datafile when full (AUTOEXTEND)
 Increment MB
 Maximum File Size Unlimited
 Value MB

4.9.2.1.4 Temporary Space Requirements As shown in Step 3 of [Section 4.9.2.1.3](#) it is recommended that you autoextend your temp files as masking requires additional space for processing. The amount of additional space you need depends on your data. Broadly speaking, masking takes up approximately two times the size of the largest table being masked.

The temp space is needed for two reasons:

- To perform sort and join operations during masking.

The space requirement is not straightforward to estimate as it depends on whether sorts and joins go to disk, which in turn depends on how much memory is available on the machine. Try to keep at least as much space as the size of your biggest masked table for sort and join processing.
- For space taken up in the default user tablespace for temporary masking tables.

The size of the temp tables is twice the size of all the columns being masked. Since these tables are dropped after processing, the space is released after masking completes.

4.9.2.1.5 Database Free Space Requirements

Ensure that sufficient free space is available to the database before executing the masking job.

Calculate the free space requirements as follows:

largest table being masked +
 total size of mapping tables for all columns in that table +
 temporary tablespace (roughly twice the size of the largest mapping table, as stated under **Temporary Space Requirements** above)

4.9.2.1.6 Role Requirements

The user executing the data masking script must have the dba role. If Virtual Private Database (VPD) security policies are used or Oracle Database Vault is enabled for the database, the user must be SYS.

Additionally, the user may need to have direct grants to objects for any PL/SQL objects provided to user-defined functions or post-processing scripts.

4.9.2.1.7 Custom Field Masks

Oracle Fusion Applications include a set of out-of-the-box masking definitions to mask common sensitive data like employee information and credit card numbers. However,

these default mask templates cannot account for custom data such as flex fields since these are specific to your application.

You can update the standard masking definitions to include additional flex fields and custom data. For details about viewing and updating masking templates, see [Section 4.9.3](#).

Note: Be aware that certain sensitive custom attributes may depend on tables in other product families. For example, Procurement masking is dependent on CRM (for Suppliers and related entities, and for Procurement Contract terms and deliverables), on SCM (for Items), on HCM (for Workers and Organizations), on FIN (for Ledgers), and on Projects (for Projects and related entities).

4.9.2.1.8 Production-to-Test Requirement

During the production-to-test process, you must replace the production user names with dummy user names. This mandatory step is needed to avoid breaking anonymization.

4.9.2.2 Sensitive Data in Oracle Fusion Applications

Oracle Fusion Applications identify the common sensitive data types for each product family.

You can use Oracle Enterprise Manager Grid Control or Oracle Enterprise Manager Cloud Control to view the sensitive attributes specified in the masking definitions. You can update the definitions with any additional desired sensitive attributes, such as flex fields or other customized fields.

For details about how to view and update masking templates, see [Section 4.9.3](#) and [Section 4.9.4](#).

For additional information about data masking, see the *Oracle Fusion Applications Security Guide*.

4.9.2.3 Masking Definitions

A masking definition specifies the columns to be masked and the format of the masked data. Out-of-the-box "template" masking definitions are provided for each family of Oracle Fusion Applications.

Masking definitions in XML format masks enable the data masking utility to identify the database tables, columns, and column formats of the data being masked.

For details, see the *Oracle Fusion Applications Security Guide*.

4.9.3 Managing the Masking Definitions

You can modify the default masks provided with Oracle Fusion Applications, and create your own masking definitions. Use Oracle Enterprise Manager Grid Control (Grid Control) to manage and maintain the data masking definitions.

Note: If using Oracle Enterprise Manager Cloud Control, see [Section 4.9.4](#).

This section contains the following topics:

- [Viewing and Modifying Data Masking Definitions](#)
- [Generating the Masking Script](#)
- [Customizing Mask Formats](#)

Oracle Enterprise Manager Grid Control online help provides more details on these topics.

Note: You must follow the instructions in [Section 4.9.2.1](#) before you can perform the operations described here.

4.9.3.1 Viewing and Modifying Data Masking Definitions

You can view and modify out-of-the-box data masking definitions as follows with Grid Control:

1. Click the **Targets** tab.
2. Click the **Databases** secondary tab.
3. Select the database whose mask definition you are configuring from the list of databases, and log in.
4. Click the **Schema** tab.
5. On the database instance page, scroll down to the **Data Masking** menu. The options include:
 - **Definitions:** Manage masking definitions.
 - **Format Library:** Manage masking formats.
6. Select **Definitions**. The Data Masking Definitions page lists the current masks for the database:

Data Masking Definitions

Data masking is the process of making sensitive information in test or non-production databases safe. It disguises sensitive information by overwriting it with realistic looking but false data of a similar type. A masking definition defines the columns to be masked and the format of masked data. You can create a new masking definition or use an existing definition for a masking operation. The Format Library contains a collection of ready-to-use masking formats.

Search

Database ▾

"db.example.com"

Go

Import

Create

View

Edit

Generate Script

Schedule Job

Delete

Actions

Clone Database ▾

Go

	Masking Definition	Database	Description ▲	Columns	Status	Most Recent Job Ended
<input checked="" type="radio"/>	HR_COMPOUND_MASK	db.example.com	Compound Mask of HR Data	6	Masking Job Succeeded	Sep 10, 2010 9:27:41 PM GMT+00:00
<input type="radio"/>	CONDITIONAL_EMPLOYEE_DATA_MASK	db.example.com	Mask Employee Data	7	Script Not Generated	

7. Select the mask definition you wish to view or update and click **Edit**.

The Edit Mask Definition page appears. It lists the columns included in the mask definition.

Edit Masking Definition: **CONDITIONAL_EMPLOYEE_DATA_MASK**

Cancel OK

* Name:

* Database:

Description:

Columns

Add columns you want to mask and define masking format for each column. Foreign key columns are automatically added to maintain referential integrity. Dependent columns are columns that do not have foreign key constraints defined, but reference a masked column due to application level constraints. You can manually add dependent columns to a masked column. Removing a column from this list will remove all foreign key and dependent columns.

Add

Remove

Select All | Select None

Select	Owner	Table	Column	Column Group	Data Type	Format	Foreign Key Columns	Dependent Columns	
								Count	Add
<input type="checkbox"/>	HR	EMPLOYEES	EMPLOYEE_ID		NUMBER		5	1	
<input type="checkbox"/>	HR	EMPLOYEES	FIRST_NAME		VARCHAR2(20)		0	0	
<input type="checkbox"/>	HR	EMPLOYEES	LAST_NAME		VARCHAR2(25)		0	0	

8. A number of operations are available on this page:
- To add another column to the definition, click **Add**.
 - To modify a column format, click the **Format** icon.
 - To remove a column from the mask definition, check the box and click **Remove**.
9. To introduce a new column into the definition, for example, click **Add**. The Add Columns page appears.
10. To locate the column, enter the schema and table name and click **Search**.
In this example, we search the **EMPLOYEES** table in the **HR01** schema:

Add Columns

Database: **db.example.com** Logged In As: **system** Cancel Add Define Format And Add

Add one or more columns for masking. Foreign key columns will be added automatically. You can define masking format at once for all selected columns if they have the same data type.

Search

Schema: Column Name:

Table Name: Column Comment:

 Enter a string in column comments.

Mask selected columns as a group

Select All | Select None

Select	Owner	Table Name	Column Name	Data Type	Comment
<input checked="" type="checkbox"/>	HR01	EMPLOYEES	CITY	VARCHAR2(30)	
<input type="checkbox"/>	HR01	EMPLOYEES	COMMISSION_PCT	NUMBER(2,2)	Commission percentage of the employee; Only employees in sales department eligible for

11. Select the checkbox for the **CITY** column and click **Add**.
The Edit Masking Definition page reappears, with **CITY** added to the column list.

Unlike the other columns, the format for **CITY** is displayed as a toolbox icon, which means that no mask format has yet been defined for the column. You must specify a format for this column to complete the definition.

12. Click on the toolbox icon. The Define Column Mask page appears.
13. You can now specify a mask format entry for **CITY** using the drop-down box. Click **Add** to complete the process.

In this example we select a random string format and specify the start and end lengths. To see how the masked data will appear in this format, click the **Sample** icon.

Define Column Mask

Owner **HR** Table **EMPLOYEES** Cancel OK
 Column **CITY** Data Type **VARCHAR2(30)**

By default all records in the table will be masked using the specified format. You can optionally identify more than one subset of records using conditions. Each subset can be masked using a corresponding masking format. The subsets will be masked in the order they are specified. A subset will not be masked again even when it matches a subsequent condition.

Add Condition

Import Format Format Entry **Random Strings** Add

Expand All | Collapse All

Select Condition	Format Entry Properties				Sample	Remove
	Property	Value	Property	Value		
▼ Conditions						
⊙ ▼ Default Condition					aaaaaagqoxrb 	
Random Strings	Start Length	<input type="text" value="2"/>	End Length	<input type="text" value="12"/>		

Cancel OK

Note: You can use the Import Format button to import an existing format entry.

14. Click **OK**. Grid Control checks that the entry is valid for the column's datatype. The Edit Masking Definition page reappears, and the **CITY** column now has a valid mask format.

Click **OK** to save the masking definition.

Likewise, starting at the Edit Mask Definition page (Step 8 above) you can modify the definition of an existing column by clicking its **Format** icon. This is useful, for example, when you wish to customize the properties of a column's format entry, or specify a different format entry for the column.

4.9.3.2 Generating the Masking Script

When you modify a data masking definition by adding or removing columns, or by modifying an existing column's mask format, you must regenerate the SQL masking script to incorporate your changes. To generate the script:

1. Follow steps 1 through 6 of [Section 4.9.3.1](#) to display the data masking definitions for the database.
2. After a masking definition is modified, its status is listed as "Script Not Generated."

Select the definition and click **Generate Script**.

3. After processing is complete, you can:
 - Integrate the masking script with the clone database process
 - Schedule execution of the script to perform the masking operation.

Script Generation Results: CONDITIONAL_EMPLOYEE_DATA_MASK

Database **db.example.com** Number of Tables **6** [Return](#)
 Logged In As **system** Columns **13**

Script Options

Use script to clone and mask the database. [Clone And Mask](#)

Schedule the data masking job. The script will be executed by the job to perform the masking operation. [Schedule Job](#)

Script

The script summary is a list of the database commands that will be used to mask the selected columns. The full script includes functions, procedures, and other commands needed during the masking operation. The full script will be executed during the masking operation.

View Script Summary Full Script

```
-- Target database: db.example.com
-- Script generated at: 13-SEP-2010 04:14
COMMIT
ALTER SESSION ENABLE PARALLEL DML
DROP TABLE "MGMT_DM_TT_32" PURGE
declare
  adj number:=0;
  num number:=0;
begin
```

Note: The results page also displays the generated SQL script, which can be modified as well.

4.9.3.3 Customizing Mask Formats

The format library is a collection of common, ready-to-use masking formats you can use in masking definitions. Oracle Data Masking enables you to extend the default mask format library by tailoring the existing formats, or creating new formats to meet your own business needs.

Take these steps to add a new format:

1. Follow steps 1 through 5 of [Section 4.9.3.1](#) to display the database instance.
2. Select Format Library.

Format Library

The Format Library contains a collection of ready-to-use masking formats which can be used when creating a masking definition.

Search [Go](#) [Export](#) [Import](#) [Create](#)

[View](#) [Create Like](#) [Edit](#) [Delete](#) [Previous](#) 1-25 of 27 [Next 2](#)

Select	Format	Data Type	Sample	Description	Owner
<input type="radio"/>	Anglo American First Name	Source Type	Not Generated	Anglo American First Name	SYSMAN
<input type="radio"/>	Anglo American Last Name	Source Type	Not Generated	Anglo American Last Name	SYSMAN
<input type="radio"/>	Bay Area Phone Number	Character	(408) 555-6964	Bay Area Phone Number	SYSMAN
<input type="radio"/>	Social Security Number	Character	585758911	Social Security Number	SYSMAN
<input type="radio"/>	Social Security String	Character	555-22-2309	Social Security String	SYSMAN
<input type="radio"/>	Expiration Date	Date	2012-11-07	Expiration Date	SYSMAN

- When adding a new format, you can take an existing format as a starting point. For example, to create a new format for credit card number fields, you can select an existing credit card format and click **Create Like**.

Create Format

Cancel OK

* Name

Description ~10 billion unique American Express credit card numbers

Format Entries

Define masking format by adding one or more format entries of different types.

Add Go

Type	Description	Edit	Remove
Random Digits	Digits Length Range: 10 - 10		

Post Processing Function DBSNMP.DM_FMTLIB.MGMT_DM_GEN_AC

The function can either be a standalone function (Example: scott.masking_func) or a function specified inside of a package (Example: scott.masking_pkg.checksum).

- Enter a descriptive name for the new format, and use the format entry drop-down list to select a masking format. You may also specify a different post-processing function if desired.

Similarly, you can modify existing formats or create new formats on the Format Library page.

4.9.4 Data Masking in Oracle Enterprise Manager Cloud Control

Data masking on Cloud Control requires the following templates, which are shipped in the emfa sysman directory:

```
ADM_Oracle_Fusion_Applications_1.0_EM_12.1.0.1.0_Combined_Template.xml
Mask_Oracle_Fusion_Applications_1.0_EM_12.1.0.1.0_Combined_Template.xml
```

Use the following procedure to set up and run data masking on Cloud Control:

- Grant Permissions.
 - Use the FUSION user for all subsequent steps. Grant select_catalog_role and 'select any dictionary' and DBA to the FUSION user.
- Deploy the TDM package.
 - Navigate to Enterprise, then Jobs, then Job Activity.
 - In the "os command" drop-down, select Deploy Test Data management packages and click **Go**.
 - Enter the job name, click **add**, and add the database where masking is to be run.
 - Click parameters. In the drop-down, select fusion driver.
 - Click credential and enter the FUSION credentials. This user must have all the privileges specified in Step 1.
 - Click **submit**.
 - When the job completes, verify that it succeeded. Review the job details and make sure there are no errors.
- Import the Application Data Model (ADM).

- a. Navigate to Enterprise, then Quality Management, then Data Discovery and Modelling menu.
 - b. Click **Action**, then **Import**. Provide the ADM xml file `ADM_Oracle_Fusion_Applications_1.0_EM_12.1.0.1.0_Combined_Template.xml`, the ADM name, and the fusion database.
 - c. If prompted for database credentials, provide the fusion user credentials.
 - d. You might see one or more warnings indicating that duplicate sensitive types were not imported; these warnings can be safely ignored.
4. Create a verification job.
 - a. Navigate to Enterprise, then Quality Management, then Data Discovery and Modeling.
 - b. Click on the ADM you just imported and click **verify**.

Note: You can also get to this task by navigating to Action, then Verify.

- c. Click **Create Verification Job**.
 - d. Provide a job name and job description.
 - e. Click New Credential, and provide the FUSION credentials.
 - f. Schedule the job to start immediately, and click **Submit**.
5. Check the job status.
 - a. When the verification job completes, navigate to Enterprise, then Job Activity, then select all job status.
 - b. Click **go** and check that your job completed successfully.
 6. Import the masking template file.
 - a. Navigate to Enterprise, then Quality Management, then Data Masking Definition. The drop-down list contains data masking definitions and formats.
 - b. Import the masking template file `Mask_Oracle_Fusion_Applications_1.0_EM_12.1.0.1.0_Combined_Template.xml`.
 - c. Provide the masking definition name.
 - d. For the ADM name, specify the name of the fusion ADM you created.
 - e. For the database name, specify the fusion database.
 7. Select the mask definition and click the **generate script** button.

Script generation can take a few hours, so consider running it on VNC or on a terminal that will be available for a while to allow the job to complete. (If you run it on a laptop and have to disconnect it during execution, you will lose the browser that was generating the script.)
 8. When the job completes, you can view the impact report and save or view the script.
 9. Submit the masking script for execution by clicking **submit job**.

4.9.5 Best Practices when Masking Test Databases

Identity data can reside in a number of repositories: the Oracle Database, the production identity store (an LDAP store), and the Oracle Identity Manager database.

In the current release, when you implement data masking, the identity attributes are only masked (anonymized) in the Oracle Fusion Applications database. To preserve masked values, ensure that these best practices are followed when using the masked data for a test database:

1. Do not run the Enterprise Scheduler Service (ESS) job to synchronize the LDAP identity store and the Oracle Fusion Applications database, since doing so would reset the identity attributes in the database to their unmasked values.
2. Set up "dummy" test users to perform the testing.
3. Do not:
 - Log in to the test database as a real user.
 - Update a user's attributes in either the LDAP identity store or the Oracle Identity Manager database.

Doing so will reset that user's attributes to their unmasked values, allowing them to see their own masked record in employee self-service, deduce other people's identities by following the hierarchy, and so on.

In short, it is essential that you maintain close access control to databases holding live cloned data, and make testers aware of their obligations and responsibilities regarding the data. As far as is possible, the processes around test data based on live data should mirror the processes around the live data itself. In particular, testers must not use privileged access to look beyond what they need to perform the required tests. For example, they must not try to work out to whom the data pertains, or try and find private information to satisfy their own curiosity.

4.9.6 References

For a detailed tutorial of data masking, see "Replacing Sensitive Data Using the Data Masking Pack" at:

http://www.oracle.com/technology/obe/11gr1_db/security/datamask/datamask.htm

For additional information about securing sensitive data, Oracle Fusion Applications data security policies, and related topics, see the *Oracle Fusion Applications Security Guide*. The Privacy chapter explains how Personally identifiable information (PII) is defined and protected.

For deployment guidelines, see *Oracle Fusion Applications Customer Relationship Management Enterprise Deployment Guide*.

4.10 Securing Web Services

Oracle Web Services Manager (WSM) provides a policy framework to manage and secure Web services consistently across your organization. Oracle WSM is available to the following users:

- Developers, at design time through Oracle JDeveloper
- System administrators in production environments by means of Fusion Middleware Control and command-line tools

The Oracle WSM policy framework secures Web services with policies, which describe capabilities and requirements such as whether and how a message must be secured, whether and how a message must be delivered reliably, and so on.

This section contains the following topics:

- [Local Policy Attachment](#)
- [Global Policy Attachment](#)
- [Web Services Security Profiles](#)
- [Key Exchange with the Domain Hosting Oracle Identity Manager](#)
- [Web Services Security Hardening](#)

For more information about the Oracle WSM policy framework, see the "Understanding Oracle WSM Policy Framework" section in the *Oracle Fusion Middleware Security and Administrator's Guide for Web Services*.

4.10.1 Local Policy Attachment

A policy subject is the target resource to which the policies are attached. Examples include Web services endpoints, Web service clients, SOA service endpoints, SOA clients, and SOA components.

Directly attaching one or more policies to a policy subject is referred to as "Local Policy Attachment" (LPA). [Table 4-6](#) lists key tasks related to LPAs.

Table 4-6 Common LPA Operations

Operation	Reference
Configure LPA using Fusion Middleware Control and WLST	"Attaching a Policy to a Single Subject" section in the <i>Oracle Fusion Middleware Security and Administrator's Guide for Web Services</i>
Attach "no behavior" policies for Web Services and Web Service Clients	"Disabling a Globally Attached Policy" section in the <i>Oracle Fusion Middleware Security and Administrator's Guide for Web Services</i>
Remove LPAs for all Web services	"Attaching a Policy to a Single Subject" section, subsection titled Attaching a Policy to a Web Service Using WLST (Step 4) in the <i>Oracle Fusion Middleware Security and Administrator's Guide for Web Services</i>
Troubleshoot Web service security	"Diagnosing Problems" chapter in the <i>Oracle Fusion Middleware Security and Administrator's Guide for Web Services</i>

4.10.2 Global Policy Attachment

A policy set, which can contain multiple policy references, enables you to attach policies globally to a range of endpoints of the same type. By attaching policies globally in this way, referred to as Global Policy Attachment (GPA), you can ensure that all subjects are secured by default. [Table 4-7](#) lists key tasks related to GPAs.

Table 4-7 Common GPA Operations

Operation	Reference
GPA concepts and usage; determining if a Web service or client is security-enabled	"Attaching Policies Globally Using Policy Sets" section in the <i>Oracle Fusion Middleware Security and Administrator's Guide for Web Services</i>
Create GPAs using Fusion Middleware Control and WLST	"Creating and Managing Policy Sets" section in the <i>Oracle Fusion Middleware Security and Administrator's Guide for Web Services</i>

Table 4–7 (Cont.) Common GPA Operations

Operation	Reference
Disable GPAs	"Enabling and Disabling a Policy Set" section in the <i>Oracle Fusion Middleware Security and Administrator's Guide for Web Services</i>
Determine what policies are enforced when both LPA and GPA are defined	"Calculating the Effective Set of Policies" section in the <i>Oracle Fusion Middleware Security and Administrator's Guide for Web Services</i>
View policies attached to a Web service	"Viewing the Policies That are Attached to a Web Service" section in the <i>Oracle Fusion Middleware Security and Administrator's Guide for Web Services</i>
Validate a policy set	"Validating a Policy Set" section in the <i>Oracle Fusion Middleware Security and Administrator's Guide for Web Services</i>

4.10.3 Web Services Security Profiles

Oracle Web Services Manager supports three Web services security profiles:

- AuthN Profile
- SSL Profile - provides transport-level security
- Message Security Profile - provides message-level security

Out-of-the box, Oracle Fusion applications are provisioned with the AuthN profile. You can move your deployment to a different profile if needed.

4.10.4 Key Exchange with the Domain Hosting Oracle Identity Manager

During provisioning, all Oracle Fusion Applications domains are set up to use a common keystore and credential store, whereas the Oracle Identity Management domain (which includes Oracle Identity Manager) uses a separate keystore, and stores credentials in a logical domain. Provisioning does not set up trust between these keystores. This section explains how to exchange trust with the Oracle Identity Manager domain, enabling Web services security when this domain is involved.

For additional background and information on the certificate exchange needed to set up Web services security trust, see the *Oracle Fusion Applications Security Hardening Guide*.

Exporting a Keystore Alias from Application Domain

Take these steps to export the keystore alias:

1. Navigate to the `DOMAIN_HOME/config/fmwconfig` directory of the domain.
2. Run the `keytool` command to export the alias into a file called `orakey.cert` using syntax like in this example:

```
JAVAHOME/bin/keytool -exportcert -alias orakey -file orakey.cert -keystore default-keystore.jks -storepass keystore-password
```

This command creates a file called `orakey.cert` containing the exported `orakey` alias.

When using this command, specify the alias name and keystore password applicable to your environment.

Importing a Keystore Alias into Oracle Identity Manager Domain

Take these steps to import the keystore alias:

1. Copy the file generated by the export procedure (`orakey.cert` in the export example) into the `DOMAIN_HOME/config/fmwconfig` directory of the Oracle Identity Manager domain where you wish to import the alias.
2. Run the `keytool` command using the following syntax to import the certificate:

```
JAVA_HOME/bin/keytool -importcert -alias orakey -file orakey.cert -keystore
default-keystore.jks -storepass keystore-password
```

This command imports the alias into the Oracle Identity Manager domain's keystore.

Similar steps are used to export the Oracle Identity Manager key and import it into the other domain.

4.10.5 Web Services Security Hardening

For detailed instructions about Web Services Security hardening, see the "Locking Down Web Services: Points to Consider" section in the *Oracle Fusion Applications Security Hardening Guide*.

4.11 Securing Oracle Fusion Middleware Products

This section describes security guidelines recommended for Oracle Fusion Middleware products in the context of Fusion Applications. These administrative tasks, typically carried out with Oracle Authorization Policy Manager, Oracle Enterprise Manager Fusion Middleware Control, or WLST scripts, apply only to a specific product.

4.11.1 Administrative Tasks and Features Specific to the IDCCS Stripe

This section describes administrative tasks and features specific to Oracle Universal Content Management (UCM) in the IDCCS stripe of the policy store, and they do not apply to any other stripe.

The grants specified in the stripe IDCCS must conform to one of the grants described in [Section 4.11.1.1](#).

[Section 4.11.1.2](#) explains why not to delete groups or accounts in that stripe.

4.11.1.1 Grants Supported by UCM

Oracle Authorization Policy Manager allows an administrator to specify grants to accounts and groups in a number of combinations, but not all these combinations are supported by UCM. The following list identifies the grants in the IDCCS stripe that are supported and unsupported.

- Grants to Accounts
 - User to account resource: supported.
 - Enterprise role to account resource: supported.
 - Application role to account resource: supported.
- Grants to Security Group Resources
 - User to security group resource: not supported.
 - Enterprise role to security group resource: not supported.

- Application role to security group resource: supported.
- Grants to Entitlements
 - None supported.

Even though Oracle Authorization Policy Manager allows an administrator to define a non-supported grant in for UCM, such grants are ignored at runtime.

4.11.1.2 Security Groups and Accounts Associated with Documents

Oracle Authorization Policy Manager allows deleting a security group or account associated with a document; but UCM does not allow deleting such artifacts.

Security administrators should be cautious never to delete a security group or account associated with a document stored in the Oracle Content Server. If, however, such deletion takes place accidentally, the deleted artifact will automatically reappear.

4.12 Extracting Data from an LDAP- Based Store to a File

This section describes the use of the OPSS script `migrateSecurityStore` to extract the contents of one stripe and save it in an XML file.

To extract the policies in a given application stripe from an LDAP-based store, proceed as follows:

1. Create a file with the following content and save it as, for example, `exported-jazn-data.xml` in the same directory where your `jps-config.xml` file is located:

```
<?xml version = '1.0' encoding = 'UTF-8' standalone = 'yes'?>
<jazn-data xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"

xsi:noNamespaceSchemaLocation="http://xmlns.oracle.com/oracleas/schema/jazn-dat
a-11_0.xsd">
  <jazn-realm default="jazn.com">
    <realm>
      <name>jazn.com</name>
    </realm>
  </jazn-realm>
</jazn-data>
```

2. Copy of your `jps-config.xml` file to, for example, `copy-jps-config.xml`.
3. Edit the file `copy-jps-config.xml` to contain an instance and two contexts like the following:

```
<serviceInstance location="./exported-jazn-data.xml"
  provider="policystore.xml.provider"
  name="export.xml" />

<jpsContext name="ldap_source">
  <serviceInstanceRef ref="policystore.ldap" />
</jpsContext>

<jpsContext name="xml_target">
  <serviceInstanceRef ref="export.xml" />
</jpsContext>
```

4. Run the script `migrateSecurityStore` as in the following invocation (the arguments are written in separate lines for the sake of clarity only):

```

migrateSecurityStore(type="appPolicies",
                    configFile="./copy-jps-config.xml",
                    src="ldap_source",
                    dst="xml_target",
                    srcApp="hcm")

```

The above illustrates the extraction of the contents of the application stripe `hcm` to the XML file `exported-jazn-data.xml`. To extract the contents of any other stripe, specify the appropriate stripe name in the argument `srcApp`. No more than one stripe can be extracted at a time.

Note that in the above example, both `copy-jps-config.xml` and `exported-jazn-data.xml` are assumed to be located in the same directory.

4.13 Customizing Security from Installation to Deployment

This section describes the typical flow of application security data from the application installation to the application deployment to a production environment, and, in particular, how to handle GUIDs in application policies throughout that flow; its phases, performed by developers and administrators, are explained in the following sections:

- [Installing a New Oracle Fusion Application](#)
- [Customizing and Testing Security with Oracle JDeveloper](#)
- [Migrating to a Staging Environment](#)
- [Migrating to a Production Environment](#)

4.13.1 Installing a New Oracle Fusion Application

In this phase, an administrator installs a new Oracle Fusion environment whereby the following application-specific artifacts are installed and provisioned:

- An identity store, containing users and security enterprise roles.
- An LDAP-based policy store, containing functional policies.
- An Oracle Fusion Data Security policy store, containing database resources and data security policies.

4.13.2 Customizing and Testing Security with Oracle JDeveloper

In this phase, a developer uses Oracle JDeveloper to customize and test functional security artifacts in the installed application. First, the developer needs an XML version of the application stripe in the installed LDAP-based store, the XML format been required by JDeveloper.

To obtain this file, the administrator proceeds as follows:

1. Runs the OPSS script `migrateSecurityStore` with the argument `preserveAppRoleGuid` set to `True` to export the contents of the application stripe in the LDAP-based policy store to the XML file (called `jazn-data.xml` throughout this section). For details on this procedure, see [Section 4.12](#).

Important Note: The use of the argument `preserveAppRoleGuid` is strictly restricted to the scenario described above; in particular, that argument should not be used with `migrateSecurityStore` in a production environment or when migrating from a staging to a production environment.

2. Edits the generated file `jazn-data.xml` to replace the group principal for enterprise roles with the group principal required by Oracle JDeveloper. Specifically, it runs a command that replaces, in that file, every instance of the string `weblogic.security.WLSGroupImpl` with the string `oracle.security.jps.internal.core.principals.JpsXmlEnterpriseRoleImpl`.
3. Delivers the edited `jazn-data.xml` to the developer, who copies it to the directory `<jdevapphome>/src/META-INF`.

Then, within Oracle JDeveloper, the developer proceeds as follows:

1. Defines a data source to point to the installed database security policy store.
2. Uses the security policy overview editor to customize function security (in the XML-based policy store), as documented in the "Implementing Function Security" chapter of the *Oracle Fusion Applications Developer's Guide*.
3. Creates test users and test groups, and run the application within Oracle JDeveloper. When the application is run, the WebLogic Server integrated in Oracle JDeveloper, merges `jazn-data.xml` with the domain file-based store `system-jazn-data.xml`.
4. Once customization and testing within Oracle JDeveloper are completed, hands the file `jazn-data.xml` back to the system administrator.

4.13.3 Migrating to a Staging Environment

In this phase, the administrator proceeds as follows:

1. If the staging environment already has a stripe for the application and that stripe is to be preserved, merges the file `jazn-data.xml` with the staging application policy store, using Oracle Authorization Policy Manager as described in the "Upgrading Oracle Fusion Applications Policies" chapter of the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)*. The relevant artifacts involved in this merge are the baseline store, represented by the original XML file `jazn-data.xml`; the production store, represented by the staging policy store stripe; and the patch store, represented by the new, customized `jazn-data.xml` file

Otherwise, if the application stripe is not present or is present but need not be preserved, the administrator uses the OPSS script `migrateSecurityStore` with the argument `overWrite` set to `TRUE` to migrate application policies in the file `jazn-data.xml` to the staging policy store.

2. Reconciles GUIDs in the Oracle Fusion Data Security policy store with GUIDs in application roles by running the command `DSDDataMigrator` as described in [Section 4.4.3](#). This operation does not modify the file `jazn-data.xml` but updates GUIDs in data security policies.

Important Note: When migrating to a staging environment, reconciling GUID's is required *only* before deploying the application in the staging environment that will merge policies (from a developer's file-based policy store) containing newly introduced custom application roles that are also referenced in data security policies.

3. Packs the file `jazn-data.xml` with the application EAR and deploys the application to the staging environment with the following system property set:

```
-Djps.deployment.handler.disabled="true"
```

The above setting disables the OPSS listeners and ensures that when the EAR is deployed, domain policies do not get overwritten or appended with those in the packed `jazn-data.xml` file.

Note: Typically, the application deployment is set so that users and groups are not migrated during deployment. For details about this setting within Oracle JDeveloper, see the "Implementing Function Security" chapter in the *Oracle Fusion Applications Developer's Guide*.

4.13.4 Migrating to a Production Environment

In this phase, the administrator proceeds as follows:

1. Migrates the staging policy store to the production policy store in one of the following two ways:
 - If the production environment has not undergone changes since the customization was tested in the staging environment, the administrator moves the staging policy store to the production store as described in [Section 17.3](#).
 - Otherwise, if the production policy store has undergone changes since the customization was tested in the staging environment and those changes must be preserved, the administrator uses Oracle Authorization Policy Manager as described in the *Upgrading Oracle Fusion Applications Policies* chapter of the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)* to merge the customized, pillar-level `jazn-data.xml` with the production policy store.
2. Reconciles GUIDs in the Oracle Fusion Data Security policy store with GUIDs in the OID LDAP-based policy store by running the command `DSDataMigrator` as described in [Section 4.4.3](#).

Important Note: When migrating to a production environment, reconciling GUID's is required *only* if application roles or data security have been modified in the production environment..

3. Packs the file `jazn-data.xml` with the application EAR and deploys the application to the production environment with the following system property set:

```
-Djps.deployment.handler.disabled="true"
```

The above setting disables the OPSS listeners and ensures that when the EAR is deployed, domain policies do not get overwritten or appended with those in the packed `jazn-data.xml` file.

Managing Oracle Enterprise Scheduler Service and Jobs

This chapter describes how to configure and manage Oracle Enterprise Scheduler and Oracle Enterprise Scheduler jobs. Oracle Enterprise Scheduler enables scheduling and running jobs within a particular time frame, or workshift, using rules to create work assignments. Oracle Enterprise Manager Fusion Applications Control allows you to define, control and manage Oracle Enterprise Scheduler job metadata, including job definitions, job requests, job sets (a collection of job requests), incompatibilities (job definitions and job sets that cannot run at the same time for a given application) and schedules governing the execution of job requests.

This chapter includes the following topics:

- [Introduction to Managing Oracle Enterprise Scheduler](#)
- [Navigating to the Oracle Enterprise Scheduler Administration Pages](#)
- [The Scheduling Service Home Page](#)
- [Configuring Oracle Enterprise Scheduler](#)
- [Searching for Configuration Changes to Oracle Enterprise Scheduler in Cloud Control](#)
- [Viewing General Oracle Enterprise Scheduler Instance Information](#)
- [Starting and Stopping Oracle Enterprise Scheduler Service Components](#)
- [Managing Oracle Enterprise Scheduler Requests](#)
- [Monitoring Oracle Enterprise Scheduler Service Activity](#)
- [Managing Purge Policies in Fusion Applications Control](#)
- [Managing Job Metadata](#)
- [Managing Work Assignments and Workshifts](#)
- [Managing Oracle ADF Connections](#)
- [Managing Application Properties](#)
- [Managing Application and Role Policies](#)
- [Managing Oracle Enterprise Scheduler Web Services](#)
- [Managing Logging for Oracle Enterprise Scheduler](#)
- [Managing Alerts for Oracle Enterprise Scheduler in Cloud Control](#)

5.1 Introduction to Managing Oracle Enterprise Scheduler

Oracle Enterprise Scheduler provides the ability to define, schedule and run different types of job, including: Java, PL/SQL, spawned jobs and binary scripts. You can run jobs on demand, or schedule them to run in the future.

Oracle Enterprise Scheduler provides scheduling services for the following purposes:

- Distributing job request processing across a grid of application servers
- Running Java, PL/SQL, spawned and binary jobs
- Processing multiple jobs concurrently

Oracle Enterprise Manager Fusion Applications Control (Fusion Applications Control) allows you to start and stop, monitor, configure and manage Oracle Enterprise Scheduler services, components and job requests.

The main Oracle Enterprise Scheduler page provides an overview of the status of scheduler components, the top running and completed scheduled job requests, as well as a performance summary of scheduled job requests. You can further monitor activity and diagnose problems by examining Oracle Enterprise Scheduler logs and comparing current to historical performance data.

Drilling down to the main components of Oracle Enterprise Scheduler, you can configure job request processors and dispatchers. A job request processor is bound to a particular Oracle Enterprise Scheduler server, and is responsible for allocating threads for job requests. A job request dispatcher polls for job requests.

Once the Oracle Enterprise Scheduler components have been configured, you can define work allocation and purge policies. Work allocation definitions allow you to configure windows of time during which jobs can run (workshifts), rules that govern their execution and bindings to a particular server and request processor (work assignments). Purge policies allow you to define for a particular request processor the conditions for retaining and purging job requests associated with that request processor. You can define a schedule for managing the execution of work assignments, workshifts and purge policies.

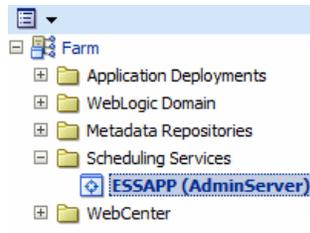
Oracle Enterprise Manager Fusion Applications Control allows you to define, control and manage Oracle Enterprise Scheduler job metadata, including job definitions, job requests, job sets (a collection of job requests), incompatibilities (job definitions and job sets that cannot run at the same time for a given application) and schedules governing the execution of job requests.

Oracle Enterprise Manager Cloud Control (Cloud Control) allows you to manage a subset of Oracle Enterprise Scheduler components, including job requests, schedules, alerts, as well as the starting and stopping of request processors and dispatchers.

5.2 Navigating to the Oracle Enterprise Scheduler Administration Pages

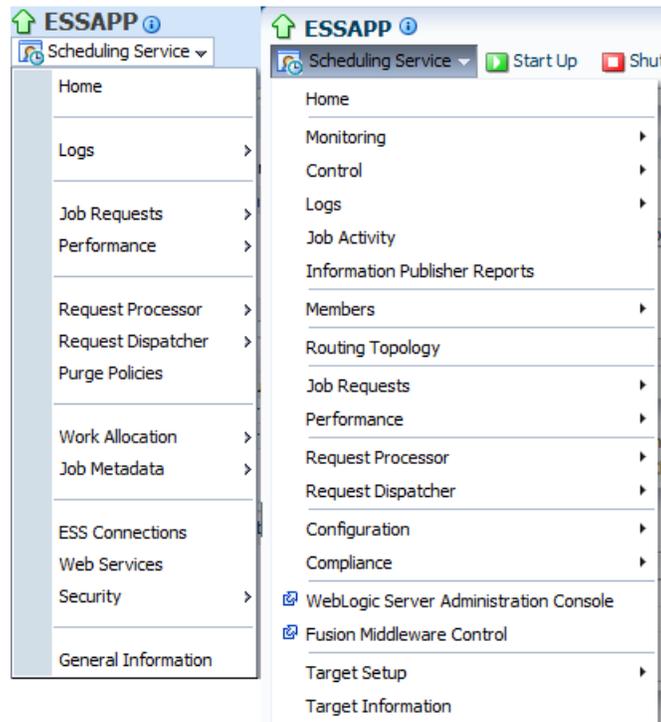
To navigate to Oracle Enterprise Scheduler administrative tasks in Fusion Applications Control or Cloud Control:

1. From the navigation pane, expand the farm and then **Scheduling Services**, and select the Oracle Enterprise Scheduler component.



The Scheduling Service home page displays. See [Section 5.3](#) for further information about the contents of the home page.

2. Select the **Scheduling Service** menu.



The **Scheduling Service** menu displays the following options.

Menu Option	Description
Home	Displays the Scheduling Service home page. For more information about the contents of this page, see Section 5.3 .
Control	Displays the following options: <ul style="list-style-type: none"> ■ Start Up: This option runs the Oracle Enterprise Scheduler server. For more information, see Section 5.7. ■ Shut Down: This option shuts down the Oracle Enterprise Scheduler server. For more information, see Section 5.7.

Menu Option	Description
Logs	<p>The View Log Message option displays the Log Messages page for viewing the contents of event log files. For more information, see Section 5.17.</p> <p>The log files are stored in the following directories:</p> <ul style="list-style-type: none"> ■ <code>DOMAIN_HOME/servers/SERVER_HOME/logs</code> ■ <code>DOMAIN_HOME/servers/SERVER_HOME/logs/apps</code> ■ <code>DOMAIN_HOME/servers/SERVER_HOME/logs/owsm/msglogging/diagnostic.log</code> ■ <code>DOMAIN_HOME/servers/SERVER_HOME/sysman/log/emoms.log</code>
Job Requests	<p>Displays the following options:</p> <ul style="list-style-type: none"> ■ Search Job Requests: Allows you to search for job requests according to particular parameters. For more information, see Section 5.8.2. ■ Submit Job Request: Allows you to submit an execution request for a particular job. For more information, see Section 5.8.1. ■ Define Schedules: Allows you to create a schedule for submitting job requests. For more information, see Section 5.8.1.2.
Performance	<p>Displays the following options:</p> <ul style="list-style-type: none"> ■ Service Summary: This option displays a brief overview of the service performance. For more information, see Section 10.4.2.1. ■ Current Activity: This option displays current Oracle Enterprise Scheduler activity. For more information, see Section 10.4.2.2. ■ Historical Reports: This option displays a chart or table depicting previous Oracle Enterprise Scheduler data. For more information, see Section 10.4.2.3.
Request Processor	<p>Displays the following options:</p> <ul style="list-style-type: none"> ■ Start/Stop: This option causes the request processor to start or stop running. For more information, see Section 5.7. ■ Configure: This option displays the Configure Request Processor page, which allows you to enable the request processor, create work assignment bindings and so on. For more information, see Section 5.4.
Request Dispatcher	<p>Displays the following options:</p> <ul style="list-style-type: none"> ■ Start/Stop: This option causes the request dispatcher to start or stop running. For more information, see Section 5.7. ■ Configure: This option displays the Configure Request Dispatcher page, which allows you to enable the request dispatcher, manage polling intervals, and so on. For more information, see Section 5.4.
Purge Policies	<p>Allows you to configure and manage job request purge policies including job request purge criteria, retention criteria and purge policy schedules. For more information, see Section 5.10. Available in Fusion Applications Control only.</p>

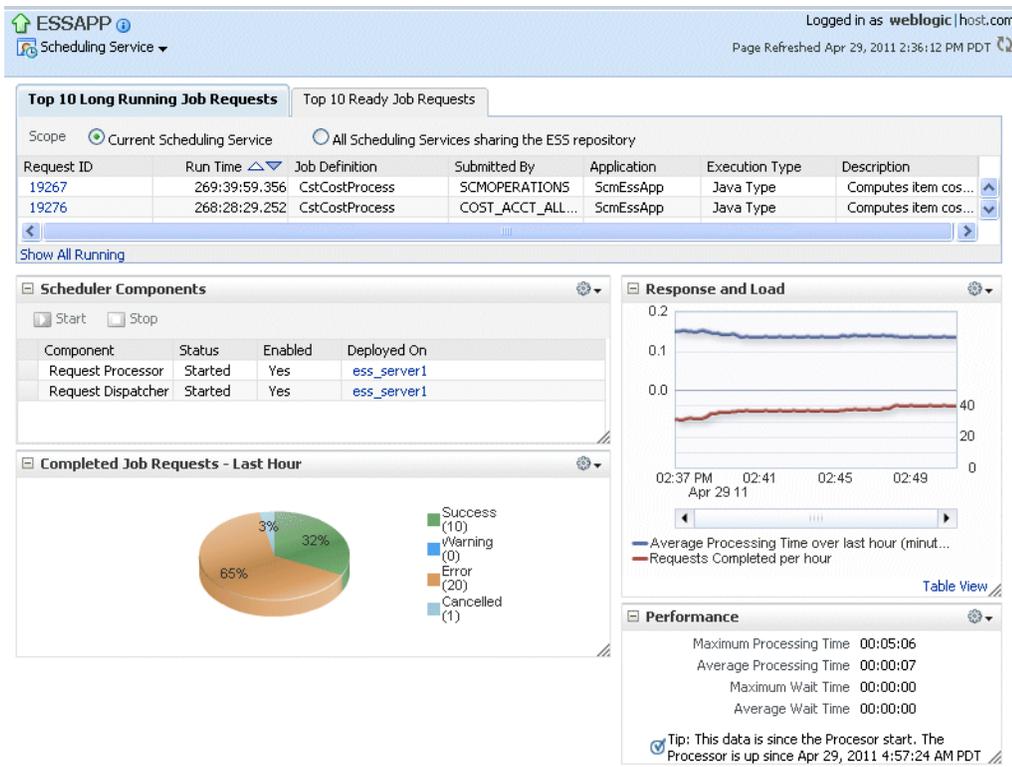
Menu Option	Description
Work Allocation	<p>Displays the following options:</p> <ul style="list-style-type: none"> ■ Work Assignments: This option displays the Work Assignments page which allows you to create, edit and delete work assignments. For more information, see Section 5.12. ■ Workshifts: This option displays the Workshifts page which allows you to create, edit and delete workshifts. For more information, see Section 5.12. ■ Schedules: This option displays the Schedules page which allows you to create, edit and delete schedules for work assignments, workshifts and job request purge policies. For more information, see Section 5.12. <p>Available in Fusion Applications Control only.</p>
Job Metadata	<p>Displays the following options:</p> <ul style="list-style-type: none"> ■ Job Definitions: This option displays the Job Definitions page, which allows you to manage job definitions. See Section 5.11.1 for further information. ■ Job Sets: This option displays the Job Sets page, which allows you to create and edit job sets for a given application. For more information, see Section 5.11.2. ■ Incompatibilities: This option displays the Incompatibilities page, which allows you to create and edit incompatibilities for a given application. For more information, see Section 5.11.3. <p>Available in Fusion Applications Control only.</p>
ESS Connections	<p>Displays a page that enables defining connections to Oracle ADF. For more information, see Section 5.13.</p> <p>Available in Fusion Applications Control only.</p>
Web Services	<p>Displays a page that provides an overview of the web services in the application including web service endpoints and application-level metrics. For more information, see Section 5.16.</p> <p>Available in Fusion Applications Control only.</p>
Security	<p>Displays the following options:</p> <ul style="list-style-type: none"> ■ Application Policies: This option displays the authorization policies that an application relies upon for controlling access to its resources. See Section 5.15 for further information. ■ Application Roles: This option displays the roles used by security aware applications specific to the application. See Section 5.15 for further information. <p>Available in Fusion Applications Control only.</p>
General Information Target Information	<p>Displays general details about the Oracle Enterprise Scheduler instance. Known as Target Information in Cloud Control. For more information, see Section 5.6.</p>

5.3 The Scheduling Service Home Page

The Scheduling Service home page provides an overview of the performance of Oracle Enterprise Scheduler components and jobs, including component status, the number of completed job requests in the last hour, as well as the processing times for running jobs.

You can use this page as a starting point for monitoring and administering Oracle Enterprise Scheduler. [Figure 5–1](#) shows a portion of the Scheduling Service home page.

Figure 5–1 Scheduling Service Home Page



These pages contain the following regions:

- [Top Ten Long Running Requests and Top Ten Ready Job Requests Regions](#)
- [Scheduler Components Region](#)
- [Completed Requests Region](#)
- [Response and Load Region](#)
- [Performance Region](#)
- [Monitoring and Diagnostics Region](#) (Cloud Control only)

5.3.1 Top Ten Long Running Requests and Top Ten Ready Job Requests Regions

The Top Ten Long Running Requests region displays the top ten long running scheduled job requests, including request ID, job run time, job definition used, executing application, job execution type and description. You can set the scope of the top ten long running requests displayed to the current scheduling service only, or all scheduling services sharing the Oracle Enterprise Scheduler repository.

The Top Ten Ready Job Requests tab displays the top ten scheduled job requests that await execution. The tab displays the same information as the Top Ten Long Running Requests tab, except the wait time displays for each job rather than the run time.

5.3.2 Scheduler Components Region

The Scheduler Components region displays the components of Oracle Enterprise Scheduler, including the job request processor and dispatcher. The tab displays the status of each component, the name of the server to which it is deployed and whether or not the component is enabled. You can start and stop each component as required.

5.3.3 Completed Requests Region

The Completed Requests region displays the scheduled jobs completed within the last hour.

5.3.4 Response and Load Region

The Response and Load region displays performance monitoring statistics regarding the time required to process to job requests.

[Table 5–1](#) describes the performance monitoring statistics in the **Response and Load** region.

Table 5–1 Response and Load Statistics

Statistic	Description
Average Processing time for Last Hour	This metric specifies the average time required to process jobs during the last hour.
Requests Completed Per Hour	This metric specifies the number of scheduled job requests completed within the last hour.

5.3.5 Performance Region

The Performance region displays performance data for job requests, such as processing times and wait times.

[Table 5–2](#) describes the performance monitoring statistics in the **Performance** region.

Table 5–2 Performance Statistics

Statistic	Description
Maximum Processing Time	This metric specifies maximum amount of time required to process a scheduled job.
Average Processing Time	This metric specifies average amount of time required to process a scheduled job.
Maximum Wait Time	This metric specifies maximum amount of time during which a scheduled job waits before running.
Average Wait Time	This metric specifies the average amount of time during which a scheduled job waits before running.

5.4 Configuring Oracle Enterprise Scheduler

You can run Oracle Enterprise Scheduler as a single instance or as a cluster of servers. Each Oracle Enterprise Scheduler server includes a request processor and dispatcher, both of which must be configured.

Configuring Oracle Enterprise Scheduler involves the following main steps:

1. **Configure a cluster.** Optionally, configure a cluster of Oracle Enterprise Scheduler servers. For more information about configuring a cluster, see the section

"Expanding an Oracle Enterprise Scheduling Service Cluster" in the chapter "Performing Basic Administration Tasks" in *Oracle Fusion Middleware Administrator's Guide for Oracle Enterprise Scheduling Service*.

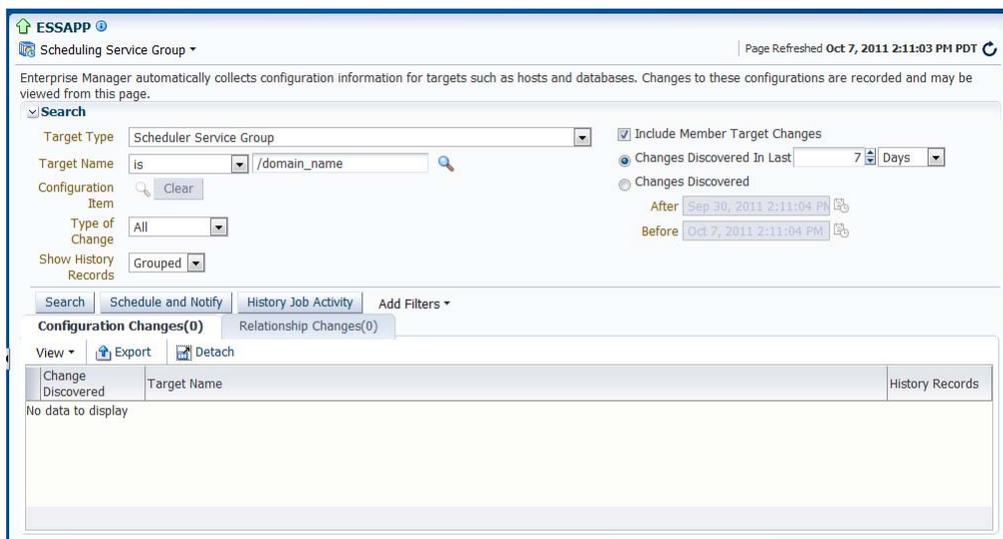
2. **Configure the request processor.** Configure the Oracle Enterprise Scheduler component which receives and manages job requests. For more information about configuring the request processor, see the section "Configuring a Request Processor" in the chapter "Performing Basic Administration Tasks" in *Oracle Fusion Middleware Administrator's Guide for Oracle Enterprise Scheduling Service*.
3. **Configure the request dispatcher.** Configure the Oracle Enterprise Scheduler component that polls the request processor for job requests and dispatches jobs. For more information about configuring the request dispatcher, see the section "Configuring a Request Dispatcher" in the chapter "Performing Basic Administration Tasks" in *Oracle Fusion Middleware Administrator's Guide for Oracle Enterprise Scheduling Service*.

5.5 Searching for Configuration Changes to Oracle Enterprise Scheduler in Cloud Control

You can search for changes made to configurations using the search page in Cloud Control.

To search for changes to configurations to Oracle Enterprise Scheduler in Cloud Control:

1. Display the alerts for Oracle Enterprise Scheduler as described in [Section 5.18.2](#).
2. In the Monitoring and Diagnostics pane, click the number of configuration changes to display the search page for configuration changes to Oracle Enterprise Scheduler alerts.



3. Configure your search as desired.
 - **Target Type:** Select the **Scheduler Service** or **Scheduler Service Group**.
 - **Target Name:** Search for or enter the name of the target you want to search. From the dropdown list, select the criterion by which you want to search for the target name: **contains**, **is**, **is not**, **does not contain**, **starts with**, **ends with**.

- **Type of Change:** From the dropdown list, select the type of configuration change for which you want to search.
 - **Show History Records:** From the dropdown list, select the way you want the records to display, **Grouped** or **Show All**.
 - **Include Member Target Changes:** Select to display configuration changes made to members of the target.
 - **Changes Discovered in Last:** Select to search for changes discovered within a given number of days, weeks or months.
 - **Changes Discovered:** Select to search for changes discovered within a particular time period. Use the **After** and **Before** fields to enter the relevant time period.
4. Click **Search** to search for the changes.

5.6 Viewing General Oracle Enterprise Scheduler Instance Information

Oracle Fusion Applications Control enables viewing information regarding the particular instance of Oracle Enterprise Scheduler. The general information popup window displays the locations of Oracle Fusion Middleware, the domain and the target application, as well as the version of Oracle Fusion Middleware currently running and the URI of the hosting server.

To view general information about Oracle Enterprise Scheduler:

1. From the Scheduling Service menu, select one of the following:
 - In Fusion Applications Control, select **General Information**.
 - In Cloud Control, select **Target Information**.
2. Oracle Fusion Applications Control displays the following information.
 - **Target Name:** The location of the target application, in this case, Oracle Enterprise Scheduler, running on the domain.
 - **Version:** The version number of Oracle Fusion Middleware.
 - **Middleware Home:** The location of the Oracle Fusion Middleware directory.
 - **Domain Home:** The full path of the domain.
 - **Host:** The URI of the hosting server.

5.7 Starting and Stopping Oracle Enterprise Scheduler Service Components

You can start up and shut down an instance of Oracle Enterprise Scheduler from Oracle Fusion Applications Control.

The following components can be started and stopped:

- Oracle Enterprise Scheduler instance.
- Job request processor and dispatcher.

Note: Stopping an Oracle Enterprise Scheduler instance or component is not recommended. Shutting down an Oracle Enterprise Scheduler component does not stop job requests from accumulating in the queue.

For more information about starting and stopping Oracle Enterprise Scheduler components, see the section "Starting and Stopping Oracle Enterprise Scheduling Service Components" in the chapter "Performing Basic Administration Tasks" in *Oracle Fusion Middleware Administrator's Guide for Oracle Enterprise Scheduling Service*.

5.8 Managing Oracle Enterprise Scheduler Requests

A job request is a request sent from a deployed application to Oracle Enterprise Scheduler to execute a job. Jobs are configured using job definitions, which specify the parameters of the job and associate them with a particular application. A job set is a collection of job requests that is grouped together.

This section contains the following topics:

- [Submitting an Oracle Enterprise Scheduler Job Request](#)
- [Searching for Oracle Enterprise Scheduler Job Requests](#)

For more information about searching for jobs, viewing job request details, holding and canceling jobs, see the section "Searching for Oracle Enterprise Scheduling Service Job Requests" in the chapter "Managing Oracle Enterprise Scheduler Requests" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Enterprise Scheduling Service*.

5.8.1 Submitting an Oracle Enterprise Scheduler Job Request

You can submit a job request for a particular application using a job definition and a schedule by which you want the job to run.

This section contains the following topics:

- [Creating a Job Request](#)
- [Specifying a Job Request Schedule](#)

5.8.1.1 Creating a Job Request

Creating a job request involves selecting a job definition for the job request, then selecting or creating a schedule.

You may want to configure system properties for the job request, such as the number of retries to attempt in the event of an execution error and a time out value for the job.

To submit a job request:

1. Use a pre-existing job definition, or create a new job definition, as described in [Section 5.11.1.2](#).
2. Navigate to the Submit job request page by clicking the Scheduling Service menu and selecting **Job Requests**. Then select **Submit Job Request**.
3. Select the application for which you want to submit the job request.
From the application dropdown list, select the J2EE application deployment name.
4. Under Job Request Details, enter details about the job request in the **Details** text field.

5. Select a job definition.
 - a. Under Job Definition, click the search button next to the Job Definition text field.
 - b. Search for the required job definition.

In the Job Definition Name text field, enter the name of the job definition you want to find.

Click the arrow next to the text field to search for the job definition.
 - c. From the search results table, select the job definition name.
6. In the Parameters region, optionally define any parameters you want to use with the scheduled job request.
 - a. In the Parameters region, click **Add** to add a parameter.

Fill in a name and value for the parameter.
 - b. You can set a value for any of the system properties available. [Table 5-3](#) lists the names and descriptions of system properties.

For example, you may want to configure the number of permissible retries for the job (`SYS_retries`), or a time out value (`SYS_request_timeout`).
7. Select a schedule by which the job is to run.
 - **Once:** Select a single date and time by clicking the calendar icon.
 - **Use existing schedule:** Click the search button to search for and select an existing schedule.
 - **Specify schedule:** Create your own schedule on the fly. Follow the instructions in [Section 5.8.1.2](#).

5.8.1.2 Specifying a Job Request Schedule

In the Submit Job Request page, selecting Specify a Schedule enables configuring your own schedule for a job request. You can configure a recurring job request using a job request schedule. Alternatively, you can configure a job request to run immediately or before a specified end date. Use a pre-existing job definition, or create a new job definition, as described in [Section 5.11.1.2](#).

To specify a schedule:

1. In the Start section, select the starting date and time for the job request.

Select **Immediately** to submit the job request right away.

Select **Later** to select a subsequent start date.
2. In the Repeating section, select the desired frequency of the job request submission schedule.

From the Repeat dropdown list, select **Every N minutes**, **Every N hours**, **Every N days** or **Every N weeks**.

In the Frequency text field, enter the number of minutes, hours, days or weeks.
3. In the End Date section, select the date and time at which the schedule job is to end.
 - **No End Date:** Select to specify a continuous schedule for the job request.
 - **Specified End Date:** Select to specify a particular end date for the schedule. Click the calendar icon next to the Date text field and select a date and time.

Optionally, save the schedule by selecting the **Save Schedule** checkbox and entering a name for the schedule in the **Name** text field. Optionally, enter the package name and description.

5.8.2 Searching for Oracle Enterprise Scheduler Job Requests

The Request Search page enables searching for job requests. Using the Request Search page, you can:

- Conduct a simple search that returns a list of job request details, including job request ID, executing application, job request status, and so on.
- Conduct an advanced search that returns the same information as the simple search, as well as the date and time of execution, the run time or wait time of the job request, the number of retries and any error type that may have occurred during execution.
- Modify the column display in the search results table.

This section contains the following topics:

- [Searching for a Job Request Using Simple Search](#)
- [Searching for a Job Request Using the Advanced Search Feature](#)
- [Configuring the Display of Columns in the Search Results Table](#)

5.8.2.1 Searching for a Job Request Using Simple Search

Simple search allows you to find a job request according to particular criteria such as job request ID, related application, job request status, and so on, or by any one of a number of pre-configured quick searches.

To search for a job request:

1. Navigate to the Request Search page by clicking the Scheduling Service menu and selecting **Job Requests**. Then select **Search Job Requests**.
2. Select the scope of the job request search by selecting one of the following options:
 - **Current Scheduling Service**: Select this option to search for job requests submitted only to the scheduling service with which you are currently working.
 - **All Scheduling Services sharing the ESS repository**: Select this option to search for job requests submitted to all scheduling services sharing the repository, for example all scheduling services in a cluster of scheduling services.
3. To run a fast search, from the **Quick Search** dropdown list, select a pre-configured search option as shown in the following list.
 - Requests submitted in the last hour
 - Pending requests submitted in the last 24 hours
 - Errored requests submitted in the last 24 hours
 - All running requests
 - All pending requests
 - Requests currently being retried
 - Requests retried in the last 24 hours

- Requests that resulted in System Error in the last 24 hours
- Requests that resulted in Business Error in the last 24 hours

To run a regular job request search, skip this step.

4. Select the criteria by which to search for job requests.
 - **Request ID:** Enter the ID of the job request for which you want to search.
 - **Application:** From the Application dropdown list, select the name of the application related to the job request for which you want to search. Alternatively, select **All** to search for job requests in all applications.
 - **Status:** Select the status of the job request for which you want to search. Alternatively, select **All** to search for job requests with all statuses. Statuses are listed in the following table.

Statuses

WAIT	VALIDATION_FAILED	EXPIRED
READY	RUNNING	SCHEDULE_ENDED
PENDING_VALIDATION	BLOCKED	WARNING

- **Execution Type:** Select `JAVA_TYPE`, `SQL_TYPE` or `PROCESS_TYPE`.
 - **Submitted:** From the dropdown list, select the time period in which the job request to be searched has been submitted. Options include **Last Hour**, **Last 24 Hours**, **Last 7 Days**, **Last 31 Days**.
 - **Submitted By:** In the text field, enter the name of the user who submitted the job request you want to find.
 - **Job Definition:** Click the search button next to the text field and select the relevant job definition name.
 - **Work Assignment:** Click the search button next to the text field and select a work assignment from the list.
5. Optionally, conduct an advanced search by clicking the **Advanced** button.
 6. Click **Search** to submit the job request search.

5.8.2.2 Searching for a Job Request Using the Advanced Search Feature

An advanced search is available in the Request Search page by clicking the **Advanced** button.

To search for a job request:

1. Enter a basic search for a job request. Navigate to the Request Search page by clicking the Scheduling Service menu and selecting **Job Requests**. Then Select **Search Job Requests**.
2. Select the scope of the job request search by selecting one of the following options:
 - **Current Scheduling Service:** Select this option to search for job requests submitted only to the scheduling service with which you are currently working.
 - **All Scheduling Services sharing the ESS repository:** Select this option to search for job requests submitted to all scheduling services sharing the

repository, for example all scheduling services in a cluster of scheduling services.

3. Select your basic search criteria. For more information, see [Section 5.8.2.1](#).
4. Click **Advanced** to display the fields for the advanced search.
5. In the Date Range section, configure the date range in which to search for job requests. The date on the left is the beginning date and the date on the right is the end date. For each date, click the calendar icon to the right of the text field to select the date and time.
 - **Submitted between:** Enter the start and end dates during which the job request was submitted.
 - **Scheduled between:** Enter the start and end dates during which the job request is scheduled to run.
 - **Completed between:** Enter the start and end dates during which the job request finished running.
6. In the Run Time/Wait Time section, select the run or wait time of the job request for which you are searching, such as long or short running requests.

- **None:** Select if no run or wait time is to be specified.
- **Long running requests:** Select to search for requests running longer than a specified number of seconds, minutes, hours or days.

In the **Minimum Run Time** text field, enter the lower limit of the time period for which the job request runs. From the dropdown list, select the required unit of time.

- **Short running requests that waited:** Select to search for job requests running longer than a specified period of time and waiting less than a specified period.

In the **Maximum Run Time** text field, enter the upper limit of the time period for which the job request runs. From the dropdown list, select **Seconds**, **Minutes**, **Hours** or **Days**.

In the **Minimum Wait Time** text field, enter the lower limit of the time period during which the job request waits to run. From the dropdown list, select **Seconds**, **Minutes**, **Hours** or **Days**.

- **Waiting requests by time:** Select to search for job requests waiting to run for a specified time period.

In the **Minimum Wait Time** text field, enter the lower limit of the time period during which the job request waits to run. From the dropdown list, select **Seconds**, **Minutes**, **Hours** or **Days**.

In the **Maximum Wait Time** text field, enter the upper limit of the time period during which the job request waits to run. From the dropdown list, select **Seconds**, **Minutes**, **Hours** or **Days**.

7. In the Retry of Failed Runs section, use the Number of Retries dropdown list to select an operator such as equal to, greater than, greater than or equal to, and so on.

In the text field, enter the number of retries.

8. In the Error Type section, use the dropdown list to select the type of error:
 - **Business:** A job ends in a business error when it must abort prematurely due to unforeseen conditions, but is otherwise able to exit with its data in a

consistent state. A job request might end in a business error as a result of an application setup or configuration condition, a functional conflict that requires an early exit or corrupt or inconsistent data. You cannot retry running a job request that ends in a business error.

- **System:** A job ends in a system error when a technical problem occurs from which the job cannot recover, but otherwise exits on its own. Alternatively, the computer running the job crashes. Examples include table space issues and unhandled runtime exceptions. You can retry running a job request that ends in a system error.

9. Click **Search** to submit the job request search.

5.8.2.3 Configuring the Display of Columns in the Search Results Table

After running a search for job requests, you can configure the display of columns in the search results table.

To configure job request search results table display columns:

1. Display the main Oracle Enterprise Scheduler Request search page and display the search interface.
2. Display the list of available column names.
Click the **View** dropdown list and select **Columns**, then **Show More Columns**.
3. Use the arrows between the columns to move column names from the **Visible** to the **Hidden** column, or vice versa.
4. Optionally, reorder the columns by selecting the relevant column names and using the vertical arrows on the right to move them up or down.
5. Alternatively, display all columns by clicking the **View** dropdown list and select **Columns**, then **Show All**.

5.9 Monitoring Oracle Enterprise Scheduler Service Activity

The main Oracle Enterprise Scheduler page displays information regarding the top ten long running scheduled jobs as well as the top ten job requests awaiting execution in two different tabs. By default, only the job requests within the current scheduling service display. However, you can change the scope such that all relevant scheduled jobs running on all scheduling services sharing the Oracle Enterprise Scheduler repository display in the tabs.

Each of the Oracle Enterprise Scheduler tabs include a table that displays a short list of top ten long running and waiting job requests, respectively. The Show All Ready or Show All Running link at the bottom of the region goes to a search page in which it is possible to search for a more comprehensive list of relevant job requests.

Each tab displays the following information about scheduled job requests:

- **Request ID:** The ID associated with the job request.
- **Run Time/Wait Time:** The period of time during which the job request has been running or awaiting execution, respectively.
- **Job Definition:** The job definition associated with the job request.
- **Submitted by:** The name of the user who submitted the job request.
- **Application:** The name of the application with which the job request is associated.
- **Execution Type:** The category of job being executed—Java, SQL or process job.

- **Description:** An optional description of the scheduled job request.

You can also monitor Oracle Enterprise Scheduler job requests related to a given product. For more information about viewing the top scheduled job requests related to a product, see [Section 9.2.4](#).

This section contains the following topics:

- [Viewing Top Ten Long Running Oracle Enterprise Scheduler Requests](#)
- [Viewing Top Ten Oracle Enterprise Scheduler Jobs Ready to Be Executed](#)

5.9.1 Viewing Top Ten Long Running Oracle Enterprise Scheduler Requests

Use the Top Ten Long Running Job Requests tab to view the job requests that have been running for the longest period of time on the Oracle Enterprise Scheduler service. You can alternatively change the scope of the region to display the top ten long running job requests on all scheduling services sharing the same repository.

The job requests displayed will all have a status of `RUNNING`.

To view the top ten long running requests:

1. From the navigation pane, expand the **Scheduling Services** folder and select the Oracle Enterprise Scheduler application.
2. At the top of the Oracle Enterprise Scheduler page, click the **Top 10 Long Running Requests** tab to display the ten job requests that have been running the longest.
3. To display all running job requests, click **Show All Running** at the bottom of the chart. This displays a search form that searches for all requests that are currently running. Click **Search** to display the requests.

5.9.2 Viewing Top Ten Oracle Enterprise Scheduler Jobs Ready to Be Executed

Use the Top Ten Ready Job Requests tab to view the job requests that are awaiting execution on the Oracle Enterprise Scheduler service. You can alternatively change the scope of the region to display the top ten ready job requests on all scheduling services sharing the same repository.

The job requests displayed will all have a status of `READY`.

To view the top ten ready job requests:

1. From the navigation pane, expand the **Scheduling Services** folder and select the Oracle Enterprise Scheduler application.
2. At the top of the Oracle Enterprise Scheduler page, click the **Top 10 Ready Requests** tab to display the ten job requests with the status `READY`.
3. To display all ready job requests, click **Show All Ready** at the bottom of the chart. This displays a search form that searches for all requests with the status `READY`. Click **Search** to display the requests.

5.10 Managing Purge Policies in Fusion Applications Control

Purge policies allow the scheduling service to remove scheduled jobs according to specified criteria. For example, a purge policy might specify the retention of all Java type job requests using a particular job definition submitted executed by a given application for three days. Another purge policy might retain a particular type of job request, say, all SQL job requests in a successful state, for only one day. You can also specify the frequency at which the purge policy is to run.

For more information about managing purge policies, see the section "Managing Purge Policies" in the chapter "Managing Oracle Enterprise Scheduler Requests" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Enterprise Scheduling Service*.

5.11 Managing Job Metadata

Oracle Enterprise Scheduler job metadata refers to the components that form the basis of scheduled job requests.

These include the following:

- **Job definition:** A job definition is the smallest unit of work performed in the context of the application that executes the job. A job definition is defined by a job type, such as a Java or SQL job type.
- **Job set:** A job set is a sequential or parallel set of job steps, where a job step can be a single job or another job set. A job set and each of its job set steps can have additional parameters, the values for which are provided when the job or job set is submitted as a job request.
- **Incompatibility:** An incompatibility allows you to specify job definitions and job sets that cannot run concurrently.

This section contains the following topics:

- [Managing Job Definitions](#)
- [Managing Job Sets](#)
- [Managing Incompatibilities](#)

You can manage job metadata only in Fusion Applications Control.

5.11.1 Managing Job Definitions

The Job Definitions page in Oracle Enterprise Manager Fusion Applications Control allows you to view, create, edit, duplicate, delete and search for job definitions.

[Figure 5–2](#) shows the Manage Job Definitions tab in the Job Definitions page. The tab displays a table of job definitions and their attributes, including name, description, location, execution type and job type.

Figure 5–2 Manage Job Definitions Tab

ESSAPP | Scheduling Service Group | Logged in as weblogic | Page Refreshed Apr 10, 2011 11:47:00 PM PDT

Scheduling Service Home > Job Definitions

Job Definitions

Select the application (J2EE application deployment name) for which you want to view the job definitions.

Application:

LedgerApp

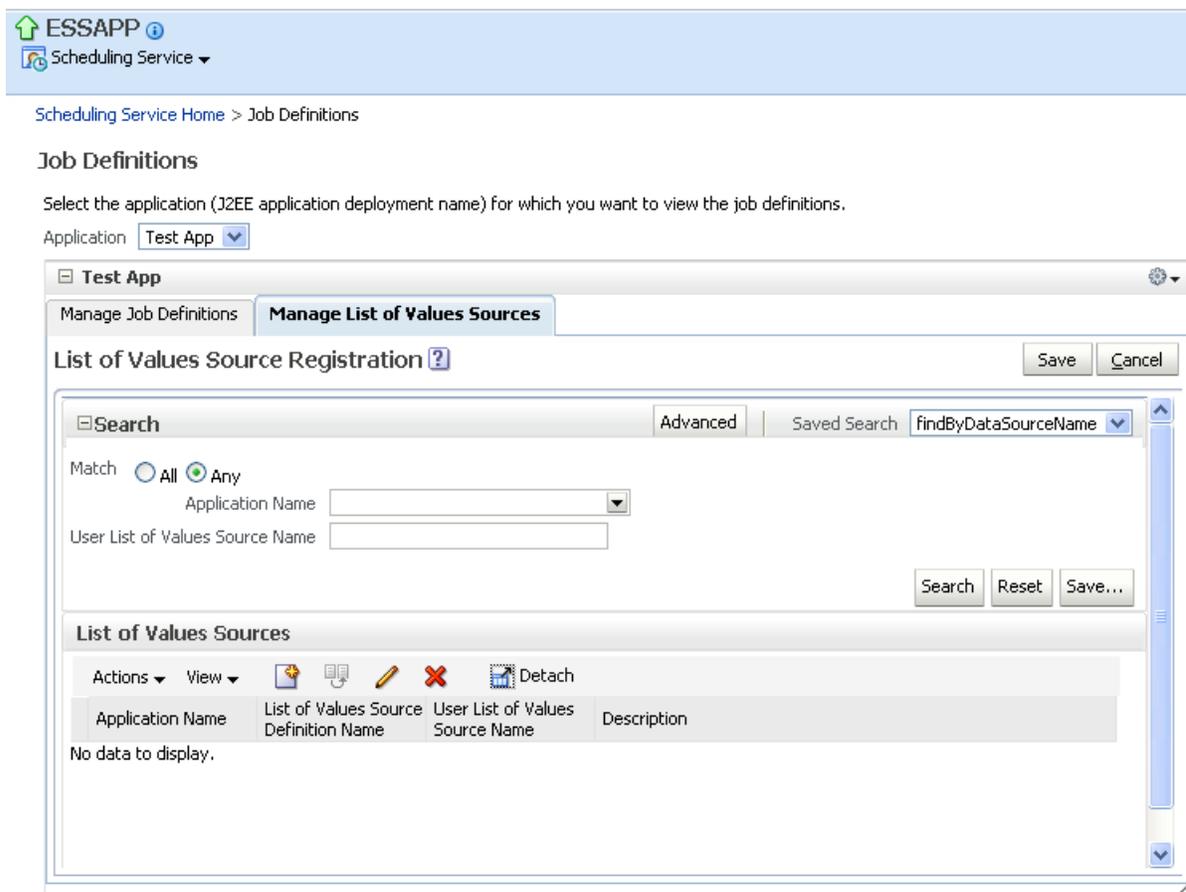
Manage Job Definitions | Manage List of Values Sources

Manage Job Definitions ?

Name	Display Name	Description	Path	Execution Type	Job Type
*APCSTTRF	Transfer Costs to Cost Management	Transfers cost	/oracle/apps/€	Java	JavaJobType
*APINVSEL	Payables Selected Installments Selection	Lists all invoice	/oracle/apps/€	Java	BIPJobType
*APMACR	Create Mass Additions	Transfers capit	/oracle/apps/€	PL SQL	PlsqlJobType
*APOPITRN	Open Items Revaluation Report	Lists changes in	/oracle/apps/€	Java	BIPJobType
*APTCMT	US 1099 Electronic Media Report	Generates sum	/oracle/apps/€	Java	BIPJobType
*APTERIAW	Terminate Approval Task Flow	Terminates the	/oracle/apps/€	PL SQL	PlsqlJobType
*APWHTXLT	Payables Withholding Tax Letter	Creates a withl	/oracle/apps/€	Java	BIPJobType
*APXAPRVL	Validate Payables Invoice	Checks the ma	/oracle/apps/€	Java	BIPJobType
*APXATR	Apply Missing Conversion Rates	Applies missing	/oracle/apps/€	PL SQL	PlsqlJobType
*APXCRRCR	Payables Cash Requirement Report	Forecasts imme	/oracle/apps/€	Java	BIPJobType
*APXIAWRE	Initiate Approval Task Flow	Initiates the ap	/oracle/apps/€	PL SQL	PlsqlJobType
*APXIIMPT	Import Payables Invoices	Creates Payabl	/oracle/apps/€	Java	BIPJobType
*APXINAGE	Payables Invoice Aging Report	Lists unpaid in	/oracle/apps/€	Java	BIPJobType
*APXINDTL	Payables Discounts Taken and Lost Repo	Identifies paym	/oracle/apps/€	Java	BIPJobType

The Job Definitions page in Oracle Enterprise Manager Fusion Applications Control also allows you to register sources for lists of values related to job definitions and the applications that run job definitions. Figure 5–3 shows the Manage List of Values Sources tab in the Job Definitions page.

Figure 5–3 Manage List of Values Sources Tab



This section contains the following topics:

- [Viewing Job Definitions](#)
- [Creating a Job Definition](#)

5.11.1.1 Viewing Job Definitions

You can view the job definitions created for a given application. The table of job definitions displays details about the jobs related to an application such as the job definition name, the full path to which it is saved, the job type, and so on.

Note: In the Job Definitions page, an asterisk displays next to those jobs in the list that are pre-seeded, and are not custom jobs. The asterisk also displays next to pre-seeded jobs that you have customized.

To display job definitions:

1. From the navigation pane, expand the **Scheduling Services** folder and select the Oracle Enterprise Scheduler application.
2. From the Scheduling Services menu, select **Job Metadata** and then select **Job Definitions**.

The Job Definitions page displays.

3. From the Applications dropdown list, select the name of the UI application whose job definitions you want to view.

The job definitions for that application display in a table below the application dropdown list.

Column Name	Description
Name	This column displays the name of the job definition.
Display Name	This column shows the name of the job definition as it displays.
Description	This column displays a description of the job definition.
Path	This column displays the full directory path to which the job definition is saved.
Execution Technology	This column displays the type of job request that you can create using the given job definition, such as a Java, C, PL/SQL, Perl or hosted script job.
Job Type Name	This column displays the name of the job type upon which the job definition is based.

Note: If there are applications with similar names, be sure to select the UI application. For example, given an application called **ProjectFinancialsApp** and **ProjectFinancialsESSApp**, select **ProjectFinancialsApp** as this is the UI application.

4. To view the details of a specific job definition, click the name of the relevant job definition.

5.11.1.2 Creating a Job Definition

You can create a new job definition, which you can then use to create a job request for a particular application. The job definition includes the directory path of the job to be run, the name of the application with which the job definition is associated and the job type used for the job definition.

Additional properties can be defined for a job definition as follows:

- Parameters. You can configure editable or read-only parameters to be submitted to the job request.
- User properties. You can configure properties to be filled in by end users at run time, such as boolean, number and string values.
- Lists of Values. You can search for and register sources for lists of values to be used for parameters.

This section contains the following topics:

- [Configuring Parameters for a Job Definition](#)
- [Configuring User Properties for a Job Definition](#)
- [Registering Sources for Lists of Values](#)
- [Searching for List of Value Sources](#)

To create a job definition:

1. From the navigation pane, expand the **Scheduling Services** folder and select the Oracle Enterprise Scheduler application.
2. From the Scheduling Services menu, select **Job Metadata** and then select **Job Definitions**.

The Job Definitions page displays.

3. From the Applications dropdown list, select the name of the UI application for which you want to create a job definition.

Note: If there are applications with similar names, be sure to select the UI application. For example, given an application called **ProjectFinancialsApp** and **ProjectFinancialsESSApp**, select **ProjectFinancialsApp** as this is the UI application.

4. Click **Create** to create a new job definition.
 5. For the new job definition, configure the following required properties:
 - **Name:** Enter a name for the job definition.
 - **Display Name:** Enter the name to be displayed for the job definition.
 - **Path:** Enter the path of the job to be run.
 - **Application Name:** From the dropdown list, select the name of the application with which you want to associate the job definition.
 - **Job Type Name:** From the dropdown list, select the job type you want to use for the job definition.
- Additional properties may display, depending on the job type selected.
6. Optionally, configure parameters, user properties, system preferences and access control. For more information, see [Section 5.11.1.2.1](#), [Section 5.11.1.2.2](#), [Section 5.11.1.2.3](#) and [Section 5.11.1.2.4](#).
 7. Click **Save and Close**.

5.11.1.2.1 Configuring Parameters for a Job Definition A parameter enables passing data of type string, boolean, number or date time to the job request. Parameters can be passed using a number of display controls, such as a text box, date picker, choice list or list of values.

To configure parameters for a job definition:

1. In the job definition creation user interface, click the **Parameters** tab.
2. Select **Actions** and then select New or click the **New** icon.
The Create Parameter dialog box displays.
3. In the Create Parameter dialog, enter the following information:
 - **Prompt:** Enter the text to be displayed at the prompt that displays during run time.
 - **Data Type:** From the dropdown list, select a data type for the parameter: String Numeric, DateTime, Number or Boolean. Required.
 - **Field Type:** Select **Editable** or **Read Only** as desired.

- **Value:** Enter a value for the parameter. This field displays for String type parameters. Required.
 - **Default Value:** Enter a default value for the parameter.
 - **Tooltip Text:** Enter documentation for the parameter.
 - **Required:** Select this checkbox if the parameter must be filled in.
 - **Do Not Display:** Select this checkbox if you want to hide the parameter in the user interface.
 - **Display Control:** From the dropdown list, select the control you want to display in the user interface. Display control options include Text Box, Choicelist, List of Values (with search), Date Picker.
4. Once you select a data type, you must configure its display.
- **Boolean:** From the Default Value dropdown list, select true or false.
 - **DateTime:** Select the **Date Only** or **Date and Time** radio buttons.
From the Date Format dropdown list, select the format in which to display the date.
From the Default Value dropdown list, select any of the following values: **Today**, **Tomorrow**, **A week from now** or **Specific date**.
If you select Specific date, use the calendar picker to select a date. This is required.
 - **Numeric:** From the Number Format dropdown list, select the format in which to display the number.
For the Data Alignment radio button, select **Left** or **Right**. This configures the alignment of the number value in the job request submission UI.
In the Default Value field, enter a default value for the number.
 - **String:** From the Display Widget dropdown list, select **Text Box**, **Choicelist** or **List of Values (with search)**. Required.
Text Box: When selecting this value from the Display Widget dropdown list, enter the default text in the **Default Value** text field.
Choicelist/List of Values (with search): From the List of Values Source dropdown list, select the source for the list of values. For the List of Values (with search), enter a default value in the field provided. In the Display Attributes shuttle list, select the attributes you want to display. For information about searching for and registering a list of values, see [Section 5.11.1.2.3](#) and [Section 5.11.1.2.4](#).
5. Click **Save and Close** or **Save and Create Another**.

5.11.1.2.2 Configuring User Properties for a Job Definition If your job definition requires additional properties to be filled in by end users at runtime, you can add these properties in the User Properties tab of the job definition creation user interface.

To configure user properties for a job definition:

1. In the job definition creation user interface, click the **User Properties** tab.
2. Select **Actions** and then select New or click the **New** icon.
The Create User Property dialog box displays.
3. In the Create User Property dialog, enter the following information:

- **Name:** Enter a name for the user property.
- **Data Type:** From the dropdown list, select a data type for the property. Data types include Boolean, Date, Number and String.
- **Default Value:** Enter a default value for the property.
- **Read Only:** Select Read Only if you want the property to be displayed only, not edited.

4. Click **Save and Close**.

5.11.1.2.3 Registering Sources for Lists of Values If you define a user property or parameter that uses a list of values, you must first register the source for the list of values.

To register a source for a list of values:

1. In the job definition creation user interface, click the **Manage List of Value Sources** tab.

2. Select **Actions** and then select **New** or click the **New** icon.

The Create List of Value Source dialog box displays.

3. In the Create List of Value Source dialog, enter the following information:

- **Application Name:** From the dropdown list, select the name of the application for the list of value. Required.
- **List of Values Source Definition Name:** Enter the name of the source definition file for the list of values. Required.
- **User List of Values Source Name:** Enter the name of the source for the list of values to be displayed to users. Required.
- **Description:** Enter a description for the list of values.

4. Click **Save and Close**.

5.11.1.2.4 Searching for List of Value Sources When defining a user property or parameter using a list of values, you must register a source for the list of values. The Search field enables searching for a source for a list of values. You can also save the search for future use.

For information about registering a source for a list of values, see [Section 5.11.1.2.3](#).

To search for list of value sources:

1. In the job definition creation user interface, click the **Manage List of Value Sources** tab.

2. In the Search field, select any of the following:

- **Match:** Select the All or Any radio button to search for any or all of the criteria.
- **Application Name:** From the dropdown list, select the name of the application you want to search. Then select an operator such as **Greater than or equal to**.
- **User List of Values Source Name:** Enter text for the user list of values source name. Then select an operator such as **Contains**.

3. Optionally, you can add additional fields to search by. Click **Add Fields** and select any of the following fields to search by:

- Application ID

- Created by
 - Creation Date
 - Data Source ID
 - Description
 - Last Update Date
 - Last Update Login
 - Last Updated By
 - List of Values Source Definition Name
 - Object Version Number
 - Product Short Name
 - User List of Values Source Name
4. Click **Search** to run the search, or **Save** to reserve the search for future use.

5.11.2 Managing Job Sets

The Job Sets page in Fusion Applications Control allows you to view, create, edit, delete and search for job sets.

A job set is a collection of job requests that can be grouped together to run as a single unit. A job set may be nested, such that it may contain a collection of job requests or one or more child job sets. Each job request or job set included within a job set is called a job set step.

Figure 5–4 shows the Results table in the Job Sets page. The Results table displays job sets and their attributes, including name, package, execution mode and description.

Figure 5–4 The Job Sets Page and Results Table

Scheduling Service Home > Job Sets

Job Sets

Filter Criteria

Application: EssDemoApp

Name:

Package:

Go

Results

Create... Edit... Delete...

Name	Package	Execution Mode	Description
ShortAndLongJob	/mypackage	Serial	

Total Rows : 1

This section contains the following topics:

- [Viewing Job Sets](#)

- [Creating or Editing a Job Set](#)
- [Deleting a Job Set](#)

5.11.2.1 Viewing Job Sets

You can view the job sets created for a given application.

To display job sets:

1. From the navigation pane, expand the **Scheduling Services** folder and select the Oracle Enterprise Scheduler application.
2. From the Scheduling Services menu, select **Job Metadata** and then select **Job Sets**.
The Job Sets page displays.
3. From the Applications dropdown list, select the name of the application whose job sets you want to view.
4. Enter the name and package name of the job set you want to find, and click **Go**.
The job sets for that application display in a table below the application dropdown list.

Column Name	Description
Name	This column displays the name of the job set.
Package	This column displays the name of the Java package associated with the job set.
Execution Mode	This column displays the execution mode of the job set, Serial or Parallel.
Description	This column displays a description of the job set.

5. To view the details of a specific job set, click the name of the relevant job set.

5.11.2.2 Creating or Editing a Job Set

A job set is defined as either a serial or a parallel job set. At run time, Oracle Enterprise Scheduler runs parallel job set steps together, in parallel. When a serial job set runs, Oracle Enterprise Scheduler runs the steps one after another in a specific sequence. Using a serial job set Oracle Enterprise Scheduler supports conditional branching between steps based on the execution status of a previous step.

For each step in a job set, you can configure an action to be take upon completion of the step, depending on the state of the step. You can also configure parameters and system properties for the job set, such as elevating access privileges to complete a particular job request or specifying the number of times a job can be retried in the event of failure.

To create or edit a job set:

1. From the navigation pane, expand the **Scheduling Services** folder and select the Oracle Enterprise Scheduler application.
2. From the Scheduling Services menu, select **Job Metadata** and then select **Job Sets**.
The Job Sets page displays.
3. Click **Create** to define a new job set, or **Edit** to modify an existing job set.

4. In the Job Set field, enter the name of the job set in the text field provided. Optionally, add a description in the **Description** text field, and the name of the relevant job set Java package in the **Package** text field.
5. In the Job Set Steps field, select **Serial** or **Parallel** to create a serial or parallel job set.
6. Add steps as required by clicking the **Add** icon. Define each step as required.
 - a. In the Step tab, in the **Step ID** field, enter a meaningful ID for the step.
In the Job field, enter search text and click the browsing button. In the window that displays, select the job definition that you want to use for this step.
 - b. In the Effective Application region, select **Insert into main diagram** or **Add to list of available steps**. If you choose to add the step to the list of available steps, use the dropdown lists that display to define an action for the possible job outcomes, namely **On Success**, **On Error** and **On Warning**.
 - c. In the Parameters tab, click the **Add** icon to define any required parameters and enter their initial value in the field provided. For more information about defining parameters, see [Section 5.11.1.2.1](#).
 - d. In the System Properties tab, click the **Add** icon to define any system parameters required for the step.
From the **Name** dropdown list, select the system property you want to specify. Possible system properties are shown in [Table 5-3](#).
In the **Initial Value** text field, enter the value you want to assign to the system property.

Table 5-3 System Properties

System Property	Description
SYS_allowMultPending	This property specifies whether multiple pending requests for the same job definition is allowed. This property has no meaning for a job set step. True or false.
SYS_application	This property specifies the logical name of the J2EE application used for request processing. This property is automatically set by Oracle Enterprise Scheduler during request submission.
SYS_bizErrorExitCode	This property specifies the process exit code for a process job request that denotes an execution business error. If this property is not specified, the system treats a process exit code of 4 as an execution business error. This property is optional for a process job type. It is not used for other job types.
SYS_className	This property specifies the Java executable for a Java job request. This should be the name of a Java class that implements the <code>oracle.as.scheduler.Executable</code> interface. This property is required for a Java job type. It is not used for other job types.
SYS_cmdLine	This property specifies the command line used to invoke an external program for a process job request. This property is required for a Process job type. It is not used for other job types.
SYS_effectiveApplication	This property specifies the logical name of the J2EE application that will be the effective application used to process the request. A job definition, job type, or a job set step can be associated with a different application by defining the <code>EFFECTIVE_APPLICATION</code> system property. This property can only be specified via metadata and cannot be specified as a submission parameter.

Table 5–3 (Cont.) System Properties

System Property	Description
<code>SYS_environmentVariables</code>	<p>This property specifies the environment variables to be set for the spawned process of a process job request. The property value should be a comma separated list of name value pairs (name=value) representing the environment variables to be set.</p> <p>This property is optional for a Process job type. It is not used for other job types.</p>
<code>SYS_executePast</code>	<p>This property specifies whether instances of a repeating request with an execution time in the past should be generated. Instances are never generated before the requested start time nor after the requested end time. To cause past instances to be generated, you must set this property to TRUE and specify the requested start time as the initial time from which instances should be generated. Note that a null requested start time defaults to the current time.</p> <p>Valid values for this property are:</p> <ul style="list-style-type: none"> ▪ TRUE: All instances specified by a schedule are generated regardless of the time of generation. ▪ FALSE: Instances with a scheduled execution time in the past (that is, before the time of generation) will not be generated. <p>If this property is not specified, the system defaults to TRUE.</p>
<code>SYS_extensionListener</code>	This property is for internal use only.
<code>SYS_external_id</code>	This property specifies an identifier for an external portion of an asynchronous Java job. For example, an asynchronous Java job usually invokes some remote process and then returns control to Oracle Enterprise Scheduler. This property can be used to identify the remote process. This property should be set by the job implementation of asynchronous Java jobs when the identifier is known. It is never set by Oracle Enterprise Scheduler.
<code>SYS_groupName</code>	This property specifies the name of the Oracle Enterprise Scheduler isolation group to which this request is bound. This property is automatically set by Oracle Enterprise Scheduler during request submission.
<code>SYS_inputList</code>	<p>This property specifies input to a job request. The input to a serial job set is forwarded as input to the first step only. The input to a parallel job set is forwarded as input to all the parallel steps.</p> <p>Oracle Enterprise Scheduler imposes no format on the value of this property.</p>
<code>SYS_inputWorkDirectory</code>	This property specifies the working directory used during job request processing for input files. Oracle Enterprise Scheduler sets the value of this property during job request processing.
<code>SYS_listener</code>	This property specifies the event listener class associated with the request. This should be the name of a Java class that implements the <code>oracle.as.scheduler.EventListener</code> interface.
<code>SYS_locale</code>	This property specifies a language code to apply to the job.
<code>SYS_logWorkDirectory</code>	This property specifies the logging working directory.
<code>SYS_outputList</code>	<p>This property specifies output from a request.</p> <p>The output of a serial job set is the <code>OUTPUT_LIST</code> of the last step. The output of a parallel job set is the concatenation of the <code>OUTPUT_LIST</code> of all the steps, in no guaranteed order, with <code>oracle.as.scheduler.SystemProperty.OUTPUT_LIST_DELIMITER</code> as a separator.</p>
<code>SYS_outputWorkDirectory</code>	This property specifies the working directory used during job request processing for output files. Oracle Enterprise Scheduler sets the value of this property during job request processing.

Table 5–3 (Cont.) System Properties

System Property	Description
<code>SYS_postProcess</code>	This property specifies the post-process callout handler class. This should be the name of a Java class that implements the <code>oracle.as.scheduler.PostProcessHandler</code> interface.
<code>SYS_preProcess</code>	This property specifies the pre-process callout handler class. This should be the name of a Java class that implements the <code>oracle.as.scheduler.PreProcessHandler</code> interface.
<code>SYS_priority</code>	This property specifies the request processing priority. The priority interval is [0..9] with 0 as the lowest priority and 9 as the highest. Default: If this property is not specified, the system default value used is 4.
<code>SYS_procedureName</code>	This property specifies the name of the PL/SQL stored procedure to be called for a SQL job request. The stored procedure should be specified using schema.name format. The property is required for a SQL job type. It is not used for other job types.
<code>SYS_product</code>	This property specifies the product within the application that submitted the request.
<code>SYS_redirectedOutputFile</code>	This property specifies the file where standard output and error streams are redirected for a Process job request. This represents the full path of the log file where the standard output and error streams are redirected for the spawned process when the request is executed. This property is optional for process job types. It is not used for other job types.
<code>SYS_reprocessDelay</code>	This property specifies the callout handler processing delay time. This represents the time, in minutes, to delay request processing when a delay is requested by a callback handler. Default: If this property is not specified, the system default used is 5. Integer type.
<code>SYS_request_timeout</code>	This property enables the job request to time out.
<code>SYS_requestCategory</code>	This property specifies an application-specific label for a request. The label, defined by an application or system administrator, allows administrators to group job requests according to their own specific needs.
<code>SYS_requestedProcessor</code>	This property specifies the request processor node on which the request should be processed. This allows processor affinity to be specified for a job request. If this property is not specified, the request can run on any available request processor node. In general, this property should not be specified. If this property is specified for a request, the request processor's work assignments <code>oracle.as.scheduler.WorkAssignment</code> (specialization) must allow the execution of such requests, otherwise the request will never be executed. If the specified node is not running, the request will remain in READY state and will not be executed until the node is restarted.
<code>SYS_requestExpiration</code>	This property specifies the expiration time for a request. This represents the time, in minutes, that a request will expire after its scheduled execution time. A expiration value of zero (0) means that the request never expires. If this property is not specified, the system default value used is 0. Request expiration only applies to requests that are waiting to run. If a request waits longer than the specified expiration period, it does not run. After a request starts running the request expiration no longer applies.
<code>SYS_retries</code>	This property specifies the retry limit for a failed request. If request execution fails, the request will be retried up to the number of times specified by this property until the request succeeds. If retry limit is zero (0), a failed request will not be retried. If this property is not specified, the system default used is 0.

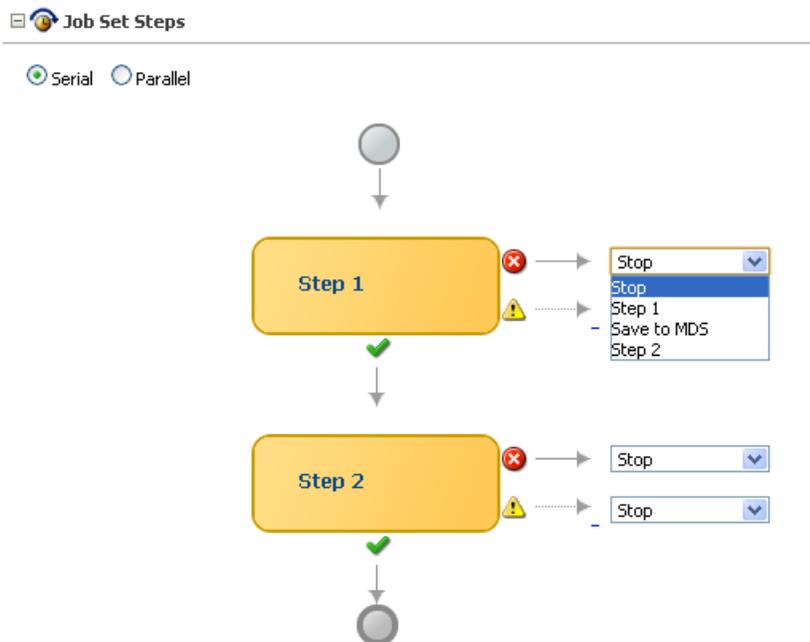
Table 5–3 (Cont.) System Properties

System Property	Description
<code>SYS_runasApplicationID</code>	<p>This property enables elevating access privileges for completing a scheduled job.</p> <p>Normally, a request runs as the submitting user. However, if this property is set in the metadata of the job associated with the request, then the request executes under the user identified by this property. This property can only be specified via metadata and cannot be specified as a submission parameter.</p>
<code>SYS_selectState</code>	<p>This property specifies whether the result state of a job set step affects the eventual state of its parent job set. In order for the state of a job set step to be considered when determining the state of the job set, the <code>SELECT_STATE</code> must be set to <code>true</code>. If <code>SELECT_STATE</code> is not specified on a job set step, the state of the step will be included in the determination of the state of the job set. Boolean type.</p>
<code>SYS_sqlJobClass</code>	<p>This property specifies an Oracle Enterprise Scheduler job class to be assigned to the Oracle Enterprise Scheduler job used to execute a SQL job request. This property need not be specified unless the job used for a job request is associated with a particular Oracle Database resource consumer group or has affinity to a database service.</p> <p>If this property is not specified, a default Oracle Enterprise Scheduler job class is used for the job that executes the SQL request. That job class is associated with the default resource consumer group. It belongs to the default service, such that it has no service affinity and, in an Oracle RAC environment, any one of the database instances within the cluster might run the job. No additional privilege or grant is required for an Oracle Enterprise Scheduler SQL job request to use that default job class.</p> <p>This property is optional for a SQL job type. It is not used for other job types.</p>
<code>SYS_submittingApplication</code>	<p>This property specifies the logical name of the J2EE application for the submitted (absolute parent) request. This property is automatically set by Oracle Enterprise Scheduler during request submission.</p>
<code>SYS_substitutionHandlers</code>	<p>This property has been deprecated.</p>
<code>SYS_successExitCode</code>	<p>This property specifies the process exit code for a process job request that denotes an execution success. If this property is not specified the system treats a process exit code of 0 as execution success.</p> <p>This property is optional for a process job type. It is not used for other job types.</p>
<code>SYS_userFileDir</code>	<p>This property specifies a base directory in the file system where files, such as input and output files, may be stored for use by the request executable.</p> <p>Oracle Enterprise Scheduler supports a configuration parameter that specifies a file directory where requests may store files. At request submission, a <code>USER_FILE_DIR</code> property is automatically added for the request if the configuration parameter is currently set and <code>USER_FILE_DIR</code> property was not specified for the request. If the property is added, it will be initialized to the value of the configuration parameter. The property will not be added if the configuration parameter is not set at the time of request submission.</p>

Table 5–3 (Cont.) System Properties

System Property	Description
SYS_userName	This property specifies the name of the user used to execute the request. Normally, this is the submitting user unless the RUNAS_APPLICATIONID property was set in the job metadata. This property is automatically set by Oracle Enterprise Scheduler during request submission.
SYS_warningExitCode	This property specifies the process exit code for a process job request that denotes an execution warning. If this property is not specified, the system treats a process exit code of 3 as execution warning. This property is optional for a process job type. It is not used for other job types.
SYS_workDirectoryRoot	This property specifies the working directory for the spawned process of a Process job request. This property is optional for a process job type. It is not used for other job types.

- e. If you configured the step as a serial step, it displays in the job set flow diagram. Configure the action to be taken upon reaching error and warning states, respectively. From the dropdown list for the error and warning icons, select **Stop** or the name of the job definition to run upon reaching an error or warning state.



7. Continue to define steps as required for the job set.
8. Define parameters and system properties as required towards the bottom of the job set window.
9. Configure access control for the job set.
 - a. In the Access Control section of the Create Job Set page, click the **Add** icon to add access control to the job set.
The Add Access Control dialog box displays.

- b. From the Role dropdown list, select the name of the role you want to apply to the job set. The role defines the users with sufficient permissions to access the job set.
 - c. Select the actions you want to allow role members to take: **Read**, **Execute**, **Update** and **Delete**.
10. Click **OK** to save the job set.

5.11.2.3 Deleting a Job Set

You can delete a job set.

To delete a job set:

1. Search for the relevant job set, as described in [Section 5.11.2.1](#).
2. In the Results table, select the job set you want to delete and click the **Delete** icon.

5.11.3 Managing Incompatibilities

An Oracle Enterprise Scheduler incompatibility specifies which job definitions are incompatible with other job definitions. At runtime, when job definitions are specified in an incompatibility, any job requests associated with the job definitions that are incompatible cannot run simultaneously.

A given incompatibility defines either a global incompatibility or a domain-specific, property-based incompatibility.

This section contains the following topics:

- [Viewing Incompatibilities](#)
- [Creating or Editing an Incompatibility](#)
- [Deleting an Incompatibility](#)

5.11.3.1 Viewing Incompatibilities

The Incompatibilities page displays information about incompatibilities including name, the Java package associated with a given incompatibility and a description for the incompatibility.

To view job incompatibilities:

1. From the navigation pane, expand the **Scheduling Services** folder and select the Oracle Enterprise Scheduler application.
2. From the Scheduling Services menu, select **Job Metadata** and then select **Incompatibilities**.

The Incompatibilities page displays.

3. From the Applications dropdown list, select the name of the application whose incompatibilities you want to view.
4. Enter the name and package name of the incompatibilities you want to find, and click **Go**.

The incompatibilities for that application display in a table below the application dropdown list.

Column Name	Description
Name	This column displays the name of the incompatibility.

Column Name	Description
Package	This column displays the name of the Java package associated with the incompatibility.
Incompatibility Type	This column displays the type of incompatibility, global or domain specific.
Description	This column displays a description of the incompatibility.

- To view the details of a specific incompatibility, click the name of the relevant incompatibility.

5.11.3.2 Creating or Editing an Incompatibility

An incompatibility consists of two or more job definitions configured as incompatible, and the resource over which they need to be incompatible. A resource is not specified for a global incompatibility. This resource is represented by the property name that forms the Incompatibility. The property name might be different across job definitions. For example, if two job definitions, JobA and JobB, are made incompatible, then the property identified for incompatibility might have different names in JobA and JobB.

- Domain-specific:** Two or more job definitions are marked as incompatible within the scope of a resource, where the resource is identified by a system property name or a user-defined parameter name. An incompatibility specifies two or more incompatible entities where each entity consists of a job definition or a job set definition and a property name that represents the resource over which the entities are incompatible. Each entity can use the same property name or a different property name. A property-based incompatibility is used when the incompatible entities are incompatible over a given property. For property-based incompatibility, a job definition or a job set and a property name must be specified for each entity. Oracle Enterprise Scheduler ensures that requests for the incompatible jobs or job sets, with a property that has the same property value, do not run at the same time.
- Global:** Two or more job definitions cannot run together at any time. A global incompatibility is used when the incompatible entities are not to be run at the same time regardless of any property. For global incompatibility, only the job definition or job set is specified for each entity. Oracle Enterprise Scheduler ensures that requests for the incompatible Jobs or job sets do not run at the same time.

Defining an incompatibility involves the following:

- Package and scope:** Select a relevant Java package to use for the incompatibility and set the scope for the incompatibility (global or domain only).
- Jobs:** Select the jobs that are incompatible.
- Parameters and properties:** Define parameters and properties as required.
- Access control:** Define access control as required.

To create or edit an incompatibility:

- From the navigation pane, expand the **Scheduling Services** folder and select the Oracle Enterprise Scheduler application.
- From the Scheduling Services menu, select **Job Metadata** and then select **Incompatibilities**.

The Incompatibilities page displays.

3. Click **Create** to define a new incompatibility, or **Edit** to modify an existing incompatibility.
The Create Incompatibilities page displays.
4. Enter the following information:
 - **Name:** Enter a name for the incompatibility.
 - **Package:** Enter the name of the relevant incompatibility Java package.
 - **Description:** Optionally, add descriptive text for the incompatibility.
 - **Type:** Configure the scope of the incompatibility by selecting **Global** or **Domain**.
5. In the Entities field, click the Add icon to add jobs to the incompatibility.
The Add Entity dialog box displays.
6. Select one or more jobs and click **OK**.
7. Define parameters and system properties as required towards the bottom of the incompatibility window. For more information about defining parameters and system properties, see [Section 5.11.2.2](#).
8. Configure access control for the incompatibility, as described in step 9 of [Section 5.11.2.2](#).
9. Click **OK** to save the incompatibility.

5.11.3.3 Deleting an Incompatibility

Deleting an incompatibility results in the incompatible job requests or job sets becoming compatible again.

To delete an incompatibility:

1. Search for the relevant incompatibility, as described in [Section 5.11.3.1](#).
2. In the Results table, select the incompatibility you want to delete and click the **Delete** icon.

5.12 Managing Work Assignments and Workshifts

Work allocation is the process of defining constraints on where and when jobs run as well as the amount of resources that can be used to process the jobs. This process includes creating a work assignment, and binding the work assignment to a request processor.

A work assignment consists of a specialization rule and one or more workshifts. The specialization rule defines restrictions on which jobs can be processed. A workshift defines the temporal windows when the jobs can be processed and what resources can be used during those periods. The resources defined in a workshift include threads, which are a local resource of the request processor, and asynchronous workers, a global resource. The number of asynchronous workers can be specified to throttle the use of a shared global resource, such as database jobs.

Binding associates a work assignment with a request processor on a server, determining where the jobs can run. An exclusive binding mode is supported to not only determine when the jobs can run but to prevent them from running anywhere else.

By default, no work assignments are bound. When there is no bound or active work assignment, a virtual default work assignment will be started to process all jobs using all available resources.

For more information about managing work assignments, workshifts and schedules, see the section "Managing Work Assignments and Workshifts" in the chapter "Managing the Work of Oracle Enterprise Scheduler Jobs" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Enterprise Scheduling Service*.

5.13 Managing Oracle ADF Connections

The ADF Connections Configuration page allows you to create connections to services such as ADF Business Components services, Oracle Business Activity Monitoring, Oracle Enterprise Scheduler, URLs and web services.

Figure 5–5 shows the ADF Connections Configuration page.

Figure 5–5 The ADF Connections Configuration Page



Use the ADF Connections Configuration page to configure the following properties for Oracle Enterprise Scheduler:

- **Request Files Store (RequestFileDirectory):** From the dropdown list, select **File System** or **Content Repository**. to store job request logs to the file system or a repository. This property is used for job requests submitted by Oracle ADF applications.
- **Request Files Location:** Configure this property to indicate the path of the folder to which job request logs are to be written. This property displays only when selecting **File System** from the **Request Files Store** dropdown list.
- **UMS Server URL (NotificationServiceURL):** Configure this property when using the notification service with Oracle Enterprise Scheduler. Set the URL of this property to that of the UMS Server NotificationServiceURL.
- **SAML Token Policy (SAMLTokenPolicyURI):** Set this property to the URI of the SAML policy used to secure job requests submitted by Oracle ADF applications.

- **Client Callback Policy (EssCallbackClientSecurityPolicyURI)**: Configure this property to enable the security policy used in the `WS-Security` headers for web service invocations from Oracle Enterprise Scheduler for web service callbacks.

You can manage ADF connections in Fusion Applications Control only.

To create or edit an ADF connection:

1. From the navigation pane, expand the **Scheduling Services** folder and select the Oracle Enterprise Scheduler application.
2. From the Scheduling Services menu, select **Configuration** and then select **ADF Connections**.

The ADF Connections Configuration page displays.

3. In the Create Connection field, click the Connection Type dropdown list and select the type of connection you want to create.
4. In the Connection Name field, enter a name for the connection.
5. Click **Create Connection** to add the connection.

The new connection displays under the relevant connection category.

6. To configure the connection, scroll down to the relevant connection category and select the new connection. Click **Edit**.

The connection window displays.

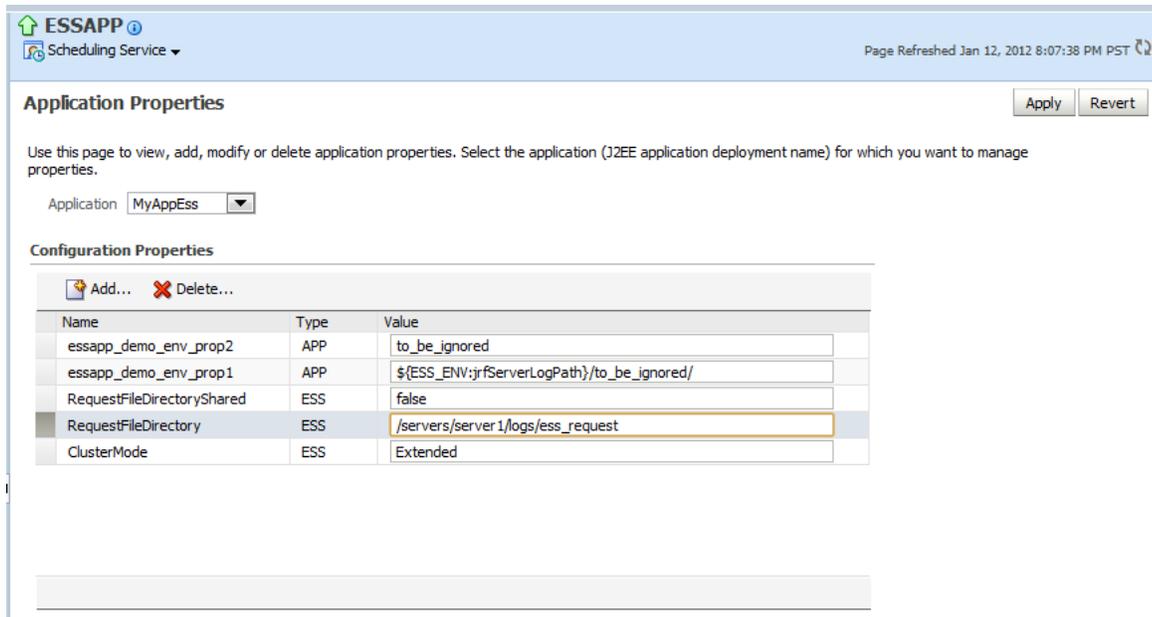
7. Configure the connection accordingly. For more information about the attributes you can configure for each type of connection, see the online help by clicking the help icon on the ADF Connections Configuration page.
8. Click **Apply** to activate the connection.

5.14 Managing Application Properties

The Managing Application Properties page allows you to define Oracle Enterprise Scheduler or application properties for a given deployed application.

Figure 5–6 shows the Applications Properties page.

Figure 5–6 The Application Properties Page



To create or edit an application property:

1. From the navigation pane, expand the **Scheduling Services** folder and select the Oracle Enterprise Scheduler application.
2. From the Scheduling Services menu, select **Configuration** and then select **Application Properties**.

The Application Properties page displays.

3. From the Application dropdown list, select the name of the application for which you want to configure a property.
4. In the Configuration Properties region, click **Add** to create a new property.

The Add Property window displays.

5. In the Add Property window, configure the following.
 - **Type:** From the dropdown list, select **ESS** or **APP**.
 - **Name:** When configuring an application property, enter a name for the property in the field provided. When configuring an Oracle Enterprise Scheduler property, select a property from the dropdown list. Properties include **PerlCommand** and **EssCallbackClientSecurityPolicyURI** (for more information about this property, see [Section 5.13](#)).
 - **Value:** Enter a value for the property.

Click **OK** to save the property.

6. Click **Apply** to activate the application property.

5.15 Managing Application and Role Policies

Oracle Enterprise Scheduler application and role policies are the same as other application and role policies. For more information about managing application and role policies, see [Chapter 4](#).

5.16 Managing Oracle Enterprise Scheduler Web Services

Oracle Enterprise Scheduler web services are the same as other Fusion application web services. For information about managing the web services related to Oracle Fusion Applications, see the "Understanding Oracle WSM Policy Framework" chapter in the *Oracle Fusion Middleware Security and Administrator's Guide for Web Services*.

5.17 Managing Logging for Oracle Enterprise Scheduler

You can search for and view log data for individual job requests, as well as set log levels for Oracle Enterprise Scheduler. You can also save job request logs to a file, and trace job requests for additional troubleshooting information.

For more information about logging, see [Chapter 13](#) and the section "Managing Logging for Oracle Enterprise Scheduling Service" in the chapter "Monitoring Oracle Enterprise Scheduler" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Enterprise Scheduling Service*. For more information about managing log files and diagnostic data, see the chapter "Managing Log Files and Diagnostic Data" in the *Oracle Fusion Middleware Administrator's Guide*.

This section contains the following topics:

- [Finding Log Information for a Job Request](#)
- [Setting Oracle Enterprise Scheduler Log Levels](#)
- [Saving Job Request Logs](#)

Note: The Oracle WebLogic Server logger (`logging.xml`) only shows logs written by Oracle Enterprise Scheduler job running in Oracle WebLogic Server. Once Oracle Enterprise Scheduler transfers control of running PL/SQL and C jobs to the PL/SQL or C process, respectively, PL/SQL and C job logging data is not written to the Oracle Enterprise Scheduler logs as they run in a separate process.

5.17.1 Finding Log Information for a Job Request

Job request logs can often be written to different files. You can find log data regarding a given job request using its Execution Context Identifier (ECID), a unique identifier that enables finding log data for the job request.

For more information about finding the ECID for a job request, see [Section 14.5.3](#)

5.17.2 Setting Oracle Enterprise Scheduler Log Levels

It is possible to configure Oracle Enterprise Scheduler server logging for a Oracle WebLogic Server by modifying the `logging.xml` file of that Oracle WebLogic Server. By default, there is no explicit logger entry for Oracle Enterprise Scheduler. Oracle Enterprise Scheduler inherits the logging level and log handlers configured for the parent logger, typically the `oracle` logger or the root logger.

By default, the log messages for the Oracle Enterprise Scheduler logger can be found in the Oracle WebLogic Server diagnostic log file for that Oracle WebLogic Server. The `logging.xml` file is located under `DOMAIN_HOME/config/fmwconfig/servers/WebLogic_Server_Name`, where `DOMAIN_HOME` is the domain home directory for the Oracle WebLogic Server domain and `WebLogic_Server_Name` is the name of the Oracle WebLogic Server that uses the `logging.xml` file.

Table 5–4 shows the Oracle Enterprise Scheduler logger names, log levels and a description for each level.

Table 5–4 Loggers and log levels for Oracle Enterprise Scheduler

Logger Name	Log Level	Description
oracle.as.ess	SEVERE	Problems encountered by Oracle Enterprise Scheduler runtime in the context of request processing that result in the request errors. Errors include exceptions thrown by the job code, unchecked exceptions when running the job code and exceptions when running Oracle Enterprise Scheduler code. Problems encountered by Oracle Enterprise Scheduler runtime outside the context of request processing, such as dispatching, system event handling, and so on.
	WARNING	Less severe problems encountered by Oracle Enterprise Scheduler runtime during or outside of request processing, which might not cause requests to enter error state.
	INFO	Messages for request state transitions. Messages related to work assignment activities. Messages about batch delete failures. Start and stop of Oracle Enterprise Scheduler resource adapter.
oracle.apps.fnd.applcp.srs	CONFIG	Application endpoint activation and de-activation for Oracle Enterprise Scheduler resource adapter.
	WARNING	Any problems that occur during UI rendering to submission.
	INFO	Messages related to job fetch and submission API calls.
oracle.apps.fnd.applcp.monitor	FINEST	Details tracing messages for the scheduled job request submission UI.
	WARNING	Any problems that occur during UI rendering to operations in the UI.
	INFO	Messages related to job request fetch and various API calls.
oracle.apps.fnd.applcp	FINEST	Details tracing messages for the job request monitoring UI.
	SEVERE	Records any errors occurred when creating a session for Oracle Fusion Middleware Extensions for Applications, or when creating a file during pre- and post-processing.
	WARNING	Records messages related to terminating sessions and closing files during pre- and post-processing
	FINE	Messages related to pre- and post-processing execution activity.

- Application Logging: Oracle Enterprise Scheduler job implementation might call standard Oracle Fusion Applications logging code. For proper operation, the AFLOG_ENABLED profile option must be set to Y. For PL/SQL jobs, set AFLOG_PLSQL_FILENAME to Y. For C jobs, set AFLOG_FILENAME to Y. For more information about

configuring log settings, see Section [Section 13.7](#) and [Section 14.2](#). For more information about setting the `AFLOG_ENABLED` property, see the section "Enabling JBO Logging in the Provisioned Environment" in the "Debugging Oracle ADF and Oracle SOA Suite" chapter in the *Oracle Fusion Applications Developer's Guide*.

- **Request Logging:** Oracle Enterprise Scheduler job implementation might write business-specific job request execution log information to the job request log file. This log file is specific to each request, and is automatically enabled by default. For more information about viewing the log file for a job request, see the section "Viewing Job Request Logs" in the chapter "Monitoring Oracle Enterprise Scheduler" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Enterprise Scheduling Service*.

For information about setting the log levels for Oracle WebLogic Server, see the section "Tracing Oracle Enterprise Scheduling Service Jobs" in the chapter "Monitoring Oracle Enterprise Scheduler" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Enterprise Scheduling Service*. For more information about Oracle Diagnostic Logging levels, see [Section 13.7.3](#).

Set the log levels for the Oracle WebLogic Server running Oracle Enterprise Scheduler as described in the Oracle WebLogic Server documentation.

5.17.3 Saving Job Request Logs

By default, job request logs are written to the `fnd_log` API if Oracle Fusion Applications logging is set to `FINER`.

It is possible to save job request log data to the server log file. Job request logs are normally stored to Oracle WebCenter Content. However, when setting the log level to `FINER`, all job request logs are copied to the server log file.

To save job request logs to the server diagnostic file in Fusion Applications Control:

1. In Fusion Applications Control, in the navigation pane, select **WebLogic Domain** and then select the domain and server names.
2. Click the **WebLogic Server** menu and select **Logs** and then select **Log Configuration**.

The Log Configuration pane displays.

3. In the Log Configuration pane, click the **Log Levels** tab.
4. In the Search text field, search for `oracle.as.scheduler`. and find the `oracle.as.scheduler.security.internal.SecurityExecutorFactory` logger.
5. Set the log level to `FINER` and click **Apply**.
6. Run a diagnostic test, as described in [Chapter 13](#). The log messages are now saved to the `<server>-diagnostic.log` file.

The logs are saved to the server diagnostic file, `<Oracle Enterprise Scheduler server name>-diagnostic.log`.

5.18 Managing Alerts for Oracle Enterprise Scheduler in Cloud Control

Cloud Control enables you to configure, view and manage alerts for Oracle Enterprise Scheduler entities including the scheduling service itself, work assignments and job definitions. Alerts are based on performance metrics for Oracle Enterprise Scheduler. Each metric includes named objects with upper and lower thresholds indicating warning or critical states for that particular metric. For example, you can configure an

alert for all job requests that have entered a state of error, such that 10 errors constitute a warning state, and 20 errors a critical state. Once you have configured alerts, you can view any existing alerts as well as any configuration changes made within a specified period of time.

Table 5–5 lists the alerts available for Oracle Enterprise Scheduler.

Table 5–5 Oracle Enterprise Scheduler Alerts

Alert	Scope	Collected Data	Alert Automatically Cleared?
Alert for job requests in ready state and the average wait time for job requests in ready state.	Work assignment	The number of job requests in ready state and average wait time of the ready job requests by work assignment.	No. The alert is only cleared when the number of ready job requests is smaller than the threshold. The same is true for alerts pertaining to the average wait time of a ready job request.
Alert for job requests that have entered a state of error.	Work assignment	The number of job requests that have entered an error state following the last collection.	Yes. The alert is automatically cleared if the delta value of the subsequent collection is below the threshold.
Alert for long running job requests.	Job definition	The longest running time of a job request listed by job definition. The data is queried for only those job definitions for which a threshold is specified.	No. The alert is only cleared when a job request runs for the job definition beyond the threshold value.
Alert for job requests that have gone to a state of <code>ERROR_MANUAL_RECOVERY</code> .	Scheduling service	The number of job requests in a state of <code>ERROR_MANUAL_RECOVERY</code> .	No. An administrator must recover the job request in order to clear the alert.
Alert for timed out job requests.	Scheduling service	The number of timed out job requests.	No. An administrator must cancel or recover the job request in order to clear the alert.
Alert for job requests that have gone into <code>ERROR</code> state.	Scheduling service	The number of job requests that have gone into <code>ERROR</code> state during the last hour.	Yes. The alert automatically clears if the delta value in the subsequent collection is below the threshold.
Alert for blocked job requests.	Scheduling service	The number of job requests in <code>BLOCKED</code> state.	No. An administrator must cancel the job request, or the job request must be selected for processing in order to clear the alert.

This section contains the following topics:

- [Configuring Metric Alerts for Oracle Enterprise Scheduler in Cloud Control](#)
- [Viewing Alerts for Oracle Enterprise Scheduler in Cloud Control](#)

5.18.1 Configuring Metric Alerts for Oracle Enterprise Scheduler in Cloud Control

You can configure the thresholds for each object monitored by the metric. For a given metric, you can configure alert rules as follows:

- Select an object that the metric monitors, such as job requests in an errored state.
- Configure the values at which the object reaches warning and critical states, including an operator value such as greater than (>) or less than (<).
- Enter a corrective action to be taken in the event that the alert criteria are fulfilled.
- Arrange the alert rules in the order you desire.

To configure metric alerts for Oracle Enterprise Scheduler in Cloud Control:

1. Navigate to the scheduling service.

- a. From the **Targets** menu, choose **Fusion Applications**.
 - b. In the table on the Fusion Applications page, click the relevant Oracle Fusion Applications instance with Oracle Enterprise Scheduler deployed.
 - c. In the **Target Navigation** pane, click the name of the domain where Oracle Enterprise Scheduler is deployed, and then click **Scheduling Services**.
2. Click the Oracle Enterprise Scheduler cluster or an individual Oracle Enterprise Scheduler server.
 3. From the **Targets** menu, select **Monitoring** and then click **Metric and Collection Settings**.
A list of metrics displays.
 4. Select the metric for which you want to configure alert thresholds and click the **Edit** icon. For example, you can configure the thresholds for the **Average Wait Time for Requests in Ready State (in seconds)** metric.
The **Edit Advanced Settings** page displays, including the Monitored Objects list.
 5. In the **Edit Advanced Settings** page, click **Add** to add a row to the list of monitored objects.
 6. Select the object you want to monitor.
 - a. In the **Name** field of the new row, click the **Search** button.
The Search and Select Objects window displays.
 - b. Search for the object you want to use and click **Select**.
 7. In the spaces provided, enter the values to be used for the Critical and Warning thresholds.
 8. Click **Continue**, then click **OK** to save your settings.
 9. Configure any additional thresholds as required.

For more information about configuring alerts, see the following topics in the Cloud Control online help:

- "Incident Rules—Common Tasks"
- "Incident Rules—Advanced Tasks"

5.18.2 Viewing Alerts for Oracle Enterprise Scheduler in Cloud Control

The Oracle Enterprise Scheduler Service home page displays a summary of alerts. You can navigate to the Incident Manager page to view additional information about incidents by clicking the number of incidents displayed in the Monitoring and Diagnostics area.

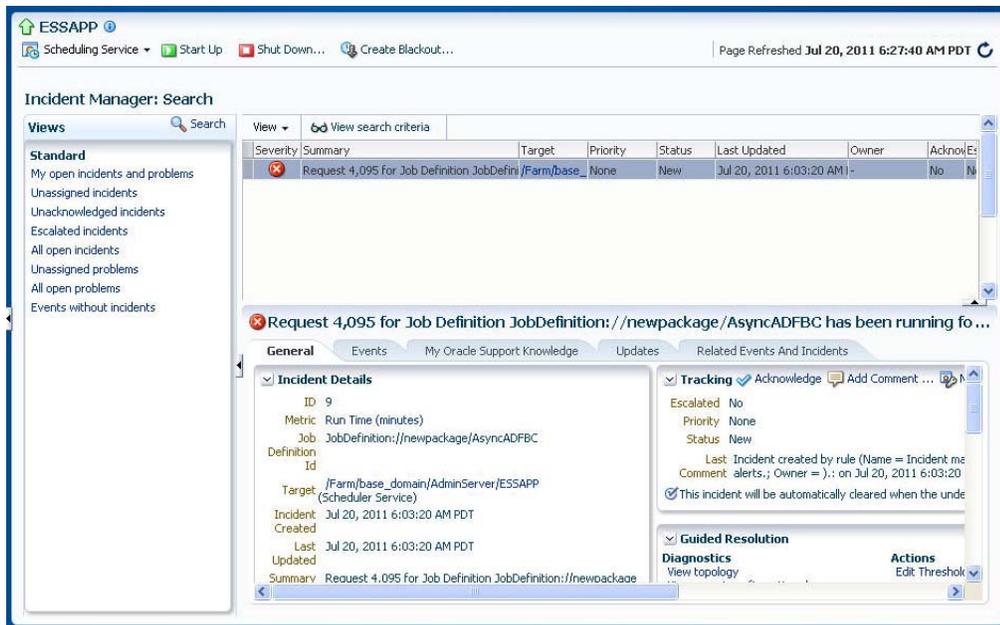
To view alerts for Oracle Enterprise Scheduler with Cloud Control:

1. Navigate to the scheduling service.
 - a. From the **Targets** menu, choose **Fusion Applications**.
 - b. In the table on the Fusion Applications page, click the relevant Oracle Fusion Applications instance with Oracle Enterprise Scheduler deployed.
 - c. In the **Target Navigation** pane, click the name of the domain where Oracle Enterprise Scheduler is deployed, and then click **Scheduling Services**.

- Click the Oracle Enterprise Scheduler cluster or an individual Oracle Enterprise Scheduler server.

The Scheduling Service summary page displays. The Monitoring and Diagnostics pane displays the following information.

- Incidents: The number of incidents that have occurred in the scheduling service instance displays here. The number of incidents is determined by those incidents collected as a result of configuring job request alerts. For more information about configuring job request alerts, see [Section 5.18.1](#). If the number of incidents is greater than zero, click the number for more information about the incidents that have occurred.
 - Descendant Target Incidents: The number of incidents that have occurred anywhere in the scheduling service cluster displays here. If the number of incidents is greater than zero, click the number for more information about the incidents that have occurred.
 - Configuration Changes: The number of changes that have been made to alert configurations displays here. Click the number of configuration changes to search for the changes made within a particular period of time. The default time period is seven days. For more information, see [Section 5.5](#).
- Click the number of incidents to display detailed information about the alerts in the Incident Manager pane.



The Incident Manager displays incidents in a table, with the following information listed for each incident: severity, summary, target, priority, status, last updated, owner, acknowledged, escalated, type and category.

You can click an incident to display its details. For more information about using Incident Manager, see the following topics in the Cloud Control online help:

- "Incident Manager—Common Tasks"
- "Incident Manager—Advanced Tasks"

Managing Report Delivery Servers

This chapter describes report delivery server administration in Oracle Fusion Applications.

This chapter contains the following topics:

- [Introduction to Managing Report Delivery Servers](#)
- [Navigating to the Oracle BI Publisher Administration Page](#)
- [Configuring Report Delivery Servers](#)
- [Understanding the Report and Delivery Processors](#)
- [Managing Report Processing and Delivery Server Load](#)
- [Scaling Out Report Processors](#)
- [Diagnosing Report Processing Issues](#)
- [Configuring System Properties for Reporting](#)

6.1 Introduction to Managing Report Delivery Servers

Oracle Business Intelligence Publisher is the report generation and delivery engine for Oracle Fusion Applications. Oracle BI Publisher receives report requests from Oracle Fusion Applications in the following ways:

- Through Oracle Enterprise Scheduler
- Through the Reports and Analytics pane
- From an application page

Requests submitted through Oracle Enterprise Scheduler are processed by the Oracle BI Publisher scheduler. Requests submitted through the Reports and Analytics pane can be either real-time online requests or scheduled requests. Requests submitted through an application may invoke Oracle Enterprise Scheduler or may return report request results directly back to the application page.

After installing Oracle Fusion Applications, Oracle BI Publisher is configured to accept requests from Oracle Fusion Applications. However, before you can deliver report documents to their destinations you must define the delivery servers in Oracle BI Publisher. Use the Oracle BI Publisher Administration page to define your delivery servers.

Once set up, you can then further configure the number of report processor and delivery threads to best handle your processing and delivery requirements. In addition, you can configure report properties for the system or at the report level to

tune performance of your reports. To diagnose report processing issues, BI Publisher provides a set of scheduler diagnostics.

6.2 Navigating to the Oracle BI Publisher Administration Page

Use the Oracle BI Publisher Administration page to:

- Configure delivery servers
- Manage report and delivery processors
- View scheduler diagnostics
- Set system properties and report runtime configuration properties

Note: You must be assigned the BI_Administrator role to access the BI Publisher Administration page.

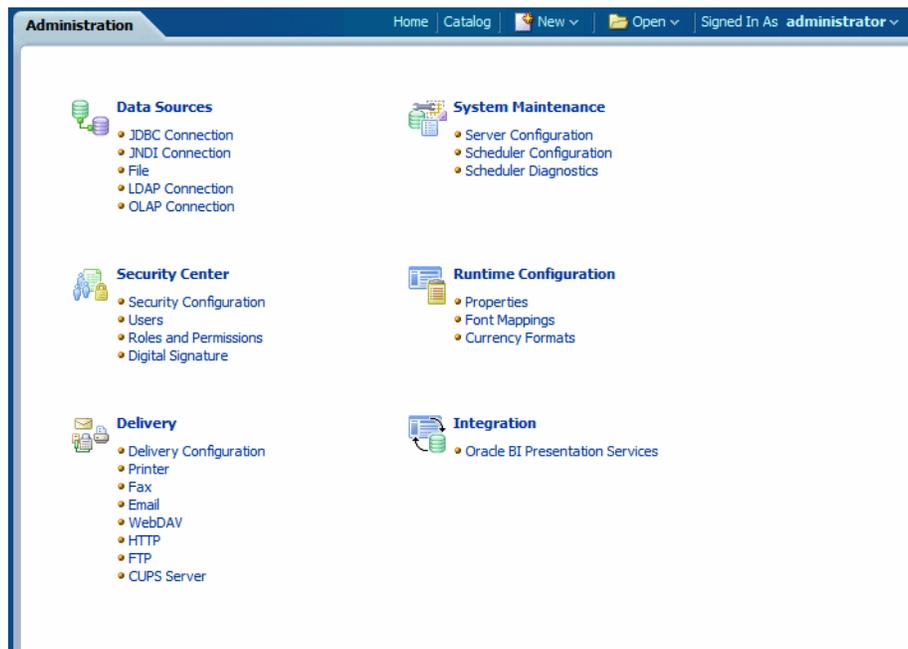
To navigate to the Oracle BI Publisher Administration page:

- From the Oracle Fusion Applications **Navigator**, under **Tools**, click **Reports and Analytics**. In the **Reports and Analytics** pane, click **Catalog** to display the Oracle Business Intelligence presentation catalog page. From here, click **Administration** and then click **Manage BI Publisher**.

Alternatively, log in to Oracle Business Intelligence directly (example: <http://example.com:port/analytics>). Click **Administration** and then click **Manage BI Publisher**.

Figure 6–1 shows the BI Publisher Administration page:

Figure 6–1 BI Publisher Administration Page



6.3 Configuring Report Delivery Servers

To configure delivery servers:

1. From the BI Publisher Administration page, click **Delivery Configuration**.
2. Enter values in the **Delivery Configuration Options** tab to set general properties for email deliveries and notifications. [Figure 6–2](#) shows the Delivery Configuration Options tab:

Figure 6–2 Delivery Configuration Options Tab

For more information about this tab see "Setting Delivery Configuration Options" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Business Intelligence Publisher*.

3. To configure a delivery server, click the appropriate tab.

The following table lists the report delivery channels supported by Oracle BI Publisher. See the corresponding section in the *Oracle Fusion Middleware Administrator's Guide for Oracle Business Intelligence Publisher (Oracle Fusion Applications Edition)* for configuration information.

Delivery Type	Section
Printer and Fax	Adding a Printer or Fax Server
E-mail	Adding an E-mail Server
WebDAV	Adding a WebDAV Server
HTTP	Adding an HTTP Server
FTP	Adding an FTP Server

Note that printing is supported through Internet Printing Protocol (IPP). If Oracle BI Publisher is operating in a UNIX environment, you must set up the Common Unix Printing Service (CUPS) and then define your CUPS server to Oracle BI Publisher. For a Windows environment, you must set up Windows Print Server for IPP. For information on setting up CUPS and Windows IPP, see "Setting Up Print Servers" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Business Intelligence Publisher (Oracle Fusion Applications Edition)*.

6.4 Understanding the Report and Delivery Processors

When Oracle Enterprise Scheduler initiates a job, it is picked up by the BI Publisher scheduler queue. The processors perform as follows:

- **Job Processor** - listens for requests from the scheduler queue. When the job information is received, the job processor executes the data model (and splits the data for bursting jobs), stores the data in the shared temporary folder, and sends the report metadata to the report queue.
- **Report Processor** - listens for requests from the report queue. When the report information is received, the report processor generates the report documents, stores it in the shared temporary folder and puts the delivery information in the delivery queue.
- Delivery processors - each delivery processor listens to the delivery queue and handles the delivery of reports for its channel. The delivery processors are:
 - **Email Processor**
 - **File Processor**
 - **FTP Processor**
 - **Print Processor**
 - **WebDAV Processor**
 - **Fax Processor**

You can configure the number of threads dedicated to each processor.

6.5 Managing Report Processing and Delivery Server Load

Manage the processors in the BI Publisher Scheduler Configuration page. By default, each processor is enabled and the thread count for each is set to five. For each managed server that is running in the BI cluster, a table for that instance's processors will display. Use the table to enable or disable processors for the instance and configure the thread counts.

To configure processor threads:

1. From the BI Publisher Administration page, click **Scheduler Configuration**.
2. On the **Scheduler Configuration** page scroll to the **Cluster Instances** region.

[Figure 6–3](#) shows the Oracle BI Publisher **Scheduler Configuration** page, highlighting the **JMS Configuration** and **Cluster Instances** regions:

Figure 6–3 BI Publisher Scheduler Configuration Page

JMS Processor	Enable	Number Threads
JobProcessor	<input checked="" type="checkbox"/>	
ReportProcessor	<input checked="" type="checkbox"/>	
EmailProcessor	<input checked="" type="checkbox"/>	
FileProcessor	<input checked="" type="checkbox"/>	
FTPProcessor	<input checked="" type="checkbox"/>	
PrintProcessor	<input checked="" type="checkbox"/>	
WebDavProcessor	<input checked="" type="checkbox"/>	
FaxProcessor	<input checked="" type="checkbox"/>	

3. Enter the **Number Threads** value in the processor configuration table.

Note that all processors will automatically be set to use the number of threads defined in the **Threads Per Processor** value of the **JMS Configuration** region of the page. Enter a value in the **Number Threads** column to change the value from this default.

6.6 Scaling Out Report Processors

To add managed servers to enable more report and delivery processors to handle your reporting throughput see the "Scaling Out the Oracle Business Intelligence Domain" chapter in the *Oracle Fusion Applications Customer Relationship Management Enterprise Deployment Guide*.

After performing the scale-out procedure, follow the steps in "[Managing Report Processing and Delivery Server Load](#)" to configure the processor threads for each cluster instance.

6.7 Diagnosing Report Processing Issues

The following tools enable you to diagnose report processing and delivery issues:

- Scheduler Diagnostics page
- Oracle Enterprise Manager Fusion Applications Control (Fusion Applications Control)

6.7.1 Using the Scheduler Diagnostics Page

The Scheduler Diagnostics page provides the runtime status of the scheduler. It provides status of its JMS configuration, JMS queues, cluster instances, scheduler Database, Toplink, and Oracle Enterprise Scheduler.

To access the Scheduler Diagnostics page:

1. Navigate to the Oracle Business Intelligence Publisher **Administration** page.
2. In the **System Maintenance** group, click **Scheduler Diagnostics**.

The Scheduler Diagnostics page is shown in [Figure 6-4](#):

Figure 6-4 BI Publisher Scheduler Diagnostics Page

The screenshot shows the 'Administration > Scheduler Diagnostics' page. It features a 'System Maintenance' section with tabs for 'Server Configuration', 'Scheduler Configuration', and 'Scheduler Diagnostics'. A 'Diagnose' button is visible. The main content area displays a table with the following data:

Diagnostic Item	Value	Status	Details
Enterprise Scheduler		Passed	
--JMS		Passed	
----JMS Cluster Config	/scratch/apphome/xmlpserver/repository/Admin/Scheduler/jms_cluster_config.properties	Passed	
-----JMS_PROVIDER_TYPE	WebLogic	Info	WebLogic JMS is selected.
-----JMS_WEBLOGIC_VERSION	10.3	Info	
-----JMS_WEBLOGIC_JNDI_FACTORY	weblogic.jndi.WLInitialContextFactory	Info	
-----JMS_WEBLOGIC_JNDI_URL	t3://example.com:7001	Info	
-----BIP_CONNECTION_FACTORY_NAME	BIP.JMS.CF	Info	
-----BIP_SYSTEM_TOPIC_NAME	BIP.System.T	Info	
-----BIP_BURST_JOB_QUEUE_NAME	BIP.Burst.Job.Q	Info	
-----BIP_BURST_REPORT_QUEUE_NAME	BIP.Burst.Report.Q	Info	
-----BIP_DELIVERY_EMAIL_QUEUE_NAME	BIP.Delivery.Email.Q	Info	
-----BIP_DELIVERY_FILE_QUEUE_NAME	BIP.Delivery.File.Q	Info	
-----BIP_DELIVERY_FTP_QUEUE_NAME	BIP.Delivery.FTP.Q	Info	
-----BIP_DELIVERY_PRINT_QUEUE_NAME	BIP.Delivery.Print.Q	Info	
-----BIP_DELIVERY_WEBDAV_QUEUE_NAME	BIP.Delivery.WebDAV.Q	Info	
-----BIP_DELIVERY_FAX_QUEUE_NAME	BIP.Delivery.Fax.Q	Info	
-----NUMBER_THREADS_PER_PROCESSOR	5	Info	
-----JMS_Shared_Temp_Directory	/scratch/apphome/xmlpserver/wljms/shared	Info	
----JMS Runtime		Passed	
-----Topic - BIP.System.T		Passed	

For more information about this page see the "Scheduler Diagnostics" topic in the *Oracle Fusion Middleware Administrator's Guide for Oracle Business Intelligence Publisher (Oracle Fusion Applications Edition)*.

6.7.2 Using Fusion Applications Control

For information on using Fusion Applications Control to diagnose issues with Oracle BI Publisher, see the "Diagnosing and Resolving Issues in BI Publisher" topic in the *Oracle Fusion Middleware Administrator's Guide for Oracle Business Intelligence Publisher (Oracle Fusion Applications Edition)*.

6.8 Configuring System Properties for Reporting

Use the Oracle BI Publisher **Runtime Configuration** page to set the system-level runtime properties for reports.

To access the **Runtime Configuration** page:

1. Navigate to the Oracle Business Intelligence Publisher **Administration** page (see ["Navigating to the Oracle BI Publisher Administration Page"](#)).
2. In the **Runtime Configuration** group, click **Properties**.

For more information about this page see the "Defining Runtime Configurations" chapter in the *Oracle Fusion Middleware Administrator's Guide for Oracle Business Intelligence Publisher (Oracle Fusion Applications Edition)*.

Managing Search with Oracle Enterprise Crawl and Search Framework

This chapter describes how to manage search using Oracle Enterprise Crawl and Search Framework (ECSF).

This chapter includes the following topics:

- [Introduction to Oracle Fusion Applications Search](#)
- [Validating the Environment for ECSF](#)
- [Getting Started with Managing ECSF with Fusion Applications Control](#)
- [Administering Search](#)
- [Performing Advanced Administration Tasks](#)

7.1 Introduction to Oracle Fusion Applications Search

Knowledge workers can spend up to 24% of their time looking for information. Good search tools help them find this information more quickly. Expectations about the ease of use of search tools are high because searches are so easy to perform on the Internet. Oracle Fusion Applications Search is designed to exceed user expectations.

Oracle Fusion Applications Search is embedded within Oracle Fusion Applications. Any application that connects to multiple data sources or manages a significant amount of unstructured (non-database) information—or both—needs advanced search capabilities so that application users can easily locate and take action on data that is relevant to them. Powered by the Oracle Enterprise Crawl and Search Framework and Oracle Secure Enterprise Search, Oracle Fusion Applications Search not only provides a seamless search experience to the Oracle Fusion Applications end user for getting the right content at the right time, but also enhances the Oracle Fusion Applications developer's experience by providing a metadata-driven, declarative design time and runtime interface to manage enterprise search capabilities.

7.1.1 Oracle Fusion Applications Search Use Case Example

Simone Sterling is a new sales representative at Cobalt Inc. a computer hardware company. She is taking over for Jim Anderson, who just left the company. Simone needs to learn about Jim's interactions with his customers and about how to sell Cobalt's new line of energy-saving displays.

Simone starts by entering "Jim Anderson" in the Oracle Fusion Applications Search box that is available at the top of every Oracle Fusion Applications page. Simone retrieves hundreds of documents pertaining to Jim Anderson and to people named

"Jim" or "Anderson". Simone is only interested in documents related to sales or customers, so she de-selects all the search categories except Sales and Customers and then re-runs her search.

Now that Simone is only looking at potentially relevant records, she decides to start exploring the results for specific customers. Simone expands the Customer Account node under the Customer node in the Available Filters tree to display a list of customer accounts. Simone can only view the search results for documents to which she has been granted access. Likewise, she can only view the search categories and filters to which she has access.

Simone clicks "ABC Corporation" to display only the records that involve both Jim Anderson and ABC Corporation. Simone hovers her mouse pointer over result items for Customer Notes and Customer Interactions to see the first 244 characters of each document. She drills down to view details for a note about the new energy-saving displays.

Simone continues to explore her search results. She filters the results by selecting new customers from the Customer Accounts node, views result items in the Search Results window, then drills into Oracle Fusion Applications to view details for interesting items.

Simone finds a particularly useful set of results by searching for the new energy-saving displays, then filtering for recent customer interactions. She decides to save her search so she can monitor customer interactions around the energy-saving displays, so she clicks the Save... button and names her search "Energy Saver Interactions". This enables her to re-run her search at any time from any page within Oracle Fusion Applications.

Thus, Simone reaps the benefits of the powerful Oracle Fusion Applications Search feature that helps her tailor the search results to meet her specific business needs. She can also save the searches that she wants to reuse often.

7.1.2 Oracle Fusion Applications Search UI

The user interface for Oracle Fusion Applications Search, shown in [Figure 7-1](#), offers a variety of useful features.

Figure 7-1 Oracle Fusion Applications Search UI



In the preceding figure, the following features are called out:

- The **Search Categories** dropdown list allows Oracle Fusion Applications users to scope their searches to one or more selected categories. Each search category corresponds to one Oracle Secure Enterprise Search (Oracle SES) data source group. Search categories (Oracle SES data source groups) are used for display purposes and for organizing searchable objects (Oracle SES data sources).
- The **Keywords** field allow users to enter terms or Oracle SES-supported operators for search.
- **Tabbed Browsing** organizes the search results for Oracle Fusion Applications and Oracle Business Intelligence (BI) Catalog into separate tabs. In Oracle Fusion Applications V1, Oracle Fusion Applications Search federates queries to the Oracle SES instance that searches the Oracle Business Intelligence (BI) Catalog. Search results from the Oracle Fusion Applications Oracle SES instance appear on the Oracle Fusion Applications tab. Search results from the Oracle Business Intelligence (BI) Catalog Oracle SES instance appear on the Business Intelligence tab.
- **Faceted Browsing** allows Oracle Fusion Applications users to filter their search results by selecting nodes in the Available Filters tree. Applied filters are shown in bold in the Available Filters tree and as tiles in the Selected Filters region. The numbers in parentheses next to the applied filters in the Available Filters tree represent the number of documents in the search result set that have that facet value.

Users can filter by category (Oracle SES data source group), subcategory (searchable object /Oracle SES data source), and values of one or more facets defined for a given subcategory. Filtering on a facet value automatically filters on its parent subcategory and category.

Filters can be removed via the delete icons in the Available Filters tree and the Selected Filters region. Removing a parent filter also removes all its children.

- Saved Searches allows Oracle Fusion Applications users to save the searches they create. Saved searches comprise all characters in the Keywords field and all the filters the user applies. Selecting a saved search runs the query and display all filters in the Available Filters tree and the Selected Filters region.
- The Main Link identifies the object or page to which a search result item links. It is the primary link and search result item title, as defined in the Default Action Title by the Oracle Fusion Applications developer.
- The Title provides key information about each search result item in addition to the main link. The title is part of the search result item's fixed content that appears regardless of where the user's keywords match the searchable document.
- The Body of each search result item includes a context snippet that comprises one or more instances of the keywords and the surrounding text extracted from the body of the searchable document.
- The Actionable Results are the list of the action links that appear below the body of each search result item. The action links allow Oracle Fusion Applications users to navigate directly to the pages for the objects represented by the search result items so that they can take specific actions on those objects. Action links can also take the user to third party web sites or any page accessed by a fully qualified URL.

7.1.3 Oracle Enterprise Crawl and Search Framework

Oracle Fusion Applications Search functionality is fundamentally made possible by the integration of three systems, each playing a role in forming the complete search platform:

- Search engine, as provided by Oracle Secure Enterprise Search (Oracle SES), which provides the fundamental search capability that includes indexing, querying, as well as some value added functionalities such as security. For more information about Oracle SES, see *Oracle Secure Enterprise Search Administrator's Guide*.
- Source system, such as a relational database, where the searchable information resides
- Search development framework, such as Oracle Enterprise Crawl and Search Framework (ECSF), that supports the integration of applications with search engines

ECSF is an Oracle Fusion Middleware search framework that enables the exposure of application context information on various business objects to enable full-text transactional search. Benefits of ECSF include:

- Transparent integration of applications with search engines, which minimizes development time and maximizes the user experience with search
- Code reuse, through use of a well designed set of abstract classes, to reduce long design cycles
- Basic platform for developing search, which helps new developers to grasp the conceptual flow of work easily

- Centralized process and control mechanism, which enhances search functionality
- Wide range of optimizations that offer better control to leverage search results

7.1.3.1 ECSF Management Features

ECSF management features include:

- Runtime server, a metadata-driven runtime engine that serves as an integration framework between enterprise data sources and Oracle SES. It enables crawling, indexing, and the security service. It also serves as the semantic engine that provides "smart" search features, such as faceted navigation, actionable results, and related search.
- Oracle Enterprise Manager Fusion Applications Control (Fusion Applications Control), an administration user interface for configuring and administering the ECSF runtime server, managing the searchable object lifecycle, and synchronizing with Oracle SES. Support for a command line administration option is also provided. For information, see the "ECSF Command Line Administration Utility" appendix in the *Oracle Fusion Applications Developer's Guide*.

7.1.3.2 Key ECSF Features

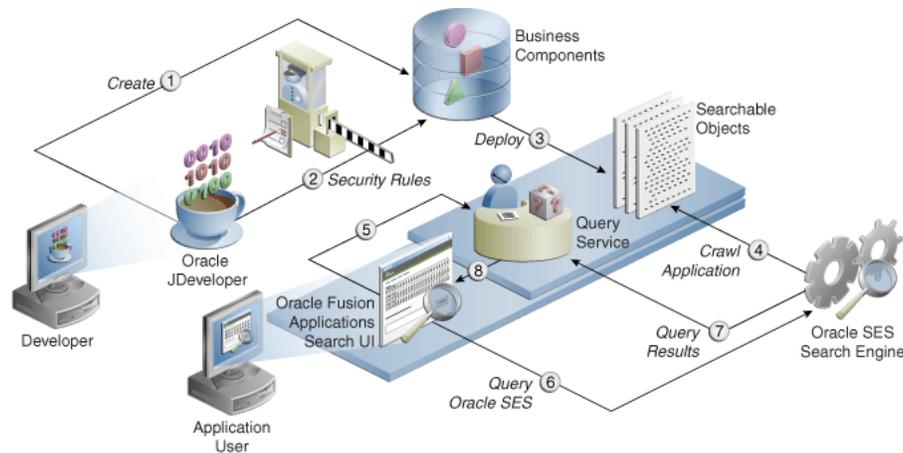
Key ECSF features that are built on top of Oracle SES and enhance the Oracle Fusion Applications user experience with search include:

- Basic search, which allows query based on keyword and search category.
- Advanced search, which allows query based on keyword, search category, and up to 100 attribute filters.
- Faceted navigation, which allows the filtering of search results based on attributes of the business objects. Users can navigate a search result set based on a set of predefined facets, or dimensions. This feature returns a list of facets and their associated set of available values with the search result. Users can select a value for each facet, which is then submitted with the search query in order to narrow down the result set.
- Actionable results, which are search results with action links associated with the searchable objects. From the search results users can either go straight to the page displaying the record they selected, or they can invoke a specific task on a search result.
- Saved searches, which allows saved search criteria for later use. Users can create new saved search entries, edit and delete existing saved search entries, and retrieve user-specified or public saved search entries.
- File attachments, which allow the crawling of attachments that are associated with Oracle Fusion Applications transactional objects or records.
- Crawling Oracle WebCenter Portal tags, which supports crawling searchable objects that contain Oracle WebCenter Portal tags.
- Crawling tree structures, which supports search functionality on source systems containing data that is organized in a tree structure (for example, Oracle Business Intelligence (BI) Catalog).
- Search support for external data sources, which allows querying against search groups that contain external data sources, which are non-ECSF related data sources, such as wiki pages and blogs, that are directly crawled by Oracle SES.

7.1.3.3 ECSF Process Flow

The process flow for search using the ECSF architecture, illustrated in [Figure 7-2](#), begins from the time the Oracle Fusion Applications developer makes a view object searchable to the time the Oracle Fusion Applications user receives the search results. The arrows indicate the direction of flow, and the numbered circles indicate the sequence of the steps, which correspond to the list immediately following the illustration.

Figure 7-2 ECSF Process Flow



The process flow includes the following steps:

1. The Oracle Fusion Applications developer creates a searchable object.
2. The Oracle Fusion Applications developer configures the security rules.
3. The Oracle Fusion Applications administrator deploys the searchable objects to the ECSF runtime environment by registering them in the Oracle Fusion Applications database.
4. The Oracle Fusion Applications administrator deploys the searchable objects to Oracle SES. The Oracle SES search engine crawls the searchable objects in the Oracle Fusion Applications and indexes them as documents.
5. The Oracle Fusion Applications administrator creates search categories containing the searchable objects and deploys them to Oracle SES. When the Oracle Fusion Applications user issues a search query, the Oracle Fusion Applications Search UI performs the query through the ECSF Query Service.
6. The Query Service applies the Oracle Fusion Applications-specific rules and sends the query to the Oracle SES search engine.
7. The Oracle SES search engine returns the result set to the Query Service.
8. The Query Service applies post query security rules and constructs the search hits, and the Oracle Fusion Applications Search UI renders the search results for the Oracle Fusion Applications user.

7.1.4 Item Business Object and Data Level Security

Oracle SES does not apply data level security to Item business objects to determine if the object can be viewed by the user. All qualified users will see a limited amount of information when the Item business object is enabled for search. The information displayed in the search results includes the item number and description. The Oracle

Product Information Management application checks the data level security when the user elects to view item details through the links within the search results.

Review your internal security policies prior to enabling the Item business object for search to determine if item information should be viewed by the end users.

7.2 Validating the Environment for ECSF

Before you begin to manage search with ECSF, make sure that the environment is set up properly for using ECSF.

To validate the ECSF setup, follow the procedures in the following tasks:

- [Task 1, "Make Sure That Oracle Fusion Applications Includes Search Functionality"](#)
- [Task 2, "Make Sure That Oracle SES Is Installed and Configured Properly"](#)
- [Task 3, "Make Sure That Fusion Applications Control Is Available"](#)
- [Task 4, "Provide Access to ECSF Pages in Fusion Applications Control"](#)
- [Task 5, "Validate the Application Identities"](#)

Task 1 Make Sure That Oracle Fusion Applications Includes Search Functionality

Oracle Fusion Applications Search should be embedded within Oracle Fusion Applications, but it must be enabled in the user interface by setting the profile option `FUSION_APPS_SEARCH_ENABLED` to `Y`. For information, see [Task 6, "Enable the Oracle Fusion Applications Search UI"](#).

To make sure that Oracle Fusion Applications includes search functionality:

1. Log in to Oracle Fusion Applications. If you cannot log in to Oracle Fusion Applications, contact your installation team.
2. Verify that the **Enterprise Search** box is available at the top of every Oracle Fusion Applications page.
3. View the **Search Categories** dropdown list. There should be no search categories listed.
4. Log out.

Task 2 Make Sure That Oracle SES Is Installed and Configured Properly

Oracle SES provides the fundamental search capability that includes crawling, indexing, and querying. For more information about Oracle SES, see *Oracle Secure Enterprise Search Administrator's Guide*.

To make sure that Oracle SES is installed and configured properly:

1. Check the administration endpoint by logging in to the Oracle SES Administration GUI with the administration username and password at the following URL.

```
http://host_name:7777/search/admin/index.jsp
```

The default port number is 7777. Make sure that you use the correct port number for your installation. If you cannot access the Oracle SES search engine, contact your installation team.

2. Make sure that the Oracle SES identity plug-in has been registered.

3. Make sure that the federated trusted entities are created. Depending on what product families are installed, you should see one to three proxy users listed. The valid values are:
 - FUSION_APPS_CRM_ECSF_SEARCH_APPID
 - FUSION_APPS_FSCM_ECSF_SEARCH_APPID
 - FUSION_APPS_HCM_ECSF_SEARCH_APPID

Task 3 Make Sure That Fusion Applications Control Is Available

Fusion Applications Control must be available for configuring and administering the ECSF runtime server, managing the searchable object lifecycle, and synchronizing with Oracle SES.

To make sure that Fusion Applications Control is available:

1. Log in to Oracle Enterprise Manager.
2. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
3. Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page, shown in [Figure 7-3](#).

The search engine types (Oracle SES) should be listed.
4. Click the Oracle SES search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page, and validate the Oracle SES search engine instance parameters.
5. From the table of search engine instances, select a search engine instance record, and then select the **Searchable Objects** tab to view the table of searchable objects, and validate the list of searchable objects for the application. For a list of seeded searchable objects, see [Appendix C](#).
6. Select the **Search Categories** tab to view the table of search categories, and validate the list of search categories and objects associated with the search categories for the application. For a list of seeded search categories, see [Appendix C](#).
7. From the navigation pane, re-select the application to open the Enterprise Crawl and Search Framework Configuration Settings page, then click the **Search Application Service Component** link to open the Search Application Service Component administration page, and validate that the search applications for the product families are installed.

Task 4 Provide Access to ECSF Pages in Fusion Applications Control

In order to access the ECSF pages in Fusion Applications Control, users must have Operator privileges in Oracle WebLogic Server. You must add the users to the Operator group and above on Oracle WebLogic Server. For information, see *Oracle Fusion Middleware Securing Oracle WebLogic Server*.

Task 5 Validate the Application Identities

Oracle Fusion Applications include seven search-related application identities that are seeded and are stored in the identity store:

- FUSION_APPS_CRM_SES_CRAWL_APPID
- FUSION_APPS_CRM_ECSF_SEARCH_APPID

- FUSION_APPS_FSCM_SES_CRAWL_APPID
- FUSION_APPS_FSCM_ECSF_SEARCH_APPID
- FUSION_APPS_HCM_SES_CRAWL_APPID
- FUSION_APPS_HCM_ECSF_SEARCH_APPID
- FUSION_APPS_ECSF_SES_ADMIN_APPID

ECSF is powered by Oracle SES. To integrate with Oracle SES, a number of integration identities known as application identities are used. For each Oracle Fusion Applications application, there are a pair of application identities, for example, FUSION_APPS_HCM_SES_CRAWL_APPID and FUSION_APPS_HCM_ECSF_SEARCH_APPID. The CRAWL application identities are used by Oracle SES to interact with ECSF for crawling and security requests, while the SEARCH application identities are used by Oracle SES to query Oracle SES as proxy users.

FUSION_APPS_ECSF_SES_ADMIN_APPID is the application identity used by ECSF to integrate with Oracle SES for administration tasks, such as deployment, scheduling, and so on.

Application identities are provisioned as users in the Oracle Fusion Applications identity store. They often have high level privileges, and their credentials are generated and stored in the credential store. These users are used mainly for machine to machine (application to application) integration.

The Lightweight Directory Access Protocol (LDAP) credential store stores the passwords for the identities that Oracle Fusion Applications and ECSF uses to retrieve passwords for Oracle SES integration.

View the LDAP credential store to make sure the application identities exist.

7.3 Getting Started with Managing ECSF with Fusion Applications Control

Fusion Applications Control is an administration user interface that can be used for configuring and administering the ECSF runtime server, managing the searchable object lifecycle, and synchronizing with Oracle SES.

ECSF also provides support for a command line administration option. The ECSF Command Line Administration Utility offers a command line interface, and has no external dependencies on Oracle Enterprise Manager. The utility is available in the ECSF Library (`ecsfc.jar`) and can be started as a Java program. For more information, see the "ECSF Command Line Administration Utility" appendix in the *Oracle Fusion Applications Developer's Guide*.

Both Fusion Applications Control and the ECSF Command Line Administration Utility require connection to the Oracle Fusion Applications database.

Note: While both Fusion Applications Control and the ECSF Command Line Administration Utility can be used concurrently, it is recommended that ECSF be administered using one or the other to facilitate data control.

Key points of administration include:

- ECSF Configuration: ECSF currently supports only the Oracle SES search engine and its different versions. The Oracle SES search engine with its versions are

stored in the `ECSF_SEARCH_ENGINE_TYPE` table, which resides in the `FUSION` schema, and is rendered in Fusion Applications Control. Every Oracle SES search engine instance must be configured using a set of Oracle SES parameters. Every instance created for search must also be configured using a set of ECSF parameters. Both sets of parameters and their values are stored in the `ECSF_PARAMETER` table and are rendered in Fusion Applications Control. Parameter names are read-only, while parameter values can be updated.

- ECSF Administration
 - Search engine instances and parameters: ECSF supports more than one search engine instance of a given search engine type (currently only Oracle SES). It supports the declarative distribution of index over several search engine instances. Each search engine instance belongs to exactly one version of the Oracle SES search engine type. All the supported search engines instances are stored in the `ECSF_ENGINE_INSTANCE` table and managed using Fusion Applications Control.
 - Searchable objects: Searchable object definitions are configured at design time, and the information is stored in metadata XML files. This information is imported into the Oracle Fusion Applications database and rendered in Fusion Applications Control in the form of searchable objects.
 - Search categories: The metadata for search categories, also called search groups (for Oracle SES), is collected and rendered in Fusion Applications Control.
 - External search categories: Fusion Applications Control enables you to import external search categories (registered Oracle SES data source groups that are crawled and created directly using Oracle SES Administration UI) to be used directly for querying.
 - Index schedule: The index schedule is created in Fusion Applications Control. A corresponding Oracle SES index schedule is then created.

Security administration for ECSF is handled outside the framework and Fusion Applications Control.

For general information about Fusion Applications Control, see [Chapter 2](#).

7.3.1 Navigating to the ECSF Administration Pages

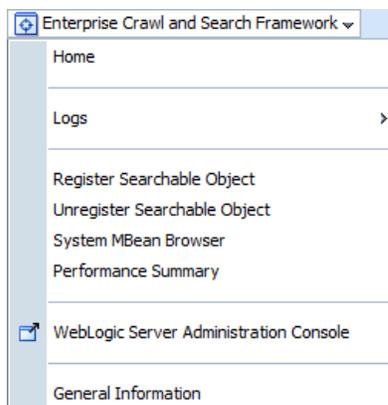
To navigate to the ECSF administration pages in Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder, and select the application engine instance that contains the searchable objects you want to manage.



The Enterprise Crawl and Search Framework Configuration Settings page displays. For information about the Configuration Settings page, see [Section 7.3.2](#).

2. Expand the **Enterprise Crawl and Search Framework** target menu.



The menu displays the options described in [Table 7-1](#).

Table 7-1 Enterprise Crawl and Search Framework Target Menu Options

Menu Option	Description
Home	Displays the Enterprise Crawl and Search Framework Configuration Settings page. For more information about this page, see Section 7.3.2 .
Logs	Displays the options to either view log messages or access the Log Configuration page to configure basic and advanced log configuration settings. For more information, see "Managing Log Files and Diagnostic Data" chapter in <i>Oracle Fusion Middleware Administrator's Guide</i> . The log files are stored in the following directory: <code>DOMAIN_HOME/sysman/log/emoms.log</code>
Register Searchable Object	Displays the Register Searchable Object page. For information about the Register Searchable Object task, see Task 2, "Register the Searchable Objects" .
Unregister Searchable Object	Displays the Unregister Searchable Object page. For information about the Unregister Searchable Object task, see Task 4, "Unregister the Searchable Objects" .
System MBean Browser	Displays the MBean browser that allows to you browse the MBeans for the selected application.
Performance Summary	Displays performance information for the selected component or application. For information, see Section 10.4.1 .
WebLogic Server Administration Console	Launches the Oracle WebLogic Server Administration Console in a separate browser window.
General Information	Displays general information about the application instance.

7.3.2 Understanding the Enterprise Crawl and Search Framework Configuration Settings Page

The Enterprise Crawl and Search Framework Configuration Settings page, shown in [Figure 7-3](#), displays the search engine types and their corresponding search engine instance parameters. This page is your starting place for administering and managing search.

Figure 7-3 Enterprise Crawl and Search Framework Configuration Settings Page

To manage and configure ECSF Service Component Application, use this link [ECSF Service Component](#)

ECSF Configuration Settings

Click the search engine type for which you want to administer the search engine instances.

Search Engine Types:

Name	Version	Connector Class Name
Oracle SES	11.1.2.0.0	oracle.ecsf.ses.v11g.impl.SearchEngineImpl
Oracle SES	10.1.8.4	oracle.ecsf.ses.impl.SearchEngineImpl

Columns Hidden

Name	Required
SES_ADMIN_SERVICE	✓
SES_ADMIN_USERNAME	✓
SES_ADMIN_PASSWORD	✓
SES_ADMIN_SESSION_TIMEOUT	✓
SES_QUERY_SERVICE	✓
SES_QUERY_PROXY_USERNAME	✓
SES_QUERY_PROXY_PASSWORD	✓
SES_QUERY_SESSION_TIMEOUT	✓
ECSF_DATA_SERVICE	✓

From the Enterprise Crawl and Search Framework Configuration Settings page, you can navigate to the Search Application Service Component Administration page to create and configure Search Application Service Component instances.

You can also navigate to the Search Engine Instance Administration page (by selecting a search engine type) to define search engine instances, manage data sources and schedules, and extend search.

7.4 Administering Search

Administer search by managing data sources and schedules and modifying data sources.

Note: Administration of searchable objects, search categories, and index schedules must be performed using Fusion Applications Control. You can use the Oracle SES Administration GUI for all other Oracle SES related administration tasks.

To access the ECSF pages in Fusion Applications Control, you must have Operator privileges in Oracle WebLogic Server. For information, see [Task 4, "Provide Access to ECSF Pages in Fusion Applications Control"](#).

You can perform the following procedures to administer search using Fusion Applications Control:

- [Making Seeded Searchable Objects Available for Search](#)
- [Managing Index Schedules](#)
- [Changing the Association of Searchable Objects with Search Categories](#)

- [Renaming Search Categories](#)
- [Deleting Search Categories](#)
- [Making External Search Categories Available for Federated Search](#)
- [Deleting External Search Categories](#)
- [Modifying the Display Name of Deployed Searchable Objects](#)
- [Undeploying Searchable Objects](#)
- [Unregistering Searchable Objects from the Oracle Fusion Applications Database](#)
- [Modifying Search Application Service Component Instance Parameters](#)

In order to avoid disrupting running schedules, it is highly recommended that you administer search when the index schedules are not running.

7.4.1 Making Seeded Searchable Objects Available for Search

The Oracle Fusion Customer Relationship Management, Oracle Fusion Human Capital Management, and Oracle Fusion Supply Chain Management product families each has its own set of seeded searchable objects that are packaged into its corresponding search application. For example, the seeded searchable objects for Oracle Fusion Customer Relationship Management are packaged in the Oracle Fusion Customer Relationship Management search application. For more information, see [Appendix C](#).

Note: While Oracle Fusion Financials searchable objects, search categories, and index schedules are seeded, they are not available for search. Do not use Oracle Fusion Financials searchable objects, search categories, and index schedules to enable search.

To manage the searchable objects for a particular product family, you must select the search application corresponding to that product family.

To make seeded searchable objects available for search, follow the procedures in the following tasks:

- [Task 1, "Deploy the Searchable Objects to the Oracle Secure Enterprise Search \(Oracle SES\) Instance"](#)
- [Task 2, "Activate the Searchable Objects"](#)
- [Task 3, "Deploy the Search Categories"](#)
- [Task 4, "Deploy the Index Schedules"](#)
- [Task 5, "Start the Index Schedules"](#)
- [Task 6, "Enable the Oracle Fusion Applications Search UI"](#)
- [Task 7, "Validate That Searchable Objects Are Available for Search"](#)

Task 1 Deploy the Searchable Objects to the Oracle Secure Enterprise Search (Oracle SES) Instance

Deploying searchable objects to the Oracle Secure Enterprise Search (Oracle SES) instance makes the objects available for the search engine instance to crawl. The searchable objects deployed to Oracle SES must have a unique and fully qualified name, for example, `oracle.apps.crm.Opportunity` or `oracle.apps.hcm.Opportunity`.

To deploy searchable objects to the Oracle SES instance using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
2. Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
3. Click the **Oracle SES** search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page, shown in [Figure 7-4](#).

Figure 7-4 Search Engine Instance Administration Page

Administration: Search Engine Instance

Define search engine instances for the selected search engine type. You can create, modify, and delete engine instances, and configure search categories.

Search Engine Type: **Oracle SES**

Actions View Register Refresh Detach

Name	Comments
CRM Search Engine Instance	Instance for CRM Searchable Objects
FSCM Search Engine Instance	Instance for FSCM Searchable Objects
HCM Search Engine Instance	Instance for HCM Searchable Objects
Help Instance	Instance for Help Portal Objects

Columns Hidden

Parameters Searchable Objects Search Categories External Search Categories Index Schedules

Actions View Detach Deploy Parameters

Name	Required	Value
SES_ADMIN_SERVICE	✓	http://hostserver:5270/search/api/admin/AdminService
SES_ADMIN_USERNAME	✓	searchsys
SES_ADMIN_PASSWORD	✓	•••••
SES_ADMIN_SESSION_TIMEOUT	✓	60
SES_QUERY_SERVICE	✓	http://hostserver:5270/search/query/OracleSearch
SES_QUERY_PROXY_USERNAME	✓	scott
SES_QUERY_PROXY_PASSWORD	✓	•••••
SES_QUERY_SESSION_TIMEOUT	✓	60

(Columns Hidden) 4

4. From the table of search engine instances, select the record of the search engine instance that contains the searchable object you want to deploy, and then select the **Searchable Objects** tab, shown in [Figure 7-5](#), to view the table of searchable objects for the selected search engine instance.

Figure 7–5 Searchable Objects Tab

Display Name	Searchable Object Name	Version	Deployed	Active	Application ID	Parameters Applied	Customized
ECM Contracts	oracle.apps.contracts.dash...	1	✓	✓	contracts	✓	<input type="checkbox"/>
This table lists the Searchable Objects for the selected Instance							
Contract Documents	oracle.apps.contracts.dash...	1	✓	✓	contracts	✓	<input type="checkbox"/>
Sourcing Contracts	oracle.apps.contracts.dash...	1	✓	✓	contracts	✓	<input type="checkbox"/>
Contracts	oracle.apps.contracts.dash...	1	✓	✓	customer	✓	<input type="checkbox"/>
References	oracle.apps.sales.baseSale...	1	✓	✓	sales	✓	<input type="checkbox"/>
Opportunities	oracle.apps.sales.opptyMg...	1	✓	✓	sales	✓	<input type="checkbox"/>

- Click the **Deploy** icon. Alternatively, select **Deploy** from the **Actions** menu or use the Ctrl+Shift+P keyboard shortcut.

A green check mark appears in the Deployed column of the selected searchable object to indicate that the source has been created in Oracle SES.

Note: When a searchable object is registered to ECSF, by default it is inactive and not deployed. You can deploy the searchable object to the search engine instance regardless of whether the searchable object is active or inactive. Once deployed successfully, the searchable object is flagged as deployed and inactive until you activate it. Once a searchable object is deployed, the search engine instance can crawl the searchable object. However, a query of that object will return no results until the object is activated.

Task 2 Activate the Searchable Objects

A query of deployed and crawled searchable objects will return results only if the objects are activated.

To activate searchable objects using Fusion Applications Control:

- From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
- Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
- Click the Oracle SES search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page.
- From the table of search engine instances, select the record of the search engine instance that contains the searchable object you want to activate, and then select the **Searchable Objects** tab, shown in to view the table of searchable objects for the selected search engine instance.
- Select the desired searchable object from the table, and click the **Activate** icon. Alternatively, select **Activate** from the **Actions** menu or use the Ctrl+Shift+T keyboard shortcut.

A green check mark appears in the Active column of the selected searchable object to indicate that the searchable object is available for search.

Note: A query of this object will return results only if the object has been deployed to a search engine instance and has been crawled.

Task 3 Deploy the Search Categories

Deploying search categories pushes them to the Oracle SES instance. Only search categories with associated searchable objects that are activated can be deployed. Only deployed search categories appear in the Search Categories dropdown list in the Oracle Fusion Applications Search UI.

To deploy search categories using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
2. Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
3. Click the Oracle SES search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page.
4. From the table of search engine instances, select the record of the search engine instance to which you want to deploy a search category, and then select the **Search Categories** tab to view the table of search categories for the selected search engine instance.
5. Select the record of the search category you want to deploy.
6. Set the scope of the search category to `GLOBAL` to display the search category on the Oracle Fusion Applications Search UI.
7. Click the **Deploy** icon. Alternatively, select **Deploy** from the **Actions** menu or use the `Ctrl+Shift+P` keyboard shortcut.

A green check mark appears in the **Deployed** column of the selected search category to indicate that the source group has been created in Oracle SES.

Note: You cannot deploy a search category that contains a searchable object that is not activated. All searchable objects in the search category must be activated. For more information, see [Task 2, "Activate the Searchable Objects"](#).

The Oracle Fusion Applications Search UI reflects this change after the ECSF cache expires and the user logs out and logs back in. The default is 30 minutes after the change is added to the cache.

Task 4 Deploy the Index Schedules

Deploying index schedules pushes them to the Oracle SES instance. Only index schedules with associated searchable objects can be deployed.

Caution: Deploying index schedules to Oracle SES automatically sets the value for `recrawlPolicy` to `Process Documents That Have Changed in Oracle SES`. Do not change the value of `recrawlPolicy`. ECSF supports only incremental crawling.

To deploy index schedules using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
2. Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
3. Click the Oracle SES search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page.
4. From the table of search engine instances, select the record of the search engine instance for which you want to deploy an index schedule, and then select the **Index Schedules** tab, shown in [Figure 7-6](#), to view the table of index schedules for the selected search engine instance.

Figure 7-6 Index Schedules Tab

Name	Status	Deployed	Frequency Type	Time Between Launches	Schedule Launch Day	Schedule Launch Time
Customers Schedule			hourly	1		
Contacts Schedule	Scheduled	✓	manual			
Customer Contracts Schedule	Scheduled	✓	manual			
Customer References Schedule	Scheduled	✓	manual			
References Schedule	Scheduled	✓	manual			
Opportunities Schedule	Scheduled	✓	manual			

5. Select the desired index schedule from the table, and click the **Deploy Index Schedule** icon. Alternatively, select **Deploy** from the **Actions** menu or use the Ctrl+Shift+P keyboard shortcut.

A green check mark appears in the **Deployed** column of the selected index schedule, and the **Status** column indicates *Scheduled*. This indicates that the schedule has been created in Oracle SES.

Task 5 Start the Index Schedules

Starting an index schedule initiates the crawling and indexing of data. You cannot start an index schedule that has not been deployed. You must first deploy the index schedule to add it to the search engine instance. For more information, see [Task 4, "Deploy the Index Schedules"](#).

To start index schedules using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
2. Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
3. Click the Oracle SES search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page.
4. From the table of search engine instances, select the record of the search engine instance for which you want to start an index schedule, and then select the **Index**

Schedules tab to view the table of index schedules for the selected search engine instance.

5. Select the desired index schedule from the table, and click the **Start Index Schedule** icon. Alternatively, select **Start** from the **Actions** menu or use the Ctrl+Shift+W keyboard shortcut.

The **Status** column of the selected index schedule indicates that the schedule is **Launching**. The date and time appear in the **Start Time** column when you start the index schedule.

You must manually refresh the page to view the status updates.

Task 6 Enable the Oracle Fusion Applications Search UI

After you have configured search and started the crawl, enable the Oracle Fusion Applications Search UI by setting the value of the `FUSION_APPS_SEARCH_ENABLED` profile option to `Y`. This setting makes the search controls visible in the user interface. The default value is `N`, which disables the search controls in the user interface.

To set the `FUSION_APPS_SEARCH_ENABLED` profile option to `Y`:

1. Sign in to Oracle Fusion Applications with a user account that is provisioned with the Manage All Application Profile Values role. Contact your security administrator for details.
2. From the **Administration** menu in the work area of Oracle Fusion Applications, choose **Setup and Maintenance**.
3. From the Setup and Maintenance Overview page, search for profile tasks. The Manage Administrator Profile Values page displays.
4. Edit profile option values.
5. Query the profile option using the code `FUSION_APPS_SEARCH_ENABLED`.
6. Set the value to **Yes** at the site level. This profile option is defined at the site level only.

For information about configuring profile options, see [Section 13.7.5](#).

Task 7 Validate That Searchable Objects Are Available for Search

Make sure that you can see valid search results by running a query on a selected search category.

To validate that searchable objects are available for search using Fusion Applications Control:

1. Make sure that the index schedule you started has completed crawling and indexing the content for search, which is indicated by `Scheduled` in the **Status** column for the index schedule.
2. Log in to your application and expand the **Categories** field in the global area. You should see all the search categories you deployed.
3. Select one search category, enter * (wildcard) in the **Search Term** field, and click the **Play** button to initiate the search.

Search results should appear, confirming the readiness for search.

7.4.2 Managing Index Schedules

You can perform the following tasks to manage index schedules:

- [Creating the Index Schedules](#)
- [Associating the Searchable Objects with Index Schedules](#)
- [Deploying the Index Schedules](#)
- [Starting the Index Schedules](#)
- [Stopping Index Schedules](#)
- [Disassociating Searchable Objects from Index Schedules](#)
- [Starting Full Indexing](#)
- [Undeploying Index Schedules](#)
- [Deleting Index Schedules](#)

Note: Repetitive deployment of index schedules is not supported.

7.4.2.1 Creating the Index Schedules

You can either create new index schedules for the new searchable objects or add the new searchable objects to existing index schedules. Creating an index schedule adds a new record to the `ECSF_INDEX_SCHEDULE` table in the Oracle Fusion Applications database.

To create index schedules using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
2. Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
3. Click the Oracle SES search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page.
4. From the table of search engine instances, select the record of the search engine instance for which you want to create an index schedule, and then select the **Index Schedules** tab, shown in [Figure 7-6](#), to view the table of index schedules for the selected search engine instance.
5. Click the **New Index Schedule** icon. Alternatively, select **New Record** from the **Actions** menu or use the Ctrl+Shift+N keyboard shortcut.

A new record is added to the top of the table of index schedules.

6. Enter a value for the **Name** column.
7. From the **Frequency Type** dropdown list, select the desired value for how often you want the index schedule to run.
8. Complete the **Time Between Launches**, **Schedule Launch Days**, and **Schedule Launch Time** columns that become enabled based on your selection in the **Frequency Type** column. [Table 7-2](#) lists the valid values for the required columns based on the frequency type.

Table 7-2 Index Schedule Frequency Values

Frequency Type	Time Between Launches	Schedule Launch Day	Schedule Launch Time
Manual	---	---	---

Table 7–2 (Cont.) Index Schedule Frequency Values

Frequency Type	Time Between Launches	Schedule Launch Day	Schedule Launch Time
Hourly	1 to 23	---	---
Daily	1 to 99	---	1:00 AM to 12:00 AM
Weekly	1 to 12	Monday to Sunday	1:00 AM to 12:00 AM
Monthly	1 to 12	1 to 31	1:00 AM to 12:00 AM

- Click the **Save Index Schedule** icon to save the selected record. Alternatively, select **Save Record** from the **Actions** menu or use the Ctrl+Shift+S keyboard shortcut.

Caution: A save will occur only for the selected record. You can only save one record at a time.

Note: To proceed with other operations, you must first select any record. Creating or deleting ECSF objects removes the focus from the window.

7.4.2.2 Associating the Searchable Objects with Index Schedules

Associating searchable objects with an index schedule adds the objects to that index schedule. You can only associate each searchable object with one index schedule. Only a searchable object that is not already associated with an index schedule can be added to an index schedule. If the searchable object is already associated with an index schedule, you must disassociate it from that index schedule before you can associate it with another index schedule. For information, see [Section 7.4.2.6](#).

You can only associate searchable objects with index schedules that are not already deployed. If the index schedule has already been deployed, you must undeploy it before you can associate searchable objects with it. For information, see [Section 7.4.2.8](#).

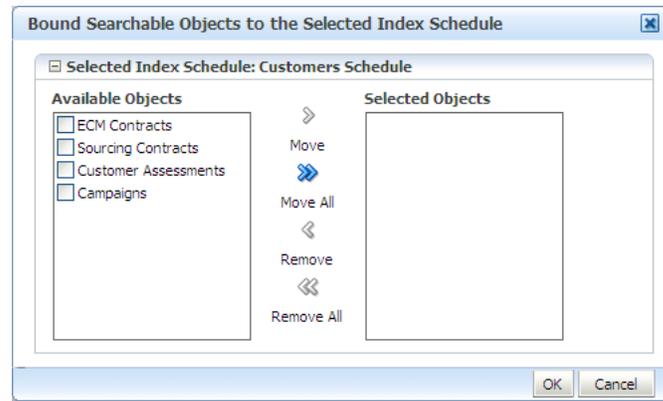
To associate searchable objects with index schedules using Fusion Applications Control:

- From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
- Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
- Click the Oracle SES search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page.
- From the table of search engine instances, select the record of the desired search engine instance, and then select the **Index Schedules** tab, shown in [Figure 7–6](#), to view the table of index schedules for the selected search engine instance.
- Select the desired index schedule to which you want to add searchable objects, and click the **Bind** icon.

The **Bind** icon is disabled if the index schedule is deployed. You cannot associate or disassociate searchable objects if the index schedule is deployed.

- In the Bound Searchable Objects to the Selected Index Schedule dialog, shown in [Figure 7-7](#), select the checkboxes of the desired searchable objects from the **Available Objects** list, and click the **Move** icon to add them to the **Selected Objects** list.

Figure 7-7 Bound Searchable Objects to the Selected Index Schedule Dialog



- Click **OK**.

The `IS_ID` column in the `ECSF_SEARCH_INDEX_OBJECT` table is populated with the Index Schedule ID.

Note: A searchable object can be assigned to only one index schedule. If a searchable object has been assigned to an index schedule, it will not be listed in the table of available searchable objects.

7.4.2.3 Deploying the Index Schedules

Deploying index schedules pushes them to the Oracle SES instance. For information, see [Task 4, "Deploy the Index Schedules"](#).

7.4.2.4 Starting the Index Schedules

Starting an index schedule initiates the crawling and indexing of data. For information, see [Task 5, "Start the Index Schedules"](#).

7.4.2.5 Stopping Index Schedules

Stopping an index schedule that has been started aborts the index process.

To stop index schedules using Fusion Applications Control:

- From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
- Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
- Click the Oracle SES search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page.
- From the table of search engine instances, select the record of the search engine instance for which you want to stop an index schedule, and then select the **Index**

Schedules tab, shown in [Figure 7-6](#), to view the table of index schedules for the selected search engine instance.

5. Select the desired index schedule from the table, and click the **Stop Index Schedule** icon. Alternatively, select **Stop** from the **Actions** menu or use the Ctrl+Shift+O keyboard shortcut.

The **Status** column of the selected index schedule indicates that the schedule has **Failed**, or that the index process was aborted. The date and time appear in the **Stop Time** column when you stop the index schedule.

7.4.2.6 Disassociating Searchable Objects from Index Schedules

Disassociating a searchable object from an index schedule removes it from that index schedule and makes it available to be added to another index schedule. You can only disassociate a searchable object from an index schedule with which it is associated. You can only disassociate a searchable object from an index schedule that is not deployed. If the index schedule has been deployed, you must undeploy it before you can disassociate a searchable object from it. For information, see [Section 7.4.2.8](#).

To disassociating searchable objects from index schedules using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
2. Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
3. Click the Oracle SES search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page.
4. From the table of search engine instances, select the record of the desired search engine instance, and then select the **Index Schedules** tab, shown in [Figure 7-6](#), to view the table of index schedules for the selected search engine instance.
5. Select the desired index schedule from which you want to remove searchable objects, and click the **Bind** icon.

The **Bind** icon is disabled if the index schedule is deployed. You cannot associate or disassociate searchable objects if the index schedule is deployed.

6. In the Bound Searchable Objects to the Selected Index Schedule dialog, shown in [Figure 7-7](#), select the checkboxes of the desired searchable objects from the **Selected Objects** list, and click the **Remove** icon to remove them from the **Selected Objects** list and make them available for adding to an index schedule.
7. Click **OK**.

7.4.2.7 Starting Full Indexing

Full indexing is the process by which the search engine creates indexes from all crawled and indexable documents for a particular data source. In contrast, incremental indexing builds indexes on top of existing indexes.

To start full indexing using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.

2. Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
3. Click the Oracle SES search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page.
4. From the table of search engine instances, select the desired search engine instance record, and then select the **Index Schedules** tab, shown in [Figure 7-6](#), to view the table of index schedules for the selected search engine instance.
5. From the table of index schedules, select the desired schedule for which you want to build a full index.
6. Click the **Start Full Indexing** icon. Alternatively, select **Full Index** from the **Actions** menu or use the Ctrl+Shift+L keyboard shortcut.

A confirmation dialog appears. Click **OK**.

The **Status** column indicates that the full indexing is **Launching**. Indexing will occur only on the active and deployed searchable objects.

7.4.2.8 Undeploying Index Schedules

Undeploying an index schedule removes it from the Oracle SES instance. Only deployed index schedules can be undeployed.

To undeploy index schedules using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
2. Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
3. Click the Oracle SES search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page.
4. From the table of search engine instances, select the record of the search engine instance for which you want to undeploy an index schedule, and then select the **Index Schedules** tab, shown in [Figure 7-6](#), to view the table of index schedules for the selected search engine instance.
5. Select the desired index schedule from the table, and click the **Undeploy Index Schedule** icon. Alternatively, select **Undeploy** from the **Actions** menu or use the Ctrl+Shift+U keyboard shortcut.

The green check mark disappears from the **Deployed** column of the selected index schedule, and **Scheduled** disappears from the **Status** column. This indicates that the schedule has been deleted from Oracle SES.

7.4.2.9 Deleting Index Schedules

Deleting an index schedule disassociates it from the search engine instance and removes its corresponding record from the `ECSF_INDEX_SCHEDULE` table in the Oracle Fusion Applications database. Only undeployed index schedules can be deleted.

To delete index schedules using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.

2. Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
3. Click the Oracle SES search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page.
4. From the table of search engine instances, select the record of the search engine instance from which you want to delete an index schedule, and then select the **Index Schedules** tab, shown in [Figure 7-6](#), to view the table of index schedules for the selected search engine instance.
5. Select the record of the index schedule you want to delete.
6. Click the **Delete Index Schedule** icon. Alternatively, select **Delete Record** from the **Actions** menu or use the Ctrl+Shift+D keyboard shortcut.

The selected record is removed from the table of index schedules.

Note: To proceed with other operations, you must first select any record. Creating or deleting ECSF objects removes the focus from the window.

7.4.3 Changing the Association of Searchable Objects with Search Categories

Changing the association of searchable objects with a search category disassociates the searchable objects from the search category and or associates additional searchable objects with the search category.

To change the association of searchable objects with search categories, follow the procedures in the following tasks:

- [Task 1, "Undeploy the Search Categories"](#)
- [Task 2, "Disassociate the Searchable Objects from Search Categories"](#)
- [Task 3, "Associate Searchable Objects with the Search Categories"](#)
- [Task 4, "Deploy the Search Categories"](#)

Task 1 Undeploy the Search Categories

Undeploying a search category removes it from the Oracle SES instance. Only deployed search categories can be undeployed.

To undeploy search categories using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
2. Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
3. Click the Oracle SES search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page.
4. From the table of search engine instances, select the record of the search engine instance in which you want to undeploy a search category, and then select the **Search Categories** tab, shown in [Figure 7-15](#), to view the table of search categories for the selected search engine instance.
5. Select the record of the search category you want to undeploy.

6. Click the **Undeploy** icon. Alternatively, select **Undeploy** from the **Actions** menu or use the Ctrl+Shift+U keyboard shortcut.

The green check mark disappears from the **Deployed** column of the selected search category to indicate that the source group has been removed from Oracle SES.

Note: The Oracle Fusion Applications Search UI reflects this change after the ECSF cache expires and the user logs out and logs back in. The default is 30 minutes after the change is added to the cache.

Task 2 Disassociate the Searchable Objects from Search Categories

Disassociating a searchable object from a search category removes it from that search category. The searchable object is still available for association to other search categories.

To disassociate searchable objects from search categories using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
2. Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
3. Click the Oracle SES search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page.
4. From the table of search engine instances, select the record of the desired search engine instance, and then select the **Search Categories** tab, shown in [Figure 7-15](#), to view the table of search categories for the selected search engine instance.
5. Select the desired search category from which you want to remove searchable objects, and click the **Bind** icon.

The **Bind** icon is disabled if the search category is deployed. You cannot associate or disassociate searchable objects if the search category is deployed.

6. In the Bound Searchable Objects to the Selected Search Category dialog, shown in [Figure 7-16](#), select the checkboxes of the desired searchable objects from the **Selected Objects** list, and click the **Remove** icon to remove them from the **Selected Objects** list.
7. Click **OK**.

Task 3 Associate Searchable Objects with the Search Categories

Associate searchable objects with a search category to add the objects to that search category. For information, see [Task 8, "Associate the Searchable Objects with Search Categories"](#).

Task 4 Deploy the Search Categories

Deploy the search categories to push them to the Oracle SES instance. For information, see [Task 3, "Deploy the Search Categories"](#).

7.4.4 Renaming Search Categories

Renaming a search category changes the name of the search category.

To change the name of a search category using Fusion Applications Control:

1. If the search category is deployed, you must undeploy it. For information, see [Task 1, "Undeploy the Search Categories"](#).
2. In the **Name** column of the Search Categories tab, edit the value corresponding to the record of the search category you want to modify.
3. Click the **Save Search Category** icon. Alternatively, select **Save Record** from the **Actions** menu or use the Ctrl+Shift+S keyboard shortcut.
4. If the search category was deployed prior to renaming it, redeploy it. For information, see [Task 4, "Deploy the Search Categories"](#).

7.4.5 Deleting Search Categories

Deleting a search category disassociates it from the search engine instance and removes its corresponding record from the ECSF_SEARCH_INDEX_GROUP table in the Oracle Fusion Applications database.

To delete search categories using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
2. Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
3. Click the Oracle SES search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page.
4. From the table of search engine instances, select the record of the search engine instance from which you want to delete a search category, and then select the **Search Categories** tab, shown in [Figure 7-15](#), to view the table of search categories for the selected search engine instance.
5. Select the record of the search category you want to delete.
6. Click the **Delete Search Category** icon. Alternatively, select **Delete Record** from the **Actions** menu or use the Ctrl+Shift+D keyboard shortcut.

The selected record is removed from the table of search categories.

Note: To proceed with other operations, you must first select any record. Creating or deleting ECSF objects removes the focus from the window.

7.4.6 Making External Search Categories Available for Federated Search

External search categories are the registered Oracle SES data source groups that contain external data sources (that is, data sources created using Oracle SES administration). External data sources are not registered in the ECSF_SEARCH_INDEX_OBJECT table. They include data sources such as Intranet, mail, database, and federated data sources (that is, data sources from another search engine instance). Oracle SES data source groups can also contain federated ECSF data sources. For more information, see [Section 7.5.4](#).

External search categories are directly used for querying. Executing the import operation pulls all the external search categories into the Oracle Fusion Applications

database. All external search categories in the Oracle Fusion Applications database are available during ECSF query time through the ECSF query API.

Note: Since external search categories are not managed by ECSF, you must use Oracle SES to modify the external search categories.

To make external search categories available for federated search, follow the procedures in the following tasks:

- [Task 1, "Import the External Search Categories"](#)
- [Task 2, "Set the Application ID Value"](#)
- [Task 3, "Associate the Application ID with an Active Search Application"](#)

Task 1 Import the External Search Categories

Importing external search categories adds new records to the `ECSF_SEARCH_INDEX_GROUP` table in the Oracle Fusion Applications database. The import also automatically sets the `IS_EXTERNAL` flag to `TRUE` to indicate that the search category is an external search category.

To import external search categories using Fusion Applications Control:

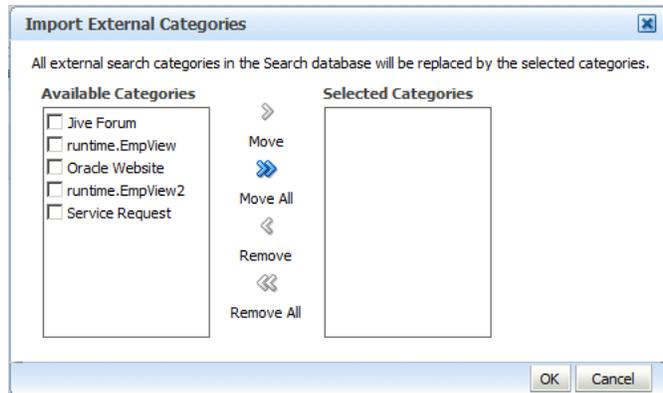
1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
2. Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
3. Click the Oracle SES search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page.
4. From the table of search engine instances, select the record of the search engine instance that points to an Oracle SES instance, and then select the **External Search Categories** tab, shown in [Figure 7–8](#), to view the table of external search categories for the selected search engine instance.

Figure 7–8 External Search Categories Tab

Name	Application ID
runtime.EmpView	<input type="text"/>
runtime.EmpView2	<input type="text"/>
Service Request	<input type="text"/>

5. Click the **Import** button.

The Import External Categories dialog, shown in [Figure 7–9](#), appears.

Figure 7–9 Import External Categories Dialog

6. In the Available Categories column, select the checkbox of the external search categories you want to import and click the **Move** icon to shuttle your selection to the Selected Categories column.
7. Click **OK** to import the selected external search categories.

Note: All existing external search categories in the Oracle Fusion Applications database are replaced by the latest import from Oracle SES.

Clicking **OK** without selecting any external search categories also deletes all the existing external search categories in the Oracle Fusion Applications.

The selected search categories from the selected Oracle SES instance display in the table of external search categories. Fusion Applications Control currently only displays the name of each external search category that is imported.

Clicking **Cancel** cancels the import operation, preserves any existing external search categories in the Oracle Fusion Applications database, and returns you to the External Search Categories page.

Task 2 Set the Application ID Value

Setting the Application ID value allows the external search categories to be queried using the GLOBAL scope, which is the default scope for Oracle Fusion Applications.

To set the Application ID value using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
2. Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
3. Click the Oracle SES search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page.
4. From the table of search engine instances, select the record of the search engine instance that points to an Oracle SES instance, and then select the **External Search Categories** tab, shown in [Figure 7–8](#), to view the table of external search categories for the selected search engine instance.

5. From the table of external search categories, select the record of the external search category that you want to modify.
6. In the **Application ID** column corresponding to the external search category you want to modify, enter an application ID (for example, CRM) based on the desired Service Component to which this category is to be assigned.
7. Click the **Save External Search Category** icon to save the selected record. Alternatively, select **Save Record** from the **Actions** menu or use the Ctrl+Shift+S keyboard shortcut.

Task 3 Associate the Application ID with an Active Search Application

The application ID set for the external search category must be associated with an active search application in order for the external search category to be available for federated search.

To associate the application ID with an active search application, update the ECSF_QUERY_SERVICE_APP_IDS parameter value to include the application ID in the comma separated string. For information, see [Section 7.4.11](#).

7.4.7 Deleting External Search Categories

You can delete individual external categories once they are imported into the Oracle Fusion Applications database. Deleting an external search category removes its corresponding record from the ECSF_SEARCH_INDEX_GROUP table in the Oracle Fusion Applications database, and will make it unavailable for querying.

To delete external search categories using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
2. Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
3. Click the Oracle SES search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page.
4. From the table of search engine instances, select the record of the search engine instance from which you want to delete an external search category from the Oracle Fusion Applications database, and then select the **External Search Categories** tab, shown in [Figure 7-8](#), to view the table of external search categories for the selected search engine instance.
5. Select the record of the external search category you want to delete.
6. Click the **Delete Search Category** icon. Alternatively, select **Delete Record** from the **Actions** menu or use the Ctrl+Shift+D keyboard shortcut.

The selected record is removed from the table of external search categories.

Note: This deletion removes the external search category from the Oracle Fusion Applications database. It does not undeploy the search group from Oracle Secure Enterprise Search (Oracle SES).

7.4.8 Modifying the Display Name of Deployed Searchable Objects

You can change a deployed searchable object's display name and application ID without first having to deactivate and undeploy the searchable object.

Caution: Do not modify the Oracle Fusion Applications Help searchable object named `TopicSearchPVO`.

To change the display name and application ID using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
2. Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
3. Click the Oracle SES search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page.
4. From the table of search engine instances, select the record of the search engine instance that contains the searchable object you want to modify, and then select the **Searchable Objects** tab, shown in [Figure 7-5](#), to view the table of searchable objects for the selected search engine instance.
5. In the **Display Name** column, edit the value corresponding to the searchable object you want to modify.
6. In the **Application ID** column, edit the value corresponding to the searchable object you want to modify.
7. Click the **Save Searchable Object** icon to save the selected record. Alternatively, select **Save Record** from the **Actions** menu or use the Ctrl+Shift+S keyboard shortcut.

7.4.9 Undeploying Searchable Objects

Undeploying searchable objects from the Oracle SES instance makes the objects unavailable for the search engine instance to crawl. While undeployed, the searchable object is still associated with the search engine instance record and can be redeployed to the physical engine. When a searchable object is undeployed, any association to a search category or index schedule is removed.

To undeploy searchable objects from the Oracle SES instance, follow the procedures in the following tasks:

- [Task 1, "Deactivate the Searchable Object"](#)
- [Task 2, "Undeploy the Searchable Object from the Oracle SES Instance"](#)

Task 1 Deactivate the Searchable Object

Deactivated searchable objects are still available for the search engine instance to crawl, but a query of the deactivated searchable objects will return no results.

To deactivate searchable objects on search engine instances using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
2. Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.

3. Click the Oracle SES search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page.
4. From the table of search engine instances, select the record of the search engine instance that contains the searchable object you want to deactivate, and then select the **Searchable Objects** tab, shown in [Figure 7-5](#), to view the table of searchable objects for the selected search engine instance.
5. Select the desired searchable object from the table, and click the **Deactivate** icon. Alternatively, select **Deactivate** from the **Actions** menu or use the Ctrl+Shift+I keyboard shortcut.

The green check mark disappears from the **Active** column of the selected searchable object to indicate that the searchable object is not available for search.

Note: When the searchable object is inactive, a query of that object will return no results until the object is activated.

Task 2 Undeploy the Searchable Object from the Oracle SES Instance

Undeploy searchable objects from the Oracle SES instance to make the objects unavailable for the search engine instance to crawl. While undeployed, the searchable object is still associated with the search engine instance record and can be redeployed to the physical engine. You can undeploy only searchable objects that are inactive.

Note: When a searchable object is undeployed, any association to a search category or index schedule is removed.

To undeploy searchable objects from the Oracle SES instance using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
2. Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
3. Click the Oracle SES search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page.
4. From the table of search engine instances, select the record of the search engine instance that contains the searchable object you want to undeploy, and then select the **Searchable Objects** tab, shown in [Figure 7-5](#), to view the table of searchable objects for the selected search engine instance.
5. Select the desired searchable object from the table, and click the **Undeploy** icon. Alternatively, select **Undeploy** from the **Actions** menu or use the Ctrl+Shift+U keyboard shortcut.

The green check mark disappears from the **Deployed** column of the selected searchable object to indicate that the source has been deleted from Oracle SES.

7.4.10 Unregistering Searchable Objects from the Oracle Fusion Applications Database

Unregistering searchable objects removes the searchable object records from the Oracle Fusion Applications database. Using Fusion Applications Control, you can remove records of searchable objects from the Oracle Fusion Applications database.

Only those searchable objects that are inactive, undeployed, and not associated with any engine instance can be unregistered.

To remove searchable object records from the Oracle Fusion Applications database, follow the procedures in the following tasks:

- [Task 1, "Deactivate the Searchable Objects"](#)
- [Task 2, "Undeploy the Searchable Objects from the Oracle SES Instance"](#)
- [Task 3, "Disassociate the Searchable Objects from Search Engine Instances"](#)
- [Task 4, "Unregister the Searchable Objects"](#)

Task 1 Deactivate the Searchable Objects

Deactivate the searchable object you want to remove from the Oracle Fusion Applications database. For information, see [Task 1, "Deactivate the Searchable Objects"](#).

Task 2 Undeploy the Searchable Objects from the Oracle SES Instance

Undeploy the searchable object you want to remove from the Oracle Fusion Applications database. For information, see [Task 2, "Undeploy the Searchable Object from the Oracle SES Instance"](#).

Task 3 Disassociate the Searchable Objects from Search Engine Instances

Disassociating a searchable object from the search engine instance removes that object from the specified instance and makes it available for association to another search engine instance.

To delete searchable objects from search engine instances using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
2. Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
3. Click the Oracle SES search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page.
4. From the table of search engine instances, select the record of the search engine instance to which you want to add a searchable object, and then select the **Searchable Objects** tab, shown in [Figure 7-5](#), to view the table of searchable objects for the selected search engine instance.
5. Select the desired searchable object from the table, and click the **Delete Searchable Object** icon. Alternatively, select **Delete Record** from the **Actions** menu or use the Ctrl+Shift+D keyboard shortcut.

The selected searchable object is removed from the table for the selected search engine instance, but it is not deleted from the Oracle Fusion Applications database. It can be assigned to another search engine instance. Once assigned, you can deploy and activate the searchable object.

Note: To proceed with other operations, you must first select any record. Creating or deleting ECSF objects removes the focus from the window.

Task 4 Unregister the Searchable Objects

Unregistering a searchable object removes its record from the Oracle Fusion Applications database.

To unregister the searchable object using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
2. Select the application engine instance that contains the searchable objects you want to manage containing searchable objects to open the Enterprise Crawl and Search Framework Configuration Settings page.
3. From the Enterprise Crawl and Search Framework target menu, select **Unregister Searchable Object**.
4. On the Unregister Searchable Objects page, shown in [Figure 7–10](#), select the searchable object you want to unregister.

Figure 7–10 Unregister Searchable Objects Page



5. Click the **Unregister** icon.
A dialog with a confirmation message appears.

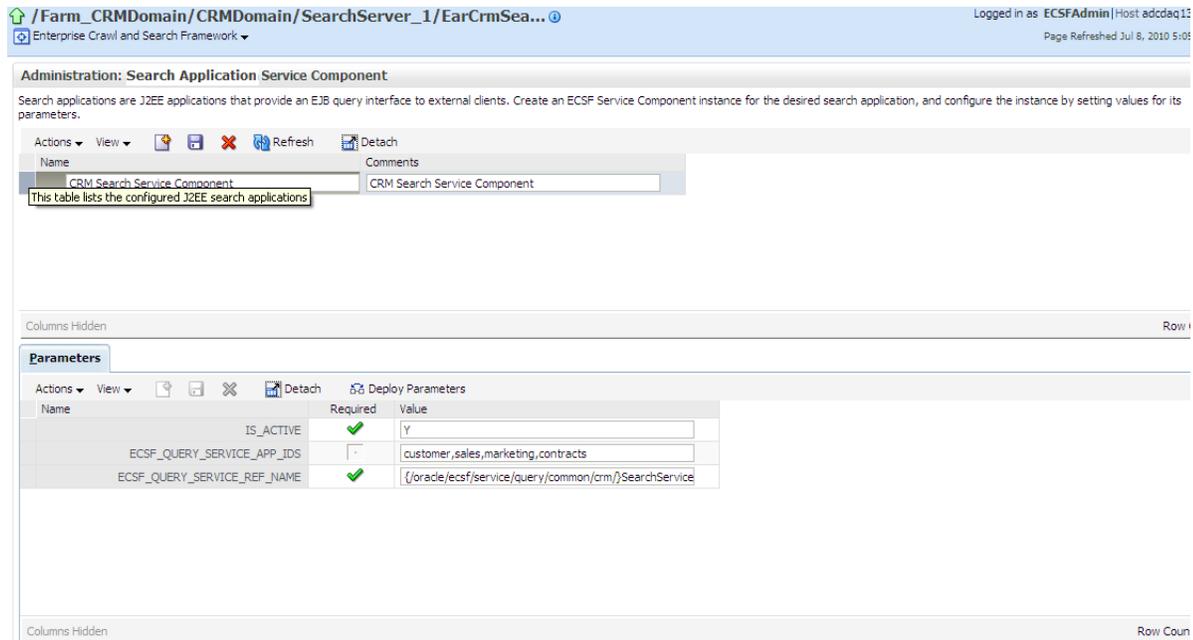
7.4.11 Modifying Search Application Service Component Instance Parameters

You can modify the Search Application Service Component instance parameters to edit the query web service reference name and activate or deactivate the search application.

To modify Search Application Service Component instance parameters using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
2. Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
3. Click the **Search Application Service Component** link.
4. On the Search Application Service Component administration page, shown in [Figure 7–11](#), select the desired Search Application Service Component instance record.

Figure 7–11 Search Application Service Component Administration Page



5. On the Parameters tab, modify the value for the desired parameter of the Search Application Service Component instance.
6. Save.

7.5 Performing Advanced Administration Tasks

You can extend search by performing any of the following tasks:

- [Modifying the Title, Body, Keyword, and Action Title of Searchable Objects](#)
- [Making New Searchable Objects Available for Search](#)
- [Defining Search Engine Instances](#)
- [Enabling Search on External Data Sources](#)
- [Adding Product Families for Federated Search](#)

Note: Administration of searchable objects, search categories, and index schedules must be performed using Fusion Applications Control. You can use the Oracle SES Administration GUI for all other Oracle SES related administration tasks. Oracle Enterprise Crawl and Search Framework only exposes the Oracle SES functionality for deploying and undeploying index schedules.

To access the ECSF pages in Fusion Applications Control, you must have Operator privileges in Oracle WebLogic Server. For information, see [Task 4, "Provide Access to ECSF Pages in Fusion Applications Control"](#).

7.5.1 Modifying the Title, Body, Keyword, and Action Title of Searchable Objects

You can customize searchable objects by changing the title, body, keywords, and search result action title using Fusion Applications Control.

Caution: Do not modify the Oracle Fusion Applications Help searchable object named `TopicSearchPVO`.

To change the title, body, keywords, and action title of searchable objects, follow the procedures in the following tasks:

- [Task 1, "Deactivate the Searchable Object"](#)
- [Task 2, "Undeploy the Searchable Object"](#)
- [Task 3, "Edit the Title, Body, Keywords, and Action Title"](#)
- [Task 4, "Enable the Customization"](#)
- [Task 5, "Deploy the Searchable Object"](#)
- [Task 6, "Activate the Searchable Object"](#)
- [Task 7, "Validate the Changes"](#)

Task 1 Deactivate the Searchable Object

If the searchable object that you want to customize is active, deactivate it. For information, see [Task 1, "Deactivate the Searchable Object"](#).

Task 2 Undeploy the Searchable Object

If the searchable object that you want to customize is deployed, undeploy it. For information, see [Task 2, "Undeploy the Searchable Object from the Oracle SES Instance"](#).

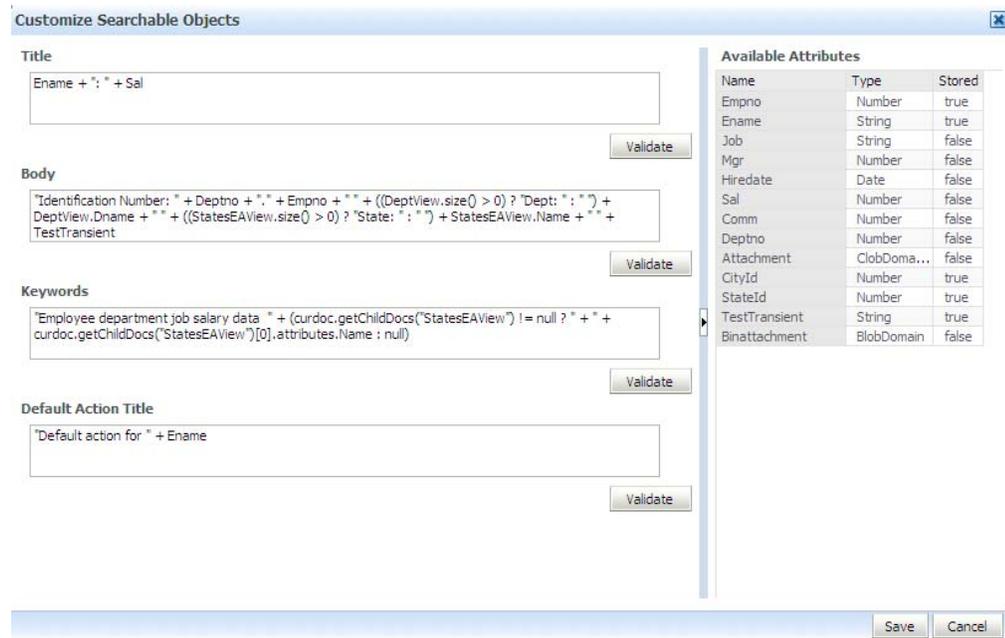
Task 3 Edit the Title, Body, Keywords, and Action Title

Edit the title, body, keywords, and search result action title to customize the searchable objects.

To edit the title, body, keywords, and action title of a searchable object using Fusion Applications Control:

1. From the **Actions** menu on the Searchable Objects tab, select **Modify**. Alternatively, use the Ctrl+Shift+Z keyboard shortcut.
2. In the Customize Searchable Objects dialog, shown in [Figure 7-12](#), edit the Groovy expression in the field for the property you want to modify.

Figure 7–12 Customize Searchable Objects Dialog



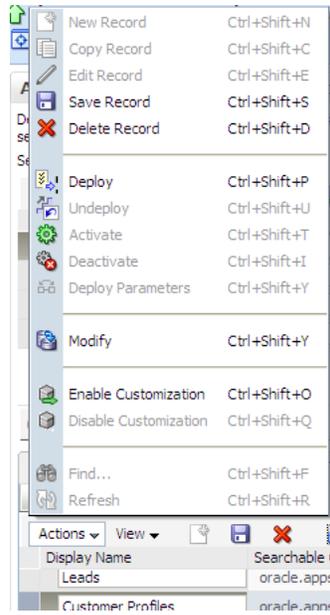
The Available Attributes table displays the attributes of the parent view object.

Note: You can use only stored attributes (for example, SRNumber) in the default action title. Using unstored attributes in the default action title results in an error during validation.

3. Click the **Validate** button to validate the Groovy expression.
4. Click the **Save** button to save your input to the database.

Task 4 Enable the Customization

Enable the customization by selecting **Enable Customization** from the **Actions** menu, or using the Ctrl+Shift+O keyboard shortcut.



This validates all the modified expressions for title, body, keywords, and action title and sets the customization flag. A green check mark appears in the Customized column to indicate that the searchable object is customized.

You can remove the customization flag by selecting **Disable Customization** from the **Actions** menu, or using the Ctrl+Shift+Q keyboard shortcut.

Task 5 Deploy the Searchable Object

Deploy the searchable object to the Oracle SES instance to make it crawlable. For information, see [Task 1, "Deploy the Searchable Objects to the Oracle Secure Enterprise Search \(Oracle SES\) Instance"](#).

Task 6 Activate the Searchable Object

Activate the searchable object so that a query of that object will return results. For information, see [Task 2, "Activate the Searchable Objects"](#).

Task 7 Validate the Changes

Validate that the Oracle Fusion Applications Search UI reflects the changes you made.

To validate your changes to the searchable objects title, body, keywords, and action title:

1. Log in to the search application.
2. Select the single category that contains the object.
3. Perform a query using the * wild card.
4. Check the search results for your changes to the title, body, keywords, and action title.

7.5.2 Making New Searchable Objects Available for Search

New searchable objects are objects in addition to the set of transactional objects that are available out of the box.

To make new searchable objects available for search, follow the procedures in the following tasks:

- Task 1, "Obtain the Necessary Information"
- Task 2, "Register the Searchable Objects"
- Task 3, "Associate the Searchable Objects with Search Engine Instances"
- Task 4, "Deploy the Searchable Objects to the Oracle SES Instance"
- Task 5, "Activate the Searchable Objects"
- Task 6, "Create the Search Categories"
- Task 7, "Undeploy the Search Categories"
- Task 8, "Associate the Searchable Objects with Search Categories"
- Task 9, "Deploy the Search Categories"
- Task 10, "Create the Index Schedules"
- Task 11, "Undeploy the Index Schedules"
- Task 12, "Associate the Searchable Objects with Index Schedules"
- Task 13, "Deploy the Index Schedules"
- Task 14, "Start the Index Schedules"

Task 1 Obtain the Necessary Information

In order to more efficiently make new searchable objects available for search, you should collect information from the Oracle Fusion Applications developer who creates the searchable objects, as well as from the business unit to which the new objects belong.

Obtain the following information before you begin the process of making new searchable objects available for search:

- Fully qualified object name of the searchable object (for example, `oracle.ecsf.demo.model.UX_CustomerVO`)
- Display name of the searchable object (for example, `Customer`)
- Name of the search category (existing or new) to which the searchable object must be associated (for example, `Customers`)
- Name of the index schedule (existing or new) to which the searchable object must be associated (for example, `CRM_Customers`)
- Name of the search engine instance with which the searchable object must be associated (for example, `CRMSES1`)
- Name of the search engine type to which the search engine instance belongs (for example, `SES11.2.1`)

Task 2 Register the Searchable Objects

Registering searchable objects creates searchable object records in the Oracle Fusion Applications database.

To create searchable object records in the Oracle Fusion Applications database using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.

2. Select the application engine instance that contains the searchable objects you want to manage containing searchable objects to open the Enterprise Crawl and Search Framework Configuration Settings page.
3. From the Enterprise Crawl and Search Framework target menu, select **Register Searchable Object**.
4. On the Register Searchable Object page, shown in [Figure 7-13](#), enter a display name for the searchable object in the **Display Name** field.

Figure 7-13 Register Searchable Object Page

The screenshot shows a web browser window with the address bar displaying the path: `/Farm_CRMDomain/CRMDomain/SearchServer_1/EarCrmSea...`. The page title is "Enterprise Crawl and Search Framework". The main content area is titled "Register Searchable Object:" and contains the following text: "Register searchable objects in the ECSF database by providing a searchable object name and a fully qualified view object as specified in your application." Below this text are two input fields: one labeled "* Display Name" and another labeled "* Searchable Object Name". At the top right of the form area, there are two buttons: "Register" and "Cancel".

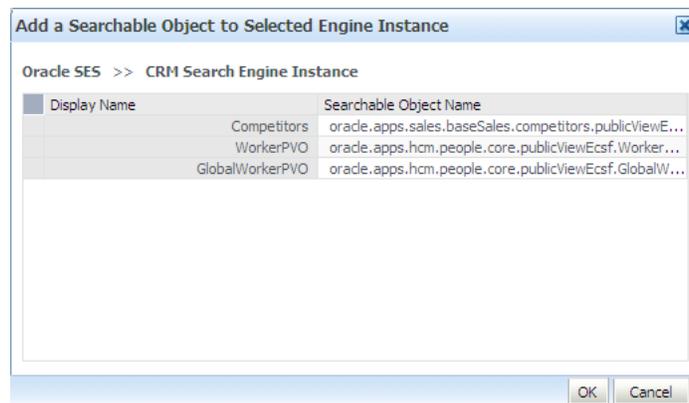
5. In the **Searchable Object Name** field, enter the fully qualified searchable object name that is defined in your application.
6. Click the **Register** button.

Task 3 Associate the Searchable Objects with Search Engine Instances

Associating a searchable object with a search engine instance adds that object to the specified instance. Only an unassigned object, which is a searchable object that is not associated with a search engine instance, can be added to a search engine instance, so a searchable object can only be associated with one search engine instance at a time.

To associate searchable objects with search engine instances using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
2. Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
3. Click the Oracle SES search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page.
4. From the table of search engine instances, select the record of the search engine instance to which you want to add a searchable object, and then select the **Searchable Objects** tab, shown in [Figure 7-5](#), to view the table of searchable objects for the selected search engine instance.
5. Click the **Add** button to open the Add a Searchable Object to Selected Engine Instance dialog, shown in [Figure 7-14](#).

Figure 7–14 Add a Searchable Object to Selected Engine Instance Dialog

6. Select the desired searchable object from the table, and click **OK**.

The searchable object you selected appears in the table of searchable objects on the Searchable Objects tab. The value in the **Application ID** column indicates the application with which the searchable object is associated.

Note: A searchable object can be assigned to only one search engine instance. If a searchable object has been assigned to a search engine instance, it will not be listed in the table of available searchable objects.

To proceed with other operations, you must first select any record. Creating or deleting ECSF objects removes the focus from the window.

Task 4 Deploy the Searchable Objects to the Oracle SES Instance

Deploy the searchable objects to the Oracle SES instance to make the objects available for the search engine instance to crawl. For information, see [Task 1, "Deploy the Searchable Objects to the Oracle Secure Enterprise Search \(Oracle SES\) Instance"](#).

Task 5 Activate the Searchable Objects

Activate the searchable objects so that a query of the deployed and crawled searchable objects will return results. For information, see [Task 2, "Activate the Searchable Objects"](#).

Task 6 Create the Search Categories

You can either create new search categories for the new searchable objects or add the new searchable objects to existing search categories. Creating a search category adds a new record to the `ECSF_SEARCH_INDEX_GROUP` table in the Oracle Fusion Applications database.

To create search categories using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
2. Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
3. Click the Oracle SES search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page.

4. From the table of search engine instances, select the record of the search engine instance for which you want to add a search category, and then select the **Search Categories** tab, shown in [Figure 7-15](#), to view the table of search categories for the selected search engine instance.
5. Click the **New Search Category** icon. Alternatively, select **New Record** from the **Actions** menu or use the Ctrl+Shift+N keyboard shortcut.
A new record named CHANGE_ME is added to the table of search categories.
6. Enter a new value in the **Name** column.
7. Select a value (LOCAL or GLOBAL) from the **Scope** column dropdown list.
8. Click the **Save Search Category** icon to save the selected record. Alternatively, select **Save Record** from the **Actions** menu or use the Ctrl+Shift+S keyboard shortcut.

Caution: A save will occur only for the selected record. You can only save one record at a time.

Note: To proceed with other operations, you must first select any record. Creating or deleting ECSF objects removes the focus from the window.

Task 7 Undeploy the Search Categories

To add the new searchable objects to existing search categories, you must first undeploy the search categories. For more information, see [Task 1, "Undeploy the Search Categories"](#).

Task 8 Associate the Searchable Objects with Search Categories

Associating searchable objects with a search category adds the objects to that search category. You can only associate deployed searchable objects with search categories. You can associate the same searchable object with multiple search categories. You cannot associate searchable objects with search categories that have been deployed. If the search category has been deployed, you must undeploy it before you can associate searchable objects with it. For information, see [Task 1, "Undeploy the Search Categories"](#).

To associate searchable objects with search categories using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
2. Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
3. Click the Oracle SES search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page.
4. From the table of search engine instances, select the record of the desired search engine instance, and then select the **Search Categories** tab, shown in [Figure 7-15](#), to view the table of search categories for the selected search engine instance.

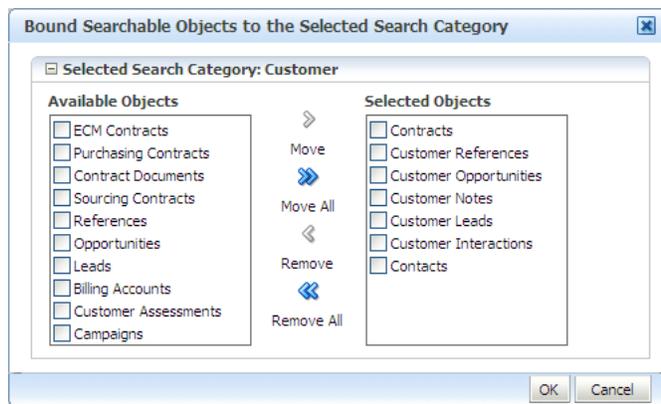
Figure 7–15 Search Categories Tab

Name	Deployed	Deployed Since	Application ID	Scope
Customer	-		customer	LOCAL
Customers	✓		customer	GLOBAL
Lead	✓		marketing	GLOBAL
Leads	✓		marketing	LOCAL
Opportunity	✓		sales	GLOBAL
Opportunities	✓		sales	LOCAL
References	✓		sales	GLOBAL

This table lists the Search Groups for the selected Instance

5. Select the desired search category to which you want to add searchable objects, and click the **Bind** icon.
The **Bind** icon is disabled if the search category is deployed. You cannot associate or disassociate searchable objects if the search category is deployed.
6. In the Bound Searchable Objects to the Selected Search Category dialog, shown in [Figure 7–16](#), select the checkboxes of the desired searchable objects from the **Available Objects** list, and click the **Move** icon to add them to the **Selected Objects** list. Only deployed searchable objects are listed as available objects.

Figure 7–16 Bound Searchable Objects to the Selected Search Category Dialog



7. Click **OK**.
A new record is created in the ECSF_SIO_SIG_INT intersection table.

Task 9 Deploy the Search Categories

Deploy the search category. For information, see [Task 4, "Deploy the Search Categories"](#).

Task 10 Create the Index Schedules

You can either create new index schedules for the new searchable objects or add the new searchable objects to existing index schedules. For information on how to create index schedules, see [Section 7.4.2.1](#).

Task 11 Undeploy the Index Schedules

To add the new searchable objects to existing index schedules, you must first undeploy the index schedules. For more information, see [Section 7.4.2.8](#).

Task 12 Associate the Searchable Objects with Index Schedules

Associate the new searchable object to either the newly created index schedule or an existing, undeployed index schedule. For information, see [Section 7.4.2.2](#).

Task 13 Deploy the Index Schedules

Deploy the index schedules. For information, see [Task 4, "Deploy the Index Schedules"](#).

Task 14 Start the Index Schedules

Start the index schedules. For information, see [Task 5, "Start the Index Schedules"](#).

7.5.3 Defining Search Engine Instances

ECSF supports multiple search engine instances for each search engine type. Each search engine instance belongs to exactly one version of the Oracle SES search engine type and is stored in the `ECSF_ENGINE_INSTANCE` table.

You can perform the following tasks to define search engine instances:

- [Creating Search Engine Instances](#)
- [Configuring Search Engine Instance Parameters](#)
- [Modifying Search Engine Instances](#)
- [Modifying Search Engine Instance Parameters](#)

7.5.3.1 Creating Search Engine Instances

Creating a search engine instance adds a new search engine instance to the specified search engine type.

To create search engine instances using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
2. Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
3. Click the Oracle SES search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page, shown in [Figure 7-4](#).
4. In the Search Engine Instance toolbar, click the **New Engine Instance** icon. Alternatively, select **New Record** from the **Actions** menu or use the Ctrl+N keyboard shortcut.

A new record is added to the top of the table of search engine instances.

5. Enter values for the **Name** and **Comments** columns.

Caution: A save will occur only for the selected record. You can only save one record at a time.

6. Click the **Save Engine Instance** icon to save the selected record. Alternatively, select **Save Record** from the **Actions** menu or use the Ctrl+S keyboard shortcut.
7. Configure the search engine instance parameters. For information, see [Section 7.5.3.2](#).
8. Click the **Register** icon to register the identity plug-in and create the federated trusted entity for the selected search engine instance. Alternatively, select **Register** from the **Actions** menu or use the Ctrl+G keyboard shortcut.

A dialog with a confirmation message appears.

Registering the identity plug-in for search engine instances allows users to perform secure searches.

7.5.3.2 Configuring Search Engine Instance Parameters

Configuring search engine instance parameters sets parameter values to the search engine instance.

To configure search engine instance parameters using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
2. Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
3. Click the Oracle SES search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page, shown in [Figure 7-4](#).
4. From the table of search engine instances, select the record of the search engine instance for which you want to configure the parameters.

The parameters of the selected search engine instance appears in the Parameters table, shown in [Table 7-3](#).

Table 7-3 Search Engine Instance Parameters

Parameter Name	Description	Sample Value
SES_ADMIN_SERVICE	URL to Oracle SES administrative web service endpoint.	http://example.com:7777/search/api/admin/AdminService
SES_ADMIN_USERNAME	Oracle SES user for logging into Oracle SES administrative web service.	searchsys
SES_ADMIN_PASSWORD	Password for user SES_ADMIN_PROXY_USERNAME.	welcome1
SES_ADMIN_SESSION_TIMEOUT	Time, in minutes, allowed by ECSF to keep administration session to Oracle Secure Enterprise Search (Oracle SES) before forcing to re-login/reestablish session with Oracle SES. This timeout is independent of Oracle SES's internal administration session timeout.	60

Table 7-3 (Cont.) Search Engine Instance Parameters

Parameter Name	Description	Sample Value
SES_QUERY_SERVICE	URL to Oracle SES Query web service endpoint.	http://example.com:7777/search/query/OracleSearch
SES_QUERY_PROXY_USERNAME	ECSF/Fusion user used by Oracle SES as a proxy user for queries.	scott
SES_QUERY_PROXY_PASSWORD	Password for ECSF user used by Oracle SES as a proxy user.	tiger
SES_QUERY_SESSION_TIMEOUT	Time, in minutes, allowed by ECSF to keep query session to Oracle SES before forcing to re-login/reestablish session with Oracle SES. This timeout is independent of Oracle SES's internal query session timeout.	60
ECSF_DATA_SERVICE	URL to ECSF servlets. This URL is used to construct configuration feeds for Oracle SES.	http://example.com:7101/ampprot/searchfeedservlet
ECSF_SECURITY_USERNAME	ECSF/Fusion user used by Oracle SES to invoke ECSF Security Service.	scott or fusion
ECSF_SECURITY_PASSWORD	Password for ECSF/Fusion user used by Oracle SES to invoke ECSF Security Service.	tiger or fusion
ECSF_SECURITY_SERVICE	URL to ECSF servlets. This URL is used to construct configuration feeds for Oracle SES. ECSF runtime URL and ECSF Security Service URL can be set to 2 different servers, which is also what Oracle SES supports.	http://example.com:7101/ampprot/searchfeedservlet
ECSF_REDIRECT_SERVICE	URL to ECSF feed servlets. This URL is used to redirect drilldowns/actionable results.	http://example.com:7101/ampprot/searchfeedservlet

5. Select the record for the desired parameter, and modify the entry in the corresponding **Value** column.

Caution: A save will occur only for the selected record. You can only save one record at a time.

6. Click the **Save Engine Instance Parameter** icon to save the selected record. Alternatively, select **Save Record** from the **Actions** menu or use the Ctrl+Shift+S keyboard shortcut.

7.5.3.3 Modifying Search Engine Instances

Modifying a search engine instance edits the properties of an existing search engine instance.

To modify search engine instances using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
2. Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
3. Click the Oracle SES search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page, shown in [Figure 7-4](#).
4. Modify the entries in the **Name** and **Comments** columns.

Caution: A save will occur only for the selected record. You can only save one record at a time.

5. Click the **Save Engine Instance** icon to save the selected record. Alternatively, select **Save Record** from the **Actions** menu or use the Ctrl+S keyboard shortcut.

7.5.3.4 Modifying Search Engine Instance Parameters

Modifying search engine instance parameters updates the parameter values. Once the parameter values are updated, you must update all the deployed searchable objects of the same search engine instance with the latest search engine instance parameters on the Oracle SES server. Only searchable objects that have been deployed can be updated with the latest search engine instance parameters on the Oracle SES server.

To modify search engine instance parameters, follow the procedures in the following tasks:

- [Task 1, "Edit the Search Engine Instance Parameters"](#)
- [Task 2, "Deploy the Parameters to All Deployed Searchable Objects"](#)
- [Task 3, "Deploy the Parameters to Individual Deployed Searchable Objects"](#)

Task 1 Edit the Search Engine Instance Parameters

You can edit the search engine instance parameters to change their values.

To change the values of the search engine instance parameters using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
2. Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
3. Click the Oracle SES search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page, shown in [Figure 7-4](#).
4. From the table of search engine instances, select the record of the search engine instance for which you want to modify the parameters.

The parameters of the selected search engine instance appears in the Parameters tab.

5. Select the record for the desired parameter, and modify the entry in the corresponding Value column.

Tip: It is recommended that you modify these parameters during system downtime. It is also recommended that you update the objects when all the parameters are set.

Task 2 Deploy the Parameters to All Deployed Searchable Objects

After you update the search engine instance parameter values, you must update all the deployed searchable objects of the same search engine instance with the latest search engine instance parameters on the Oracle SES server.

To deploy the new values of the search engine instance parameters to all the deployed searchable objects in the search engine instance using Fusion Applications Control:

1. From the Parameters tab, click the **Deploy Parameters** icon to update all the deployed searchable objects with the latest search engine instance parameter values. Alternatively, select **Deploy Parameters** from the **Actions** menu or use the Ctrl+Shift+Y keyboard shortcut.

A confirmation dialog appears.

2. Validate that all deployed searchable objects have been updated with the latest parameter values by selecting the **Searchable Objects** tab and viewing the Parameters Applied column.

A green check mark indicates that the searchable object is updated with the latest parameter values. A red X indicates that the searchable object is not updated with the latest parameter values. A gray box indicates that the searchable object is not deployed.

Task 3 Deploy the Parameters to Individual Deployed Searchable Objects

For any deployed searchable objects that were not updated with the latest parameter values, update them individually.

To deploy the parameters to individual deployed searchable objects using Fusion Applications Control:

1. From the Searchable Objects tab, select the record corresponding to a searchable object that was not updated with the new parameter value.

A green check mark indicates that the searchable object is updated with the latest parameter values. A red X indicates that the searchable object is not updated with the latest parameter values. A gray box indicates that the searchable object is not deployed.

2. Click the **Deploy Parameters** icon to update the parameter values for that deployed searchable object. Alternatively, select **Deploy Parameters** from the **Actions** menu or use the Ctrl+Shift+Y keyboard shortcut.

A confirmation dialog appears.

A green check mark will replace the red X when the latest search engine instance parameters are updated for the searchable object.

7.5.3.5 Deleting Search Engine Instances

Deleting search engine instances removes them from the specified search engine type.

To delete search engine instances using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
2. Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
3. Click the Oracle SES search engine type name link in the Search Engine Types table to open the Search Engine Instance administration page, shown in [Figure 7-4](#).
4. From the table of search engine instances, select the record you want to delete.
5. Click the **Delete Engine Instance** icon to delete the selected record. Alternatively, select **Delete Record** from the **Actions** menu or use the Ctrl+D keyboard shortcut.

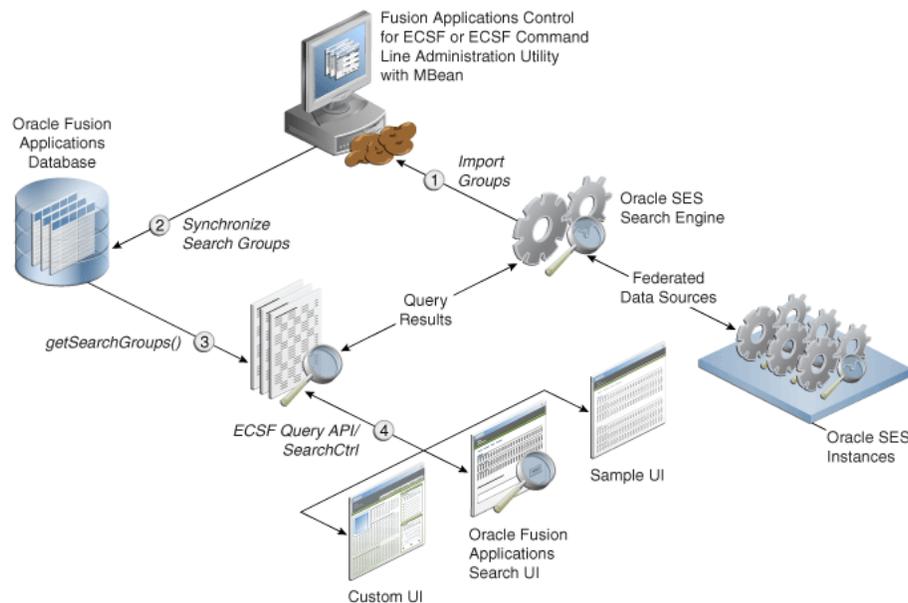
7.5.4 Enabling Search on External Data Sources

In addition to providing search against ECSF data sources (that is, view object based data sources deployed from ECSF administration) and data sources that are crawled via ECSF, ECSF supports search against all registered Oracle SES data source groups that contain external data sources. External data sources include data sources such as Intranet, mail, database, and federated data sources (data sources crawled in separate engine instances). Oracle SES data source groups can also contain federated ECSF data sources (data sources created in different logical engine instances via the ECSF administration tools).

ECSF Administration synchronizes categories with the data source groups on Oracle SES. The data sources are not synchronized.

[Figure 7-17](#) illustrates the process flow for enabling search on external data sources.

Figure 7-17 Process Flow for Search on External Data Sources



The process flow includes the following steps:

1. Data source groups are imported from Oracle SES via the Fusion Applications Control. For information, see [Task 1, "Import the External Search Categories"](#).

2. The Oracle Fusion Applications database is synchronized with the data source groups from Oracle SES.

ECSF sets the `IS_EXTERNAL` flag to `true` on records in the `ECSF_SEARCH_INDEX_GROUP` table to identify them as external search categories.

3. ECSF-created search categories and external search categories are both accessible via `getSearchGroups()`.
4. `SearchCtrl` can query any search category, and the underlying data may actually be federated to other Oracle SES instances.

Note: ECSF supports multiple search engine instances for each application. However, each query can only contain search categories deployed on the same search engine instance.

Search categories containing ECSF data sources are represented in the `ECSF_SEARCH_INDEX_GROUP` table and linked to one or more searchable objects through an intersection table. Search results are usually grouped by the searchable object name stored in the `ECSF_SO_NAME` custom attribute that is returned in each search result.

Note: ECSF currently does not support querying search categories that contain mixed data sources. Each search category must contain either only ECSF data sources or only external data sources.

External search categories are not linked to searchable objects. Instead, ECSF creates the `ExternalSearchableObject` searchable object type for external search categories. Search results from external search categories are assigned to the external searchable object, named `oracle.ecsf.meta.ExternalSearchableObject`, that is created during runtime.

[Example 7-1](#) queries an external category called "Oracle Website."

Example 7-1 Sample Code for Querying External Categories

```
public void federationTest()
{
    SearchCtrl searchCtrl = new SearchCtrl();
    SearchHits searchHits = null;
    SearchContext searchContext = null;
    SearchEngineInstance engineInstance = null;
    SearchGroup[] sgs = new SearchGroup[1];

    QueryMetaDataImpl queryMetaData = new QueryMetaDataImpl();
    queryMetaData.setQueryString("");
    queryMetaData.setPageSize(10);
    queryMetaData.setCurrentPage(1);

    ArrayList<SearchEngineInstance> engineInstances =
        (ArrayList<SearchEngineInstance>) searchCtrl.getEngineInstances();
    for (int i = 0; i < engineInstances.size(); i++)
    {
        if (engineInstances.get(i).getId() == 1)
        {
            engineInstance = engineInstances.get(i);
            break;
        }
    }
}
```

```
    }

    ArrayList<SearchGroup> groups =
(ArrayList<SearchGroup>)engineInstance.getSearchGroups();
    for (int j = 0; j < groups.size(); j++)
    {
        if (groups.get(j).getName().equals("Oracle Website"))
        {
            sgs[0] = groups.get(j);
            break;
        }
    }

    queryMetaData.setSearchGroups(sgs);
    searchContext = ContextFactory.getSearchContext();
    searchContext.bindUser("scott");
    try
    {
        searchHits = searchCtrl.runQuery(searchContext, queryMetaData);
    }
    catch (Exception e)
    {
        fail("Failed runQuery due to "+e.getMessage());
    }

    IndexedDocument doc = searchHits.getDocument(0);
    String url = null;
    try
    {
        url = doc.getDefaultAction().getRedirectURL();
        ActionURLResolver resolver = new ActionURLResolver();
        url = resolver.resolveURL(searchContext, doc.getDefaultAction());
    }
    catch (Exception e)
    {
    }
}
```

Since there is no view object or search object definition file where actions can be defined, the URL attribute of the resulting object returned from Oracle SES will be used to generate a default action with the name Default.

Advanced search and faceted navigation are not supported for external data sources.

7.5.5 Adding Product Families for Federated Search

To include a new product family for federated search, you must add a Search Application Service Component instance to engine type Oracle Search Application and configure the instance.

Note: Manage credentials via the MBean using JConsole. For information, see *Oracle Fusion Middleware Fusion Developer's Guide for Oracle Application Development Framework*.

Each product family (Oracle Fusion Customer Relationship Management, Oracle Fusion Human Capital Management, and Oracle Fusion Supply Chain Management) has its own search application. Each search application provides the data service, query web service, authentication service, and authorization service for the product

family. Each search application also contains all the searchable objects and its dependencies for the product family as well as the references to the ECSF shared libraries.

To create and configure a Search Application Service Component instance using Fusion Applications Control:

1. From the navigation pane, expand the farm and then the **Enterprise Crawl and Search Framework** folder.
2. Select the application engine instance that contains the searchable objects you want to manage to open the Enterprise Crawl and Search Framework Configuration Settings page.
3. Click the **Search Application Service Component** link.
4. On the Search Application Service Component administration page, shown in [Figure 7-11](#), click the **New Engine Instance** icon. Alternatively, select **New Record** from the **Actions** menu or use the Ctrl+Shift+N keyboard shortcut.

A new Search Application Service Component instance record is added to the top of the table of configured search applications.

5. Enter a value for the **Name** column.
6. On the Parameters tab, set the value for the `ECSF_QUERY_SERVICE_REF_NAME` parameter of the Search Application Service Component instance. Use the following format:

```
{/oracle/ecsfc/service/query/common/componentname/}SearchService
```

replacing *componentname* with the name of the Search Application Service Component.

7. Set the value for the `IS_ACTIVE` parameter of the Search Application Service Component instance.

Set the value to `Y` to specify that the Search Application Service Component is active. The component is skipped during global search related functions when you set the value to `N`. If the `IS_ACTIVE` parameter value is not defined for the component in the database, the component is treated as active.

8. Save.

Managing Oracle WebLogic Communication Services for Click-to-Dial Functionality

Oracle WebLogic Communication Services provides click-to-dial functionality for applications primarily through contextual actions. Contextual actions provide related information and actions to users within the immediate context of the object instances upon which they act. This chapter explains how to enable and configure these services.

It includes the following topics:

- [Introduction to Third Party Call Web Service](#)
- [Modifying Third Party Call Web Service Default Settings Using the System MBean Browser](#)
- [Call Routing Through a Session Border Controller](#)

8.1 Introduction to Third Party Call Web Service

The Third Party Call Web service in Oracle WebLogic Communication Services enables click-to-dial functionality in applications by providing an interface to initiate a call between two parties. The Web service is an implementation of the Parlay X 2.1 Third Party Call Interface.

The Third Party Call Web service can:

- Set up a call between two parties
- Query the status of an ongoing call
- Cancel a call it is creating
- Terminate an ongoing call it created

The Third Party Call Web service is provided through the `thirdpartycall` application, which is installed with default configuration settings in the domain for the each product family.

8.2 Modifying Third Party Call Web Service Default Settings Using the System MBean Browser

You can use the System MBean Browser to modify the default configuration settings for the `thirdpartycall` application for the product family domains.

To modify the MBeans for the `thirdpartycall` application with the System MBean Browser from the Fusion Applications Control:

1. From the target navigation pane, expand the farm and then **Application Deployments**.
2. Select the application `thirdpartycall`.
3. From the **Application Deployment** menu, select **System MBean Browser**.
Fusion Applications Control displays the **System MBean Browser** page.
4. Scroll down to the node that displays `oracle.ucs.thirdpartycall`. Expand the `oracle.ucs.thirdpartycall` node and drill down to the node that displays the `ThirdPartyCallConfiguration` MBean.
5. Click the `ThirdPartyCallConfiguration` MBean node. The configuration settings for the `thirdpartycall` application appear on the right pane. Select the **Attributes** tab to display the details for this group of attributes.

For more information about the configurable attributes for the `thirdpartycall` application, see [Table 8-1](#).
6. Update values of the attributes to modify by entering the new value in the **Value** field.
7. Click **Apply**.

For general information about the System MBean Browser, see the "Getting Started Using Oracle Enterprise Manager Fusion Middleware Control" section in *Oracle Fusion Middleware Administrator's Guide*.

[Table 8-1](#) contains a list of attributes for configuring the Third Party Call Web service

Table 8-1 Configuration Attributes for Third Party Call Web Service

Attribute	Format	Description	Values
<code>ChargingAllowed</code>	String	Specifies if charging is allowed.	Default: False Note: No server restart is required if you change the value.
<code>ControllerUri</code>	String in URI format	Specifies the Third Party controller Session Initiation Protocol (SIP) URI that is used to establish the third party call. If this value is set, a call appears to the callee to come from this URI.	Default: empty By default, no controller URI is used to establish the call. In this case, the call appears to the callee to come from the caller. Note: You must restart the server if you change this value.
<code>MaximumCallDuration</code>	int	Specifies the maximum time (in minutes) for an ongoing call. If this time expires, the call is terminated.	Default: 1440 (24 hours) Minimum: -1 (0 and -1 means that call-length supervision is disabled) Maximum: Unbounded Note: No server restart is required if you decrease this value. You must restart the server if you increase the value or set the value to 0 or -1.

Table 8–1 (Cont.) Configuration Attributes for Third Party Call Web Service

Attribute	Format	Description	Values
PAssertedIdentityUri	String in URI format	Specifies the SIP URI used in the P-Asserted-Identity header added by the Third Party Call service. If left blank the P-Asserted-Identity header is not added.	Default: empty Note: You must restart the server if you change this value.
StatusRetentionTime	int	Specifies the time (in minutes) for retaining the status information about a call after it was terminated.	Default: 5 Minimum: -1 (0 and -1 means that the status information is removed when the call is terminated). Maximum: Unbounded Note: No server restart is required if you change the value.
PreloadedRouteUri	String in URI format	Specifies the SIP URI to put in a route header of outgoing SIP INVITE requests during a call setup.	Default: empty Note: You must restart the server if you change this value.

8.3 Call Routing Through a Session Border Controller

The Third Party Call Web service interfaces on the backend with SIP infrastructure, such as a SIP-enabled switch or a SIP-to-PSTN gateway. When the SIP infrastructure lies outside the enterprise network, such as in the case of a hosted SIP trunking service or if the switch is in an external network, a Session Border Controller (SBC) is typically used to secure and control the interaction. A SBC is a piece of equipment that is deployed at the border of the network and controls the signaling and media across network address translation (NAT) devices and firewall boundaries. In addition, it prevents unauthorized access and denial of service (DoS) attacks.

You can route the Third Party Call Control (TPCC) requests through an SBC either by setting up an outbound proxy server or by configuring one TPCC attribute.

This section contains the following topics:

- [Setting Up an Outbound Proxy Server](#)
- [Configuring the PreLoadedRouteUri Attribute](#)

8.3.1 Setting Up an Outbound Proxy Server

You can set up the SBC as an outbound proxy server and route TPCC requests through the outbound proxy server. If you define an outbound proxy server (SBC) for a SIP Servlet container, any request from all SIP servlet applications (including the `thirdpartycall` application) on this SIP Servlet container is routed through the SBC.

Oracle WebLogic SIP Server provides Oracle WebLogic Administration Console extensions that allow you to modify the SIP Servlet container properties. To set up an outbound proxy:

1. Enter the following URL in a browser:

```
http://hostname:port_number/console
```

where the *hostname* and the *port_number* correspond to the Administration Server for the domain.

2. Log in using the Oracle Fusion Middleware administrative user name and password.
3. Select the **SipServer** node in the left pane. The right pane of the console displays the tabbed pages that are used for configuring and monitoring WebLogic SIP Server.
4. Select the **Proxy** tab.
5. Select the **Proxy Routing Policy** as proxy.
6. In the **Proxy Routing URIs** section, enter the URI as `sip:SBC_hostname:SBC_portnumber`, where `SBC_hostname` and `SBC_portnumber` represent the host and port of the Session Border Controller, respectively.
7. Click **Save**.

For more information, see the "SIP Server: Configuration: Proxy" topic in the Oracle WebLogic Administration Console Help.

Note: Oracle recommends that you use the Administration Console to modify `sipserver.xml` indirectly, rather than editing the file by hand. Using the Administration Console ensures that the `sipserver.xml` document always contains valid XML.

8.3.2 Configuring the PreLoadedRouteUri Attribute

To route requests only from the `thirdpartycall` application through SBC, set the value of the `PreLoadedRouteUri` attribute as `sip:SBC_hostname:SBC_portnumber`, where `SBC_hostname` and `SBC_portnumber` represent the host and port of the SBC, respectively.

For more information about modifying the `PreLoadedRouteUri` attribute, see [Section 8.2](#).

Part III

Monitoring

This part provides information about how to monitor the Oracle Fusion Applications environment. It explains how to find information about the cause of an error and its corrective action, and how to view and manage log files to assist in monitoring system activity and to diagnose problems.

Part III contains the following chapters:

- [Chapter 9, "Monitoring Oracle Fusion Applications"](#)
- [Chapter 10, "Monitoring the Oracle Fusion Applications Middle Tier"](#)
- [Chapter 11, "Diagnosing the Oracle Fusion Applications Middle Tier Performance"](#)
- [Chapter 12, "Monitoring and Tuning Oracle Database for Oracle Fusion Applications"](#)
- [Chapter 13, "Managing Oracle Fusion Applications Log Files and Diagnostic Tests"](#)

Monitoring Oracle Fusion Applications

This chapter describes how to manage Oracle Fusion web applications using Oracle Enterprise Manager Fusion Applications Control (Fusion Applications Control) and Oracle Enterprise Manager (Cloud Control).

This chapter includes the following topics:

- [Introduction to Monitoring Oracle Fusion Applications](#)
- [The Product Family Home Page and Sections in Fusion Applications Control](#)
- [The Product Home Page and Sections in Fusion Applications Control](#)
- [The Fusion Instance Overview Page and Sections in Cloud Control](#)
- [The Product Family Home Page and Sections in Cloud Control](#)
- [The Product Home Page and Sections in Cloud Control](#)
- [Monitoring Product Families and Products](#)
- [Monitoring the Detailed Status of Oracle Fusion Applications](#)
- [Monitoring Business Performance Metrics for a Product Family](#)
- [Monitoring Business Performance Metrics for a Product](#)
- [Viewing the Routing Topology of an Oracle Fusion Applications Instance, Product Family, or Product](#)
- [Viewing the Configuration Topology of an Oracle Fusion Applications Instance, Product Family, or Product](#)

9.1 Introduction to Monitoring Oracle Fusion Applications

Oracle Enterprise Manager Fusion Applications Control (Fusion Applications Control) is a tool used to monitor Oracle Fusion Applications product families and products deployed to Oracle WebLogic Server. Oracle Enterprise Manager Cloud Control (Cloud Control) is a tool used in much the same way, but viewing Oracle Fusion Applications product families and products within the larger context of Oracle Fusion Middleware.

Oracle Enterprise Manager Fusion Applications Control provides a large picture view of your deployments, while also allowing you to drill down for more detailed information. You can monitor the status of the servers to which product families are deployed and the databases that store product information, the status of all scheduled job requests running in the system, as well as the overall state and performance of product families and products. The customizable performance summary pages allow you to monitor product performance and diagnose problems. You can modify these

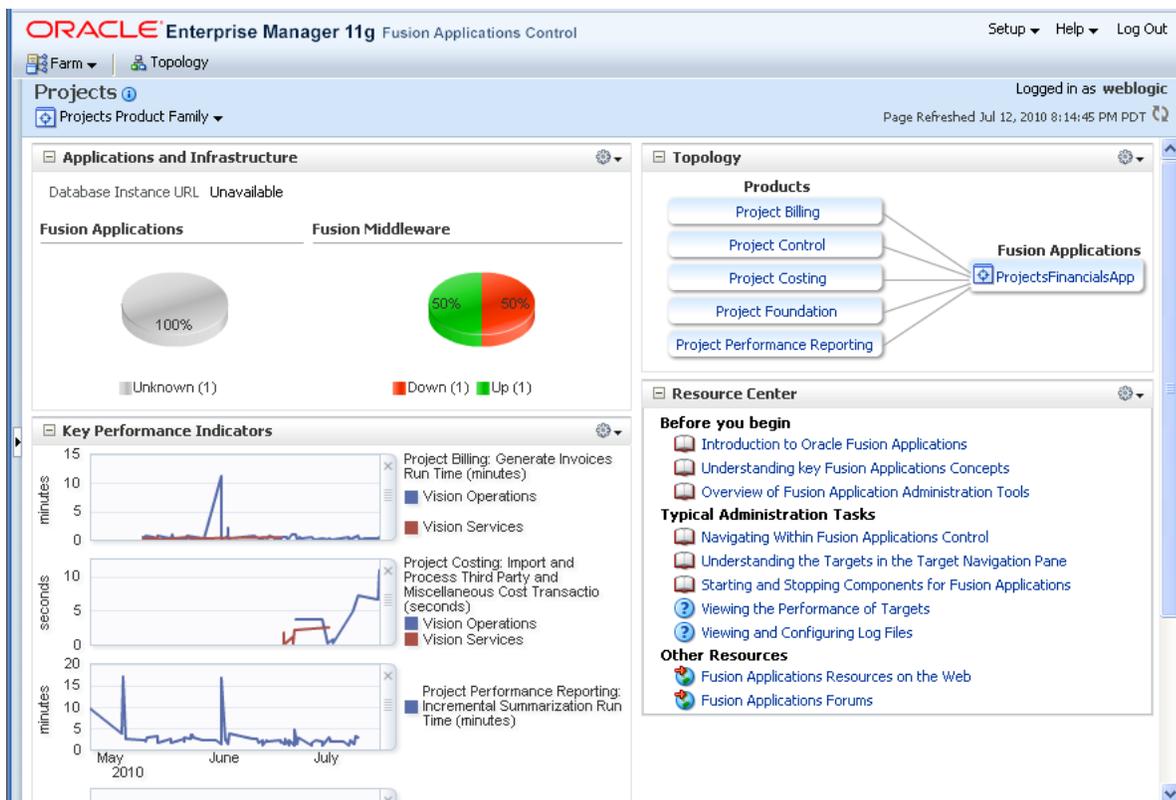
charts to display statistics that are relevant to your deployment, or add an additional product to the chart so as to compare the performance information for two components in one chart.

Using a topographical chart, you can examine the relationship between a deployed product family and its products, and drill down to view the status of individual product instances. Out-of-the-box performance metrics provided for each Oracle Fusion Applications product allow detailed monitoring of the state and performance of all product families and each individual product.

9.2 The Product Family Home Page and Sections in Fusion Applications Control

The product family home page provides an overview of the products belonging to a product family. An example of a product family home page is shown in [Figure 9-1](#).

Figure 9-1 The Product Family Home Page in Fusion Applications Control



The Product Family home page displays the following sections:

- [Applications and Infrastructure Section](#)
- [Key Performance Indicators Section](#)
- [Topology Section](#)
- [Oracle Enterprise Scheduler Sections](#)

9.2.1 Applications and Infrastructure Section

The Applications and Infrastructure section displays the status of the database, deployed Oracle Fusion Applications and Oracle Fusion Middleware itself. The Fusion Applications and Fusion Middleware pie charts display the percentage of Oracle Fusion Applications and Oracle Fusion Middleware servers that are currently up or down. The number of applications and servers that are running or down is shown below each pie chart.

9.2.2 Key Performance Indicators Section

The Key Performance Indicators section displays the main performance metrics related to the product family or product. By default, these metrics display in graphical format. Alternatively, you can display the performance metrics in tabular format instead. Mousing over a particular metric causes the performance graph to display only the data for that metric.

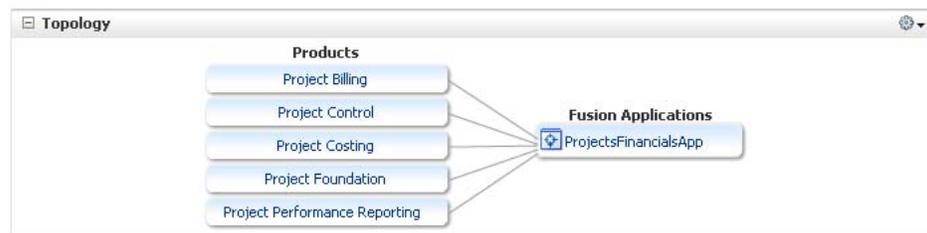
You can add or remove metrics from the Key Performance Indicators section. You can also change the order in which metrics display. For more information about monitoring business metrics for a product family or product, see [Section 9.10](#)

9.2.3 Topology Section

The Topology section displays the components of the product family or product. In the Product home page, the Topology section displays the same information, with the irrelevant products and Java EE applications grayed out.

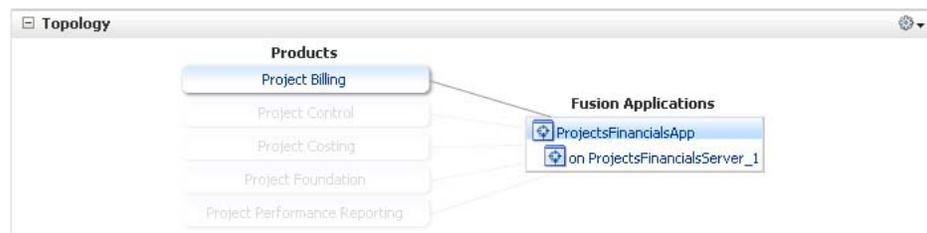
The Product Family home page Topology section is shown in [Figure 9-2](#).

Figure 9-2 Product Family Home Page Topology Region



The Product home page Topology section is shown in [Figure 9-3](#).

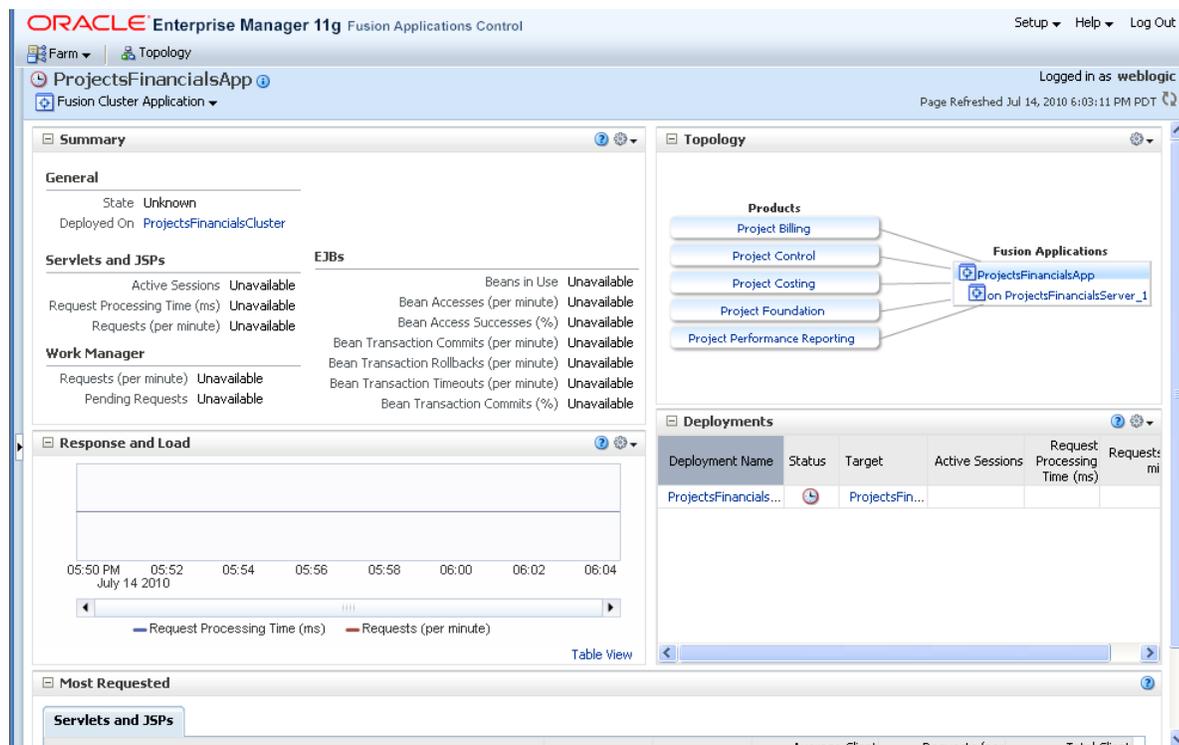
Figure 9-3 Product Home Page Topology Section



Mousing over a component under the Oracle Fusion Applications heading highlights the related product under the Product heading, and vice versa. In the Product Family home page, clicking a product in the Topology section displays the Product home page for that product. Clicking the name of the Oracle Fusion Applications product family in the Topology section displays the Oracle Fusion application instances running on

servers in the application cluster. The Fusion Cluster Application home page displays, as shown in [Figure 9-4](#).

Figure 9-4 Fusion Cluster Application Home Page



For more information about viewing topology information, see [Section 10.2.3.1](#)

9.2.4 Oracle Enterprise Scheduler Sections

The Oracle Enterprise Scheduler sections display information pertaining to scheduled jobs running in relation to the product or product family. The Oracle Enterprise Scheduler sections are as follows.

- **Top Long Running Job Requests:** This section displays information about the top ten long running scheduled job requests, including request ID, job definition used, the product and time elapsed since the job executed.
- **Recently Completed Job Requests (Last 24 Hours):** This section displays information about the most recent scheduled jobs completed within the last 24 hours, including job definition used, the product, the number of completed and failed requests, the percentage of successfully completed job requests and the time elapsed from execution.
- **Ready Job Requests:** This section displays the scheduled job requests that are awaiting execution, including the job definition used, the product and the number of jobs that are ready to run.

For more information about this section, see the online help.

9.2.5 SOA Composite Section

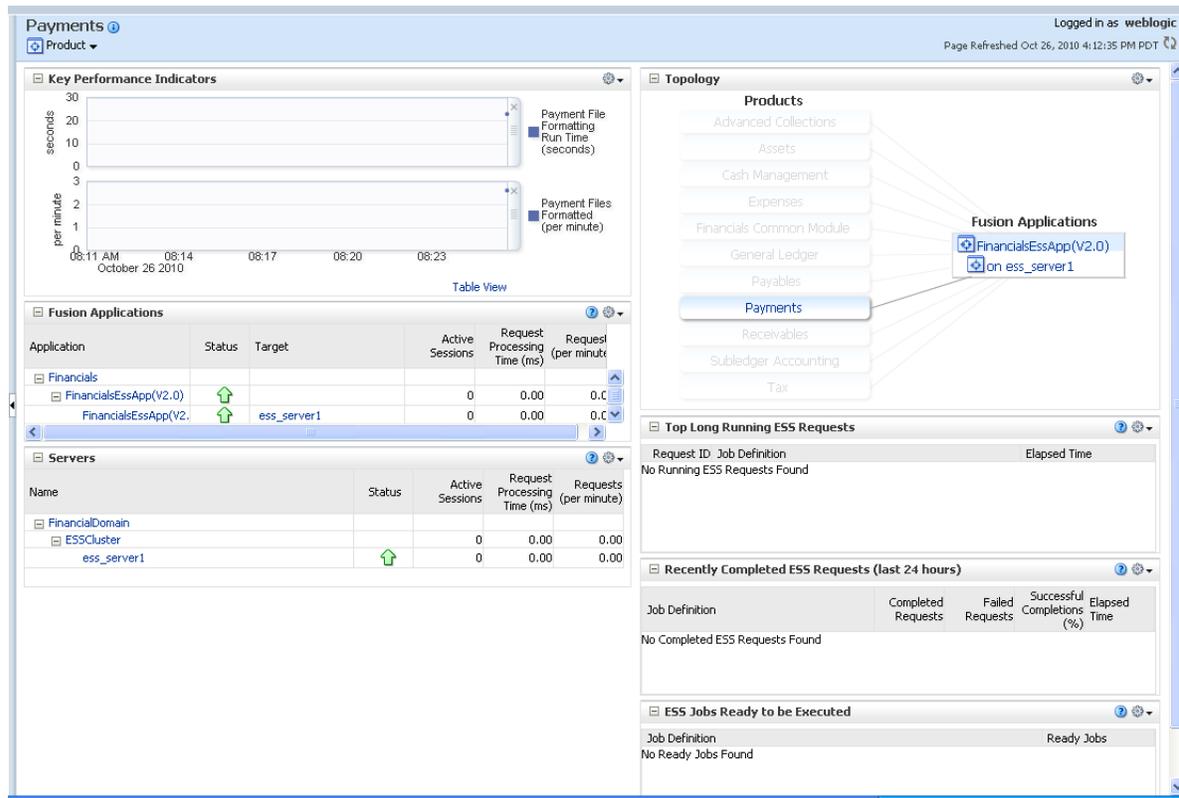
The Recently Completed SOA Composite Requests (Last 24 Hours) section displays information about the most recent SOA composite execution requests completed

within the last 24 hours, including the name of the SOA composite, the name of the partition to which the composite is deployed, the number of completed instances, the percentage of successfully completed instances and the response time in milliseconds for synchronously executed SOA composites.

9.3 The Product Home Page and Sections in Fusion Applications Control

The product home page provides an overview of the Oracle Fusion Applications required to run a given product. An example of a product home page is shown in Figure 9-5.

Figure 9-5 The Product Home Page in Fusion Applications Control



The Oracle Fusion Applications Product Home page displays the following sections:

- [Key Performance Indicators Section](#)
- [Fusion Applications Section](#)
- [Servers Section](#)
- [Topology Section](#)
- [Oracle Enterprise Scheduler Sections](#)

9.3.1 Fusion Applications Section

The Oracle Fusion Applications section displays the Oracle Fusion Applications that are part of the product and have been installed. Each application may include instances running on other servers in the domain.

For more information about this section, see the online help.

9.3.2 Servers Section

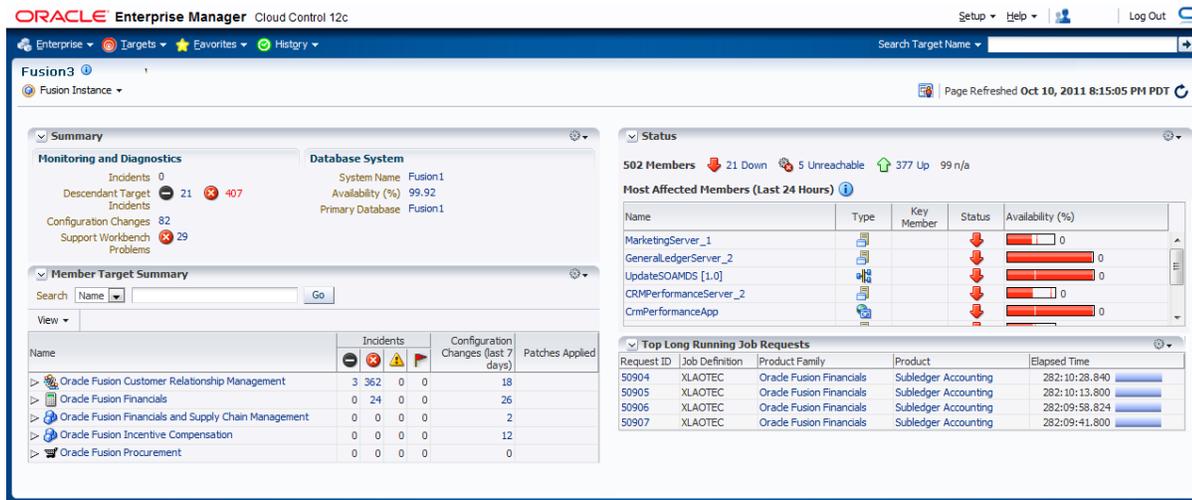
The Servers section displays the names of the domain, clusters and servers running the Oracle Fusion Application. The table listing the clusters and servers displays the status of each item, indicating whether they are up or down. Additionally, the table displays other data such as the number of sessions currently active, the number of requests received per minute and the time required to process a request.

For more information about this section, see the online help.

9.4 The Fusion Instance Overview Page and Sections in Cloud Control

The Fusion Instance home page allows you to monitor a whole instance at a glance, including diagnostic values, a summary of member targets, the status of members and top long running job requests. The Fusion Instance menu provides several administrative options for managing all the product families.

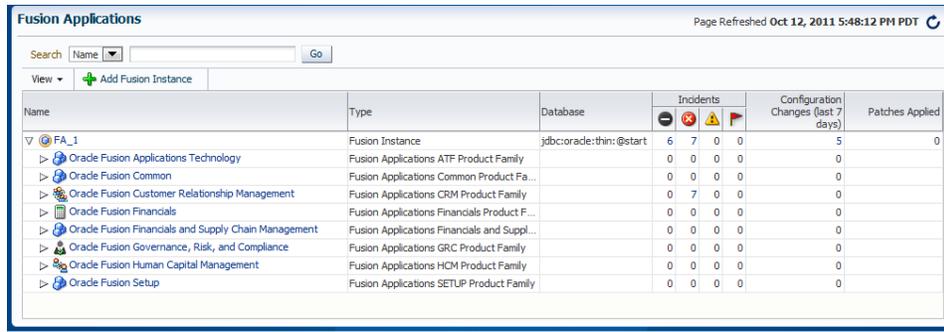
Figure 9–6 The Fusion Instance Overview Page in Cloud Control



To display the Cloud Control Fusion Instance overview page:

1. From the **Targets** menu, choose **Fusion Applications**.

The Fusion Applications target home page displays.



2. In the table on the Fusion Applications page, click the **Fusion Instance** whose overview page you want to display.

The Fusion Instance page displays, as shown in Figure 9–6.

The Fusion Instance overview page displays the following sections:

- [Summary Section](#)
- [Member Target Summary Section](#)
- [Status Section](#)
- [Top Long Running Job Requests](#)

9.4.1 Summary Section

The Summary section displays an overview of the Fusion instance as a whole, including diagnostic and database information.

The following sub-sections display in the Summary section:

- [Monitoring and Diagnostics Region](#)
- [Database System Region](#)

9.4.1.1 Monitoring and Diagnostics Region

The Monitoring and Diagnostics sub-section displays major incidents and changes that have occurred in the Fusion instance, including incidents at the instance level, incidents that have occurred in descendant targets, any configuration changes that have occurred over the last seven days and any support workbench problems. Incidents are listed by severity, namely fatal, critical, warning and escalation levels.

9.4.1.2 Database System Region

The Database System sub-section displays database information related to the Fusion instance, including the name of the database system, the percentage of available space on the system and the name of the database to which the Fusion instance writes its data.

9.4.2 Member Target Summary Section

The Member Target Summary section displays an overview of the member targets of the Fusion instance, including product families and products. The type of information displayed in this sub-section is similar to that displayed in the Cloud Control Fusion Applications overview page described in [Section 2.7.4](#).

9.4.3 Status Section

The Status section displays status information regarding the members of the Fusion instance. The numbers at the top of the Status section indicate how many members of the instance are up, down or unavailable. You can click a number or status for more information.

The table in the Status section displays members that have been affected the most within the last 24 hours. This includes Oracle WebLogic Servers, application deployments, clustered application deployments, Fusion Java EE applications, and so on. The table displays the status and percent availability for each member. You can click individual member items for more information.

9.4.4 Top Long Running Job Requests

This section displays the same information described in [Section 9.2.4](#).

9.5 The Product Family Home Page and Sections in Cloud Control

When selecting a product family from the Fusion Applications overview page or the Fusion Instance page, Cloud Control directs you to the product family home page for the selected product family, as shown in [Figure 9-7](#). The Cloud Control Product Family home page displays an overview of the product family at a glance, including the products and Oracle Fusion applications comprising the product family.

Figure 9-7 The Product Family Home Page in Cloud Control



The Product Family home page includes the following sections:

- [Applications and Infrastructure Section](#)
- [Key Performance Indicators Section](#)
- [SOA Composite Section](#)
- [Topology Section](#)
- [Oracle Enterprise Scheduler Sections](#)

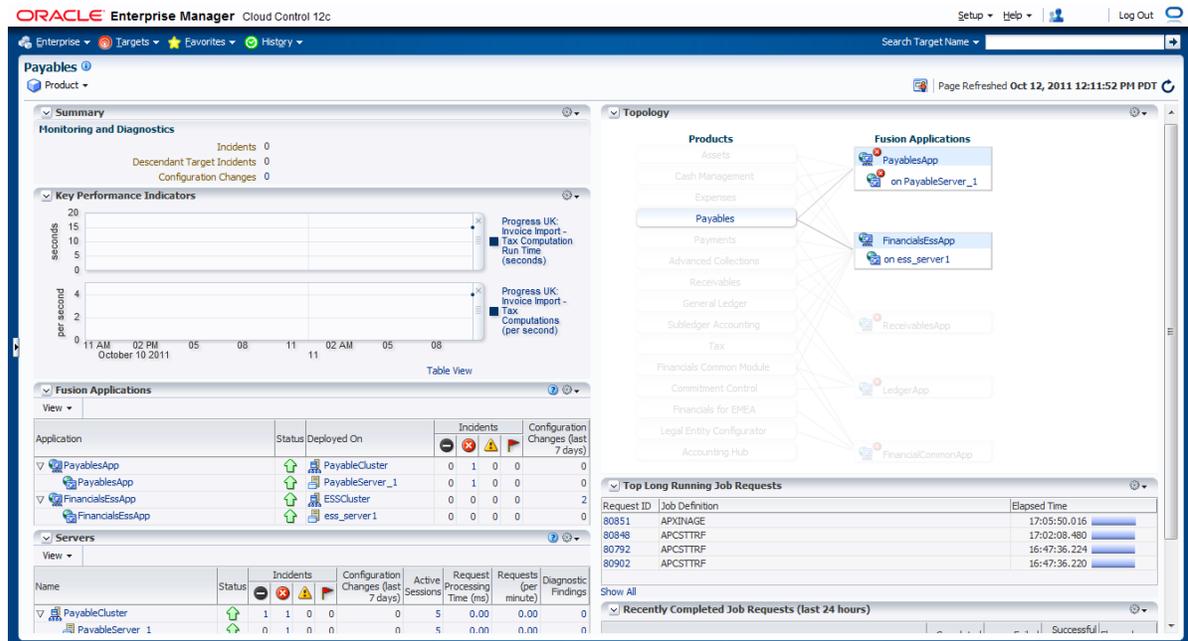
The Product Family menu provides several administrative options for managing the product family, which are similar to those in Fusion Applications Control.

You can click a particular product to display the Cloud Control product home page.

9.6 The Product Home Page and Sections in Cloud Control

When selecting a product from the Fusion Applications overview page, Fusion Instance page, or Product Family home page, Cloud Control directs you to a home page for the selected product, as shown in [Figure 9-8](#). The Cloud Control Product home page displays an overview of a product at a glance, including the Oracle Fusion applications that compose the product and the servers on which they run.

Figure 9–8 The Product Home Page in Cloud Control



The Cloud Control Product home page displays the following sections:

- [Summary Section](#)
- [Key Performance Indicators Section](#)
- [Fusion Applications Section](#)
- [Servers Section](#)
- [Topology Section](#)
- [Oracle Enterprise Scheduler Sections](#)

The Product menu provides several administrative options for managing the product, which are similar to those in Fusion Applications Control.

9.7 Monitoring Product Families and Products

The product family and product home pages allow you to monitor various aspects of your servers and applications.

The main steps in monitoring product families and products are as follows:

- Verify component status.
- View and modify performance metrics.
- Monitor the status of any scheduled jobs.

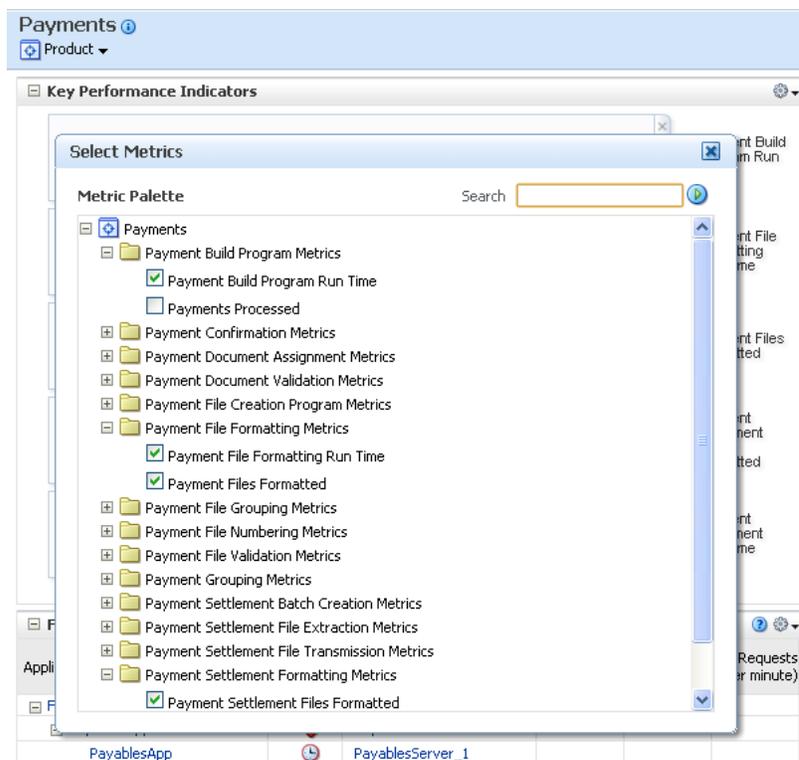
To monitor a product family or product in Fusion Applications Control or Cloud Control:

1. Navigate to the product family home page:
 - **Fusion Applications Control:** From the navigation pane, select the name of the product family.
 - **Cloud Control:**

- a. From the **Targets** menu, choose **Fusion Applications**.
The Fusion Applications target home page displays.
 - b. In the table on the Fusion Applications page, click the **Fusion Instance** whose overview page you want to display.
The Fusion Instance page displays.
 - c. In the Fusion Instance overview page, expand the Fusion instance and click the name of the product family.
 2. Alternatively, navigate to the product home page:
 - **Fusion Applications Control**: From the navigation pane, expand the product family, then **Products**, and then select the product.
 - **Cloud Control**:
 - a. From the **Targets** menu, choose **Fusion Applications**.
The Fusion Applications target home page displays.
 - b. In the table on the Fusion Applications page, click the **Fusion Instance** whose overview page you want to display.
The Fusion Instance page displays.
 - c. In the Fusion Instance overview page, expand the Fusion instance and the product family. Click the name of the product whose home page you want to display.
 3. Use the following sections on the main product family or product page to verify the status of various components.
 - **Applications and Infrastructure**: Verify the availability of Fusion Applications and Oracle Fusion Middleware.
 - **Topology**: View the complete topology of the product or product family. Click a node to view the home page of the selected component.
 - **Fusion Applications**: Verify the status of Fusion cluster applications, Fusion applications and the relationship between domain and server.
 - **Servers**: Verify the status of Fusion cluster applications and all related servers.

For more information about the individual sections on the product family and product home pages, see [Section 9.2](#) and [Section 9.3](#).

For more information about monitoring Fusion applications, see [Section 9.8](#).
 4. In the **Key Performance Indicators** section, view the product or product family main performance metrics. Click the **View Actions** menu to carry out any of the following tasks.
 - **Modifying performance metrics**: You can view and modify the performance metrics in this section by clicking the **View Actions** menu in the section and selecting **Select Metrics**.
The Select Metrics popup window displays.



In the Select Metrics dialog, you can search for a particular metric in the Search field and select or remove the metrics you want to display or hide in the **Key Performance Indicators** section.

For more information about monitoring and modifying performance metrics, see [Section 9.9](#).

- **Saving charts:** You can save metric charts to a file by selecting **View Actions > Save charts**.
 - **Restoring default charts:** You can restore the metric charts to their defaults such that only the default metrics display, by selecting **View Actions > Restore default charts**.
5. Monitor the status of any scheduled jobs by using the Oracle Enterprise Scheduler sections described [Section 9.2.4](#). You can view more scheduled jobs displayed in these sections by clicking **Show All** at the bottom of each section.

Clicking **Show All** displays the Oracle Enterprise Scheduler section in a larger pane such that a larger number of relevant jobs display.

9.8 Monitoring the Detailed Status of Oracle Fusion Applications

The Applications and Infrastructure page allows you to verify the status of a given Oracle Fusion application.

To verify the status of a given Oracle Fusion application:

1. Display the Product Family home page.
 - In Fusion Applications Control, from the navigation pane, select the product family.
 - In Cloud Control, from the **Targets** menu, choose **Fusion Applications**.

In the **Fusion Applications** table, click the desired product family.

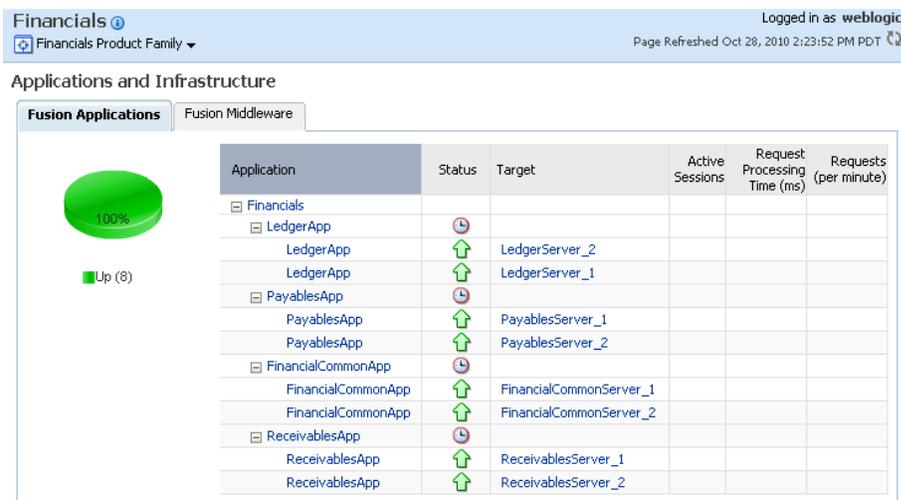
2. Display the Applications and Infrastructure page.
 - In Fusion Applications Control, from the **Product Family** menu, choose **Applications and Infrastructure**
 - In Cloud Control, from the **Product Family** menu, choose **Monitoring** and then select **Applications and Infrastructure**.

The Applications and Infrastructure page displays.

3. Click the **Fusion Applications** tab.

This tab shows the status of a given clustered Oracle Fusion J2EE application. When expanding an application, the individual instances of the Oracle Fusion display, each running on different servers. A status displays for each of these individual application deployments.

A pie chart indicates the percentage of Fusion applications in the product family that are up and running within the farm domain.



9.9 Monitoring Business Performance Metrics for a Product Family

The Key Performance Indicators page displays an overview of the performance metrics of a given product family. Typically, several default metrics display on the page in the form of graphs. A number of metrics are tracked, and these can be added to the Key Performance Indicators page. Each product belonging to a given product family has its own set of metrics.

You can modify the Key Performance Indicators page as follows:

- Add or remove metrics from the page.
- Display performance data in tabular format.
- Save or delete a set of performance charts.
- Re-organize the display of charts, for example move charts up or down.
- Display the performance metrics of another instance of the product on top of the current chart sets, or display earlier metrics of the same product instance on top of current metrics.

- Toggle on or off the grid display on the performance chart set.
- Modify the performance time period.

To view and modify a product family's business performance summary page in Fusion Applications Control or Cloud Control:

1. Navigate to the product family home page.
 - **Fusion Applications Control:** From the navigation pane, choose the product family.
 - **Cloud Control:**
 - a. From the **Targets** menu, choose **Fusion Applications**.
The Fusion Applications target home page displays.
 - b. In the table on the Fusion Applications target home page, click the product family.
2. Display the Key Performance Indicators page.
 - **Fusion Applications Control:** From the **Product Family** menu, select **Key Performance Indicators**.
 - **Cloud Control:** From the **Product Family** menu, select **Monitoring** and then select **Key Performance Indicators**.

The product family home page displays. You can view the key metrics for the product family in the **Key Performance Indicators** page.

3. Display the metrics pane by clicking **Show Metric Palette**.
4. You can modify the metrics that display on the Key Performance Indicators pane.
In the Metrics Palette, select the metrics you want to display in the Key Performance Indicators pane.
5. You can search for a particular metric.
In the **Metrics Palette Search** field, enter the name of the metric you want to find and click the **Search** button.
6. You can change the order in which metrics display on the Key Performance Indicators pane, or delete a metric from the display.
 - Click the bar to the left of the metric name and drag the metric up or down.
 - Click the **Close** button on the upper right corner of the metric to delete it from the display.
7. You can optionally click a metric name for more information about that metric.
The Additional Information dialog displays. From this dialog, you can view more details about the metric using the following links:
 - **Problem Analysis:** Directs to the Problem Analysis page. This page enables you to see a metric with other related metrics, which are affecting the metrics being analyzed. It also shows the related targets and the related targets topology. The Problem Analysis is only available for request-related metrics, such as the Request Processing Time (ms) and Request (per minute).
 - **Analyze Logs:** Displays a message that enables you to view the logs for a target on which the metric is viewed.
 - **Metric Details:** Enables you view historical metric data.

- **Target Home:** Directs you to the home page.

When used together, Problem Analysis and Analyze Logs can help you inspect metrics, target status information, and logs during troubleshooting. These links are not available for all metrics. For more information about these links, see [Section 14.4.6](#).

For information about viewing and modifying the display of product family performance indicators, see [Section 2.5.7](#).

9.10 Monitoring Business Performance Metrics for a Product

The Performance Summary page displays an overview of the performance metrics for a given product. Typically, several default metrics display on the page in the form of graphs. A number of metrics are tracked, and these can be added to the Performance Summary page. Each product belonging to a given product family has its own set of metrics.

You can modify the performance page as follows:

- Add or remove metrics from the page.
- Display performance data in tabular format.
- Save or delete a set of performance charts.
- Re-organize the display of charts, for example move charts up or down.
- Display the performance metrics of another product family on top of the current chart sets, or display earlier metrics of the same product family on top of current metrics.
- Toggle on or off the grid display on the performance chart set.
- Modify the performance time period.

To view and modify a product's business performance summary page in Fusion Applications Control or Cloud Control:

1. Navigate to the product home page.
 - **Fusion Applications Control:** From the navigation pane, expand the product family, then **Products**, and then select the product.
 - **Cloud Control:**
 - a. From the **Targets** menu, choose **Fusion Applications**.
The Fusion Applications target home page displays.
 - b. In the table on the Fusion Applications target home page, expand a product family and click an individual product.
2. Display the Key Performance Indicators page.
 - **Fusion Applications Control:** From the **Product** menu, choose **Key Performance Indicators**.
 - **Cloud Control:** From the **Product** menu, choose **Monitoring** and then select **Key Performance Indicators**.

The product family home page displays. You can view the key metrics for the product family in the **Key Performance Indicators** page.

3. To modify the display of metrics in the Key Performance Indicators page, follow the procedure in [Section 9.9](#), beginning with Step 3.

9.11 Viewing the Routing Topology of an Oracle Fusion Applications Instance, Product Family, or Product

Cloud Control provides a routing topology viewer, which provides a graphical representation of routing relationships across targets, components and elements. It shows you the routing details for the Oracle Fusion Applications environment, from Oracle HTTP Server to the Oracle WebLogic Servers to the Oracle Database. Targets are annotated to indicate status and incidents. Tool tips show additional information about entities and relationships. You can pan and zoom the graph, as in popular map programs, and also expand and collapse desired areas of the graph. The graph can be filtered to show only desired nodes. The topology graph is highly interactive. You can hover over and click on various parts of the graph. The topology graph provides visibility into the runtime environment.

You can also print the topology or save it to a .png file.

To view the topology with Cloud Control:

1. Navigate to a Fusion Instance, product family, or product home page:
 - a. From the **Targets** menu, choose **Fusion Applications**.
The Fusion Applications target home page displays.
 - b. In the table on the Fusion Applications page, click the appropriate target, the instance for the Oracle Fusion Applications or product family. You can also expand a product family to select an individual product.
2. From the **Fusion Instance, Product Family, or Product** menu, choose **Routing Topology**.
The Routing Topology page displays.
3. From the **View** menu, choose one of the following topology views:
 - **Routing: Group by Application:** Choose to view the topology of the applications as they relate to the Oracle WebLogic Server domains and servers, as well as the Oracle Database.
 - **Routing: Group by Middleware:** Choose to view the topology for the middleware targets, such as Oracle WebLogic Server.
 - **Routing: Group by Host:** Choose to view the topology of the targets on the hosts in the enterprise.
 - **Products:** Choose this view the topology of Oracle Fusion Applications products.
4. To see information about a particular target, place your mouse over the target. To view additional information, click **More**. For example, for an Oracle WebLogic Server domain, you can view general information about the domain home location, incidents, and configuration.
5. From the **Products** menu, select and deselect specific products to narrow or expand the scope of the graph.
6. From the **Node Type** menu, select or deselect specific node types to narrow or expand the focus of the graph.
7. From the **Options** menu, you can save or print the image, expand or collapse all of the nodes, change the orientation of the topology to be left to right or top to bottom, or reload the routing navigator.

In addition, you can refresh the status and the metrics or update the topology. To refresh the status and metrics, click **Refresh Target Status and Metrics**. To update the topology shown in the viewer, click **Recreate Topology**. If a target has been added or deleted, the target list and relationships are updated. This option also updates the status and metrics.

8. With Topology Viewer, you can also:
 - Search for specific target or node or type of target or node within the topology with the **Toggle Find Toolbar** widget from the toolbar at top of the page. This makes it easier to find a target if you have many targets. After the target is found, it is highlighted in the topology and the topology is repositioned. This enables you to see the target, if it was not previously visible in the viewing area.
 - View the targets by status. Choose **Show Up Targets**, **Show Down Targets**, **Show Targets with Fatal Incidents**, and **Show Targets with Critical Incidents** from the toolbar at top of the page.
 - Navigate to the home page of a target. Right-click the target, and choose **Home**.
 - Hide or show the status or metrics. From the **Annotations** menu, choose **Status** or **Metrics**.

If you select **Metrics**, one key performance metric for the component is displayed. (You cannot change the metric that is displayed.)
 - View the routing relationships between components.
 - Perform operations directly on the target by right-clicking. The right-click target menu is displayed.
9. Interact with a topology graph.

For more information, see "Interactive Gestures" in the Cloud Control online help.

9.12 Viewing the Configuration Topology of an Oracle Fusion Applications Instance, Product Family, or Product

With Cloud Control, you can view the configuration topology of the IT environment in a graphical format for the Fusion instance, product family, or individual product. The inter-relationships between targets are used to answer questions such as:

- On what targets does this target rely?
- What targets will be affected if this target goes down?
- How are the members of my system related?

To view a configuration topology:

1. Navigate to a Fusion Instance, product family, or product home page:
 - a. From the **Targets** menu, choose **Fusion Applications**.

The Fusion Applications target home page displays.
 - b. In the table on the Fusion Applications page, click the appropriate target, the instance for the Oracle Fusion Applications or product family. You can also expand a product family to select an individual product.
2. From the **Fusion Instance**, **Product Family**, or **Product** menu, choose **Configuration > Topology**.

The Configuration Topology page displays.

3. From the **View** menu, choose one of the following topology views:
 - **Routing: Group by Application:** Choose to view the topology of the applications as they relate to the Oracle WebLogic Server domains and servers, as well as the Oracle Database.
 - **Routing: Group by Middleware:** Choose to view the topology for the middleware targets, such as Oracle WebLogic Server.
 - **Routing: Group by Host:** Choose to view the topology of the targets on the hosts in the enterprise.
 - **Products:** Choose this view the topology of Oracle Fusion Applications products.
4. Use the configuration topology:
 - View existing configuration topologies. See "Viewing a Configuration Topology" in the Cloud Control online help.
 - Create custom configuration topologies. See "Creating a Custom Topology View" in the Cloud Control online help

Monitoring the Oracle Fusion Applications Middle Tier

This chapter describes how to monitor the middle tier using Oracle Enterprise Manager Fusion Applications Control (Fusion Applications Control) and Oracle Enterprise Manager (Cloud Control).

This chapter includes the following topics:

- [Introduction to Monitoring the Middle Tier](#)
- [Monitoring and Viewing Performance Status of Oracle Fusion Middleware](#)
- [Monitoring System Performance Metrics for a Fusion Instance, Product Family, or a Product](#)
- [Monitoring Performance for Key Oracle Fusion Middleware Components with Fusion Applications Control](#)
- [Performing Proactive Performance Monitoring with Cloud Control](#)
- [Tuning Middle Tiers for Oracle Fusion Applications](#)

The procedures referenced in the Oracle Fusion Middleware guides describe using Fusion Middleware Control. These procedures also apply to Fusion Applications Control.

10.1 Introduction to Monitoring the Middle Tier

Both Oracle Enterprise Manager Fusion Applications Control (Fusion Applications Control) and Oracle Enterprise Manager Cloud Control (Cloud Control) automatically and continuously measure run-time performance. The performance metrics are automatically enabled; you do not need to set options or perform any extra configuration to collect them.

You can monitor the following aspects of the middle tier with this control:

- Monitor the state and performance of the Oracle Fusion Applications products and applications and an Oracle Fusion Middleware targets by providing out-of-the-box performance metrics
- Monitor CPU usage, heap usage, Work Manager, JMS servers, and JDBC and JTA usage for Oracle WebLogic Server
- Monitor JVM performance in terms of heap versus non-heap usage, garbage collection, and threads performance
- Monitor applications and web services deployed to an Oracle WebLogic Server

- Monitor a wide range of application metrics for servlets, JSPs, and EJBs are available, as well as web services metrics for faults, invocations, and violations. Such metrics are accessible from a target's home page.
- Access customizable performance summary pages to help administrators monitor performance and diagnose problems. These charts can be modified to display content that is relevant to your domain. A target or component might be added to the chart so that you can compare the performance information for two targets in one chart.

You can also monitor a specific set of metrics for performance tuning information. With Cloud Control, you can also view historical data. See [Section 10.5](#) for more specific information about usage of Cloud Control.

Use the following sections for specifics on using Fusion Applications Control:

- [Section 10.2, "Monitoring and Viewing Performance Status of Oracle Fusion Middleware"](#)
- [Section 10.3, "Monitoring System Performance Metrics for a Fusion Instance, Product Family, or a Product"](#)
- [Section 10.5, "Performing Proactive Performance Monitoring with Cloud Control"](#)

Use the following sections for specifics on using Cloud Control:

- [Section 10.2, "Monitoring and Viewing Performance Status of Oracle Fusion Middleware"](#)
- [Section 10.3, "Monitoring System Performance Metrics for a Fusion Instance, Product Family, or a Product"](#)
- [Section 10.4, "Monitoring Performance for Key Oracle Fusion Middleware Components with Fusion Applications Control"](#)

10.2 Monitoring and Viewing Performance Status of Oracle Fusion Middleware

Monitoring the health of the underlying Oracle Fusion Middleware environment and ensuring that it performs optimally is an important task.

This section contains the following topics:

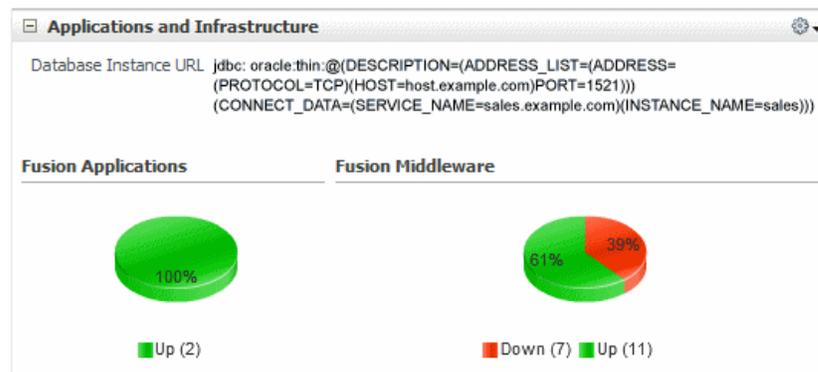
- [Viewing the General Status of Oracle Fusion Middleware Components for a Product Family](#)
- [Viewing the Detailed Status of Oracle Fusion Middleware Components for a Product Family](#)
- [Viewing the Routing Topology of Oracle Fusion Middleware](#)
- [Viewing the Performance of an Oracle WebLogic Server](#)
- [Monitoring the System Performance of Oracle Fusion Applications Deployed to a Cluster](#)

10.2.1 Viewing the General Status of Oracle Fusion Middleware Components for a Product Family

You can view the overall status of the underlying Oracle Fusion Middleware components from the home page of the product family.

To view the general status with either Fusion Applications Control or Cloud Control:

1. Navigate to a product family home page:
 - From Fusion Applications Control, in the navigation pane, select the product family.
 - From Cloud Control:
 - a. From the **Targets** menu, choose **Fusion Applications**.
The Fusion Applications target home page displays.
 - b. In the table on the Fusion Applications target home page, click the appropriate **Product Family** target.
2. View the **Applications and Infrastructure** section. The following figure shows an example of the tab from Fusion Applications Control.



The pie charts in the **Fusion Middleware** section displays the percentage of Oracle Fusion Middleware components that are currently up or down.

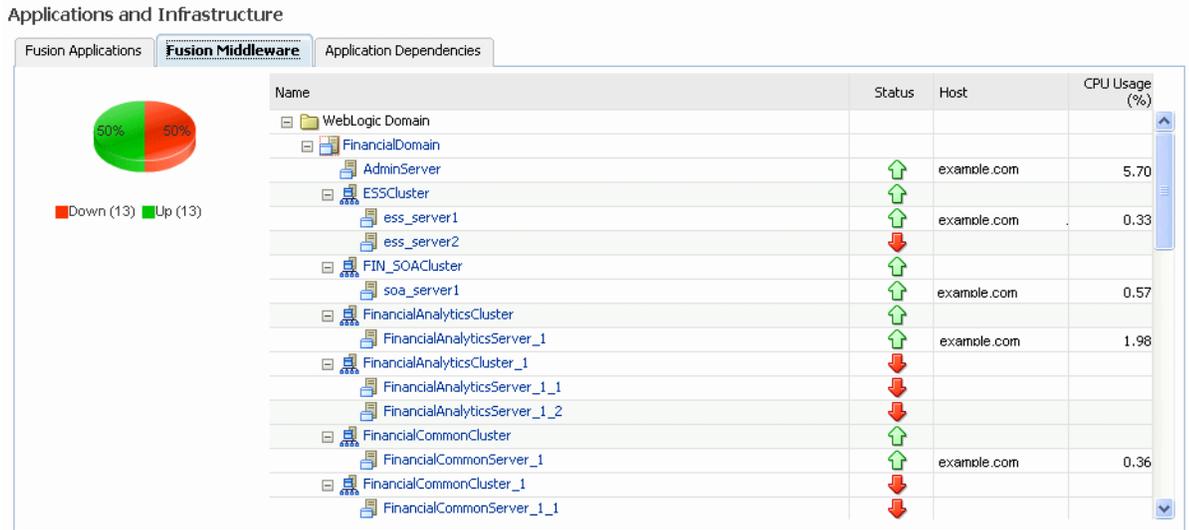
10.2.2 Viewing the Detailed Status of Oracle Fusion Middleware Components for a Product Family

You can view the detailed status of Oracle Fusion Applications and Oracle Fusion Middleware components, including the Administration Server and the Managed Servers in the Oracle WebLogic Server domain and specific Oracle Fusion Middleware components.

To view detailed status with Oracle Fusion Middleware components:

1. Navigate to a product family home page:
 - From Fusion Applications Control, from the navigation pane, select the product family.
 - From Cloud Control:
 - a. From the **Targets** menu, choose **Fusion Applications**.
The Fusion Applications target home page displays.
 - b. In the table on the Fusion Applications target home page, click the appropriate **Product Family** target.
2. From the **Product Family** menu, choose **Applications and Infrastructure**.
3. From the Applications and Infrastructure page, click the **Fusion Middleware** tab.

This tab shows the status of the various components, associated host, and the percentage of CPU that is being used the component. The following figure shows an example of the tab from Fusion Applications Control.



4. Select a specific component to gain more insight.

For example, if you select the Oracle WebLogic Server domain, Fusion Applications Control directs you to the domain home page. This page shows the following.

- A general summary of the domain, along with a link to the Oracle WebLogic Server Administration Console
- Information about the servers, both the Administration Server and the Managed Servers in the domain
- Information about the clusters in the domain
- Information about the deployments in the domain

10.2.3 Viewing the Routing Topology of Oracle Fusion Middleware

This section contains the following topics:

- [Viewing the Routing Topology of Oracle Fusion Middleware Using Fusion Applications Control](#)
- [Viewing the Routing Topology of Oracle Fusion Middleware Using Cloud Control](#)

10.2.3.1 Viewing the Routing Topology of Oracle Fusion Middleware Using Fusion Applications Control

Fusion Applications Control provides a Topology Viewer for the farm. The Topology Viewer is a graphical representation of routing relationships across components and elements of the farm. You can easily determine how requests are routed across components.

The Topology Viewer enables you to easily monitor your Oracle Fusion Applications environment. You can see which entities are up and which are down.

You can also print the topology or save it to a .png file.

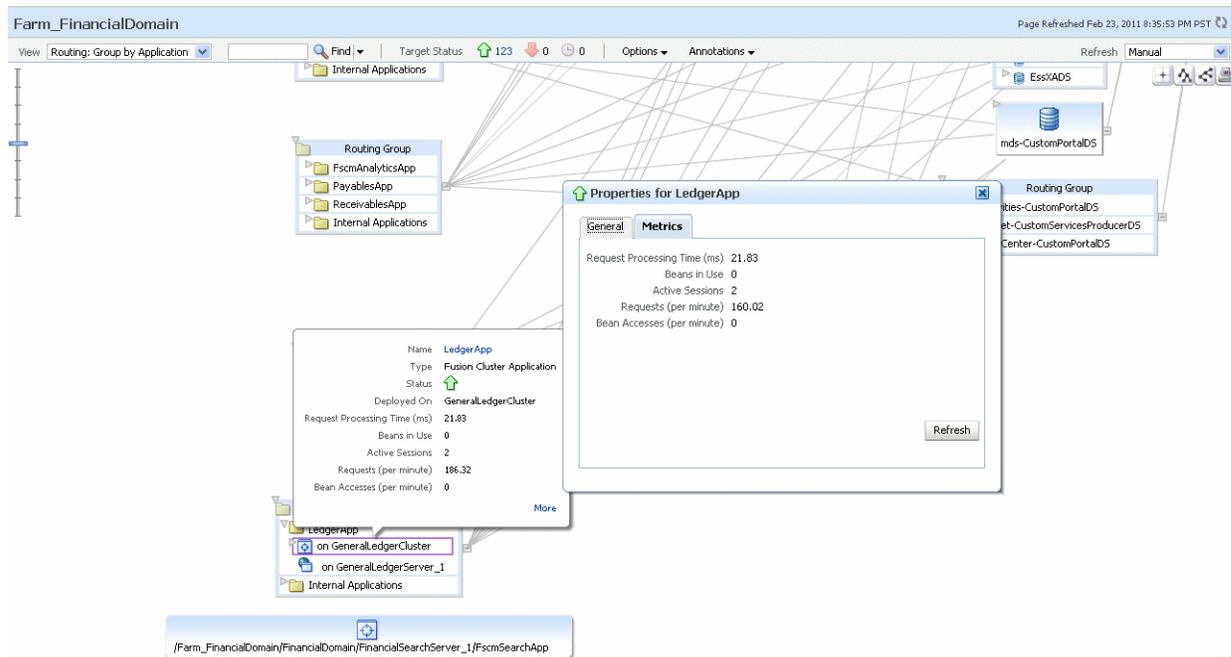
To view the topology with Fusion Applications Control:

1. Click **Topology**.

The Topology Viewer is displayed in a separate window.

2. From the **View** menu, choose **Routing: Group By Application** to view the topology of applications or choose **Routing: Group by Middleware** to view the topology for the middleware.

3. To see information about a particular target, place your mouse over the target. To view additional information, click **More**. The following figure shows additional information for the **LedgerApp** target.



4. From the **Options** menu, you can save or print the image, expand or collapse all of the nodes, change the orientation of the topology to be left to right or top to bottom, or show and hide the routing navigator.

In addition, you can refresh the status and the metrics or update the topology. To refresh the status and metrics, click **Refresh Target Status and Metrics**. To update the topology shown in the viewer, click **Recreate Topology**. If a target has been added or deleted, the target list and relationships are updated. This option also updates the status and metrics.

5. With Topology Viewer, you can also:

- Search for a target within the topology. This makes it easier to find a target if you have many targets. Enter the name in the **Find** box. The target is highlighted and the topology is repositioned so you can see the target if it was not previously visible in the viewing area.

You can also specify criteria for the search. From **Find**, choose the one or more types of **Status** or one or more of **Target Type**, or both.

- View the targets by status. Choose **Up**, **Down**, or **Unknown** from the Target Status at the top of the page.
- Navigate to the home page of a target. Right-click the target, and choose **Home**.

- Hide or show the status or metrics. From the **Annotations** menu, choose **Status** or **Metrics**.
If you select **Metrics**, one key performance metric for the component is displayed. (You cannot change the metric that is displayed.)
 - View the routing relationships between components.
 - Perform operations directly on the target by right-clicking. The right-click target menu is displayed.
6. To change what is visible in the topology view, drag the shaded section in the navigator window, which is located in the bottom right.

Notes:

- If you use Mozilla Firefox, when you click an entity in Topology Viewer to take you back to the main Fusion Applications Control window, focus is not returned to the main window. For example, if you right-click an entity and choose logs from menu, the focus remains on the Topology Viewer window. (If you go back to the main window, the Logs page is correctly displayed.)

To workaroud this problem, make the following change in Firefox:

From the **Tools** menu, choose **Options**, and then **Content**. Click **Advanced**. In the Advanced JavaScript Settings dialog box, select **Raise and lower windows**.

- If you use Internet Explorer, turn off the **Always Open Popups in New Tab** option.
-

10.2.3.2 Viewing the Routing Topology of Oracle Fusion Middleware Using Cloud Control

See [Section 9.11](#).

10.2.4 Viewing the Performance of an Oracle WebLogic Server

See the following resources in the *Oracle Fusion Middleware Administrator's Guide* to monitor the Oracle WebLogic Server:

- "Monitoring an Oracle WebLogic Server" section to view the status of a domain, including the servers, clusters, and deployments in the domain
- "Monitoring an Oracle WebLogic Server Administration or Managed Server" section to view the status of a WebLogic Server Administration Server or Managed Server in Fusion Applications Control
- "Monitoring a Cluster" section to view the status of a cluster, including the servers and deployments in the cluster

For information on locating specific metrics to monitor, see [Section 10.4.4.1](#).

10.2.5 Monitoring the System Performance of Oracle Fusion Applications Deployed to a Cluster

During installation, Oracle Fusion Applications automatically deploys each application to a Managed Server in a cluster. For each Managed Server in the cluster, there is an instance of the application on each Oracle WebLogic Server.

There are times when you want to monitor the system performance of the application on an individual server, and times when you want to monitor the overall performance of the application across all the servers in the cluster.

For example, normally, you would manage the overall performance of the application to determine if there are any performance issues affecting all users of the application, regardless of which instance users access. If you notice a performance problem, you can then drill down to a specific instance of the application to determine if the problem is affecting one or all of the application instances in the cluster. Cloud Control and Fusion Applications Control provides monitoring pages for both of these scenarios.

This section contains the following topics:

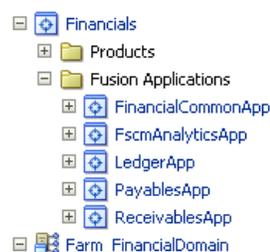
- [Monitoring the Performance of Oracle Fusion Applications Cluster Applications Using Fusion Applications Control](#)
- [Monitoring the Performance of Oracle Fusion Applications Cluster Applications Using Cloud Control](#)
- [Monitoring the Performance of Oracle Fusion Applications Cluster Applications Using Fusion Applications Control](#)
- [Monitoring the Performance of Oracle Fusion Application Instances Using Fusion Applications Control](#)

10.2.5.1 Monitoring the Performance of Oracle Fusion Applications Cluster Applications Using Fusion Applications Control

To monitor the overall performance of the Oracle Fusion Applications cluster applications with Fusion Applications Control:

1. From the navigation pane, expand the product family, then **Fusion Applications**.

Fusion Applications Control lists the applications deployed in the current domain, as shown in the following figure:



2. Monitor the overall performance of the applications on a cluster by clicking the cluster application name, for example, **PayablesApp**.

The Fusion Cluster Application page displays.

3. View performance metrics. From the **Fusion Cluster Application** menu, choose **Performance Summary**.

The Performance Summary page displays. For more information about using this page, see [Section 2.5.7](#).

10.2.5.2 Monitoring the Performance of Oracle Fusion Applications Cluster Applications Using Cloud Control

To monitor the overall performance of the Oracle Fusion Applications cluster applications with Cloud Control:

1. Navigate to a product family home page:
 - a. From the **Targets** menu, choose **Fusion Applications**.
The Fusion Applications target home page displays.
 - b. In the table on the Fusion Applications page, click the appropriate **Product Family** target.
2. From the navigation pane, expand **Fusion Applications**.

Cloud Control lists the applications deployed in the current domain, as shown in the following figure:



3. Monitor the overall performance of the applications on a cluster by clicking the cluster application name, for example, **PayablesApp**.
The Fusion Cluster Application page displays.
4. View performance metrics. From the **Fusion Cluster Application** menu, choose **Monitoring > Performance Summary**.
The Performance Summary page displays. For more information about using this page, see the topic "Using Performance Summary" in the Cloud Control online help.
5. From the Performance Summary page, you can optionally click on a metric name.
The Additional Information dialog displays. From this dialog, you can view more details about the metric with the following links:
 - **Problem Analysis:** Directs to the Problem Analysis page. This page enables you to see a metric with other related metrics, which are affecting the metrics being analyzed. It also shows the related targets and the related targets topology. The **Problem Analysis** is only available for request-related metrics, such as the **Request Processing Time (ms)** and **Request (per minute)**.
 - **Analyze Logs:** Displays a message that enables you to view the logs for a target on which the metric is viewed.
 - **Metric Details:** Enables you view historical metric data.
 - **Target Home:** Directs you to the home page.

When used together, **Problem Analysis** and **Analyze Logs** can help you inspect metrics, target status information, and logs during troubleshooting. These links are not available for all metrics. For more information about these links, see [Section 14.4.6](#).

10.2.5.3 Monitoring the Performance of Oracle Fusion Application Instances Using Fusion Applications Control

To monitor a specific instance of the application with Fusion Applications Control:

1. From the navigation pane, expand the product family, then **Fusion Applications**.
2. Expand the cluster application you want to monitor to show each instance of the application, as shown in the following figure:



3. Monitor the overall performance of an application instance in cluster by clicking one of the application deployment instances, for example, **PayablesApp (PayablesSever_1)**.

The Fusion J2EE Application page displays.

4. View performance metrics. From the **Fusion J2EE Application** menu, choose **Performance Summary**.

The Performance Summary page displays. For more information about using this page, see [Section 2.5.7](#).

10.2.5.4 Monitoring the Performance of Oracle Fusion Application Instances Using Cloud Control

To monitor a specific instance of the application with Cloud Control:

1. Navigate to a product family home page:
 - a. From the **Targets** menu, choose **Fusion Applications**.
The Fusion Applications target home page displays.
 - b. In the table on the Fusion Applications target home page, click the appropriate **Product Family** target.
2. From the navigation pane, expand the product family, then **Fusion Applications**.
3. Expand the cluster application you want to monitor to show each instance of the application, as shown in the following figure:



4. Monitor the overall performance of an application instance in cluster by clicking one of the application deployment instances, for example, **PayablesApp (PayablesSever_1)**.

The Fusion J2EE Application page displays.

5. View performance metrics. From the **Fusion J2EE Application** menu, choose **Performance Summary**.

The Performance Summary page displays. For more information about using this page, see the topic "Using Performance Summary" in the Cloud Control online help.

6. From the Performance Summary page, you can optionally click on a metric name.

The Additional Information dialog displays. From this dialog, you can view more details about the metric with the following links:

- **Problem Analysis:** Directs to the Problem Analysis page. This page enables you to see a metric with other related metrics, which are affecting the metrics being analyzed. It also shows the related targets and the related targets topology. The **Problem Analysis** is only available for request-related metrics, such as the **Request Processing Time (ms)** and **Request (per minute)**.
- **Analyze Logs:** Displays a message that enables you to view the logs for a target on which the metric is viewed.
- **Metric Details:** Enables you view historical metric data.
- **Target Home:** Directs you to the home page.

When used together, **Problem Analysis** and **Analyze Logs** can help you inspect metrics, target status information, and logs during troubleshooting. These links are not available for all metrics. For more information about these links, see [Section 14.4.6](#).

10.3 Monitoring System Performance Metrics for a Fusion Instance, Product Family, or a Product

If you encounter a problem, such as a product running slowly or hanging, you can view more detailed performance information, including aggregate system performance metrics for all the product families or an individual product family. You can also view metrics specific to a particular product, to find out more information about the problem. You can only view aggregate system performance metrics for all the product families with Cloud Control.

This section contains the following topics:

- [Monitoring System Performance Metrics for a Product Family or a Product with Fusion Applications Control](#)
- [Monitoring System Performance Metrics for a Fusion Instance, Product Family, or a Product with Cloud Control](#)

10.3.1 Monitoring System Performance Metrics for a Product Family or a Product with Fusion Applications Control

With Fusion Applications Control, you can view detailed aggregate system performance metrics for the product family or view metrics specific to a particular product.

To view the performance metrics for the product family with Fusion Applications Control:

1. From the navigation pane, select the name of the product family.
2. From the **Product Family** menu, choose **System Performance Summary**.
The System Performance page displays. For more information about using this page, see [Section 2.5.7](#).

To view the performance metrics for a particular product with Fusion Applications Control:

1. From the navigation pane, expand the product family, then **Products**.
2. Select the product you want to monitor.
3. From the **Product** menu, choose **System Performance Summary**.
The System Performance page displays. For more information about using this page, see [Section 2.5.7](#).

10.3.2 Monitoring System Performance Metrics for a Fusion Instance, Product Family, or a Product with Cloud Control

With Cloud Control, you can view aggregate system performance metrics for all the product families by monitoring the Fusion instance or an individual product family. You can also view metrics specific to a particular product.

To view the performance metrics for the product family with Cloud Control:

1. Navigate to a Fusion Instance, product family, or product home page:
 - a. From the **Targets** menu, choose **Fusion Applications**.
The Fusion Applications target home page displays.
 - b. In the table on the Fusion Applications target home page, click the appropriate target, the instance for the Oracle Fusion Applications or product family. You can also expand a product family to select an individual product.
2. From the **Fusion Instance**, **Product Family**, or **Product** menu, choose **Monitoring > System Performance**.

The Performance Summary page displays. For more information about using this page, see the topic "Using Performance Summary" in the Cloud Control online help.

3. From the Performance Summary page, you can optionally click on a metric name.

The Additional Information dialog displays. From this dialog, you can view more details about the metric with the following links:

- **Problem Analysis:** Directs to the Problem Analysis page. This page enables you to see a metric with other related metrics, which are affecting the metrics being analyzed. It also shows the related targets and the related targets topology. The **Problem Analysis** is only available for request-related metrics, such as the **Request Processing Time (ms)** and **Request (per minute)**.
- **Analyze Logs:** Displays a message that enables you to view the logs for a target on which the metric is viewed.
- **Metric Details:** Enables you view historical metric data.
- **Target Home:** Directs you to the home page.

When used together, **Problem Analysis** and **Analyze Logs** can help you inspect metrics, target status information, and logs during troubleshooting. These links are not available for all metrics. For more information about these links, see [Section 14.4.6](#).

10.4 Monitoring Performance for Key Oracle Fusion Middleware Components with Fusion Applications Control

Monitoring Oracle Fusion Middleware performance involves monitoring any Oracle WebLogic Server Managed Servers for a component and as available, performance metrics. The Oracle Fusion Middleware layer of Oracle Fusion Applications automatically and continuously measures run-time performance for middle tier components. Use [Table 10–1](#) to find information on where to find related documentation for the main Oracle Fusion Middleware components in most Oracle Fusion Applications environments. As mentioned previously in this chapter, references to Fusion Middleware Control in the Oracle Fusion Middleware documents apply to Fusion Applications Control.

Table 10–1 Monitoring Oracle Fusion Middleware Components

Component	Documentation
Oracle Business Intelligence Suite	<p>Status of an Oracle WebLogic Server Managed Server for Oracle Business Intelligence Suite: "Monitoring an Oracle WebLogic Server Administration or Managed Server" section in the <i>Oracle Fusion Middleware Administrator's Guide</i> for instructions on accessing the server home page. Specifically, review the Summary section of the page.</p> <p>Performance metrics: "Monitoring Service Levels" section in the <i>Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition</i></p>
Oracle Enterprise Content Management Suite, Content Server	<p>Status of an Oracle WebLogic Server Managed Server for Content Server: "Monitoring an Oracle WebLogic Server Administration or Managed Server" section in the <i>Oracle Fusion Middleware Administrator's Guide</i> for instructions on accessing the server home page. Specifically, review the Summary section of the page.</p> <p>Performance Metrics: "Viewing Performance Information for Oracle Content Server" section in the <i>Oracle WebCenter Content System Administrator's Guide for Content Server</i></p>
Oracle Enterprise Crawl and Search Framework	Performance metrics: Section 10.4.1
Oracle Enterprise Scheduler Service	<p>Status of an Oracle WebLogic Server Managed Server for Oracle Enterprise Scheduler: "Monitoring an Oracle WebLogic Server Administration or Managed Server" section in the <i>Oracle Fusion Middleware Administrator's Guide</i> for instructions on accessing the server home page. Specifically, review the Summary section of the page.</p> <p>Performance metrics: Section 10.4.2.4 and Section 10.4.4.3</p>
Oracle Identity Management, Oracle Access Manager	<p>Status of an Oracle WebLogic Server Managed Server for Oracle Access Manager: "Monitoring an Oracle WebLogic Server Administration or Managed Server" section in the <i>Oracle Fusion Middleware Administrator's Guide</i> for instructions on accessing the server home page. Specifically, review the Summary section of the page.</p> <p>Performance metrics: "Viewing Performance in Fusion Middleware Control" section in the <i>Oracle Fusion Middleware Administrator's Guide for Oracle Access Manager with Oracle Security Token Service</i></p>

Table 10–1 (Cont.) Monitoring Oracle Fusion Middleware Components

Component	Documentation
Oracle Identity Management, Oracle Identity Manager	Status of an Oracle WebLogic Server Managed Server for Oracle Identity Management: "Monitoring an Oracle WebLogic Server Administration or Managed Server" section in the <i>Oracle Fusion Middleware Administrator's Guide</i> for instructions on accessing the server home page. Specifically, review the Summary section of the page. Performance metrics: Section 10.4.3
Oracle Identity Management, Oracle Internet Directory	Status of an Oracle WebLogic Server Managed Server for Oracle Identity Management: "Monitoring an Oracle WebLogic Server Administration or Managed Server" section in the <i>Oracle Fusion Middleware Administrator's Guide</i> for instructions on accessing the server home page. Specifically, review the Summary section of the page. Performance Metrics: "Monitoring Oracle Internet Directory" in the <i>Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory</i>
Oracle Identity Management, Oracle Virtual Directory	Status of an Oracle WebLogic Server Managed Server for Oracle Virtual Directory: "Monitoring an Oracle WebLogic Server Administration or Managed Server" section in the <i>Oracle Fusion Middleware Administrator's Guide</i> for instructions on accessing the server home page. Specifically, review the Summary section of the page. Performance Metrics: "Monitoring Oracle Virtual Directory Using Fusion Middleware Control Metrics" in the <i>Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory</i>
Oracle SOA Suite	Performance metrics: "Monitoring Processing Requests" section in the <i>Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite</i> and Section 10.4.4.2
Oracle WebCenter Portal	Status of an Oracle WebLogic Server Managed Server for Oracle WebCenter Portal: "Monitoring an Oracle WebLogic Server Administration or Managed Server" section in the <i>Oracle Fusion Middleware Administrator's Guide</i> for instructions on accessing the server home page. Specifically, review the Summary section of the page. Performance metrics: "Monitoring Oracle WebCenter Performance" section in the <i>Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter Portal</i>
Oracle Web Tier, Oracle HTTP Server	Performance metrics: "Monitoring Oracle HTTP Server Performance" section in the <i>Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server</i>
Oracle WebLogic Server	Section 10.2.4 and Section 10.4.4.1

The section contains the following topics:

- [Monitoring Oracle Enterprise Crawl and Search Framework](#)
- [Monitoring the Performance of a Single Oracle Enterprise Scheduler Instance](#)
- [Monitoring Oracle Identity Manager Performance Metrics](#)
- [Locating Key Performance Metrics for Specific Components](#)

10.4.1 Monitoring Oracle Enterprise Crawl and Search Framework

Fusion Applications Control automatically and continuously measures the runtime performance of Oracle Enterprise Crawl and Search Framework. If you encounter a problem, you can view the performance metrics to find out more information about the problem.

To monitor performance metrics for Oracle Enterprise Crawl and Search Framework with Fusion Applications Control:

1. From the navigation pane, expand **Enterprise Crawl and Search Framework**.
2. Select the Oracle Enterprise Crawl and Search Framework application for the product family.
3. From the **Enterprise Crawl and Search Framework** menu, choose **Performance Summary** to view detailed performance metrics.

The Performance Summary page displays. For more information about using this page, see [Section 2.5.7](#).

10.4.2 Monitoring the Performance of a Single Oracle Enterprise Scheduler Instance

You can monitor the performance of an Oracle Enterprise Scheduler instance, specifically current and historical performance data, as well as a summary of all performance data. Performance data displays in a chart set, which you save, delete and reorder. Additionally, you can modify the time period for which performance data displays.

This section contains the following topics:

- [Viewing a Summary of Oracle Enterprise Scheduler Activity](#)
- [Viewing Current Oracle Enterprise Scheduler Activity](#)
- [Viewing Historical Oracle Enterprise Scheduler Reports](#)
- [Viewing Oracle Enterprise Scheduler Metrics](#)
- [Managing Chart Sets](#)
- [Selecting the Performance Time Period](#)

10.4.2.1 Viewing a Summary of Oracle Enterprise Scheduler Activity

The Oracle Enterprise Scheduler Performance Summary page displays current charts of job requests such as requests completed per hour, running requests, requests by processing time and requests by status.

To view a summary of Oracle Enterprise Scheduler activity:

1. From the **Scheduling Service** menu, choose **Performance > Service Summary**.

By default, this page displays job requests running during a selected time period in the following categories:

- Requests completed per hour
- Processing times, average, minimum and maximum
- Running requests and maximum running requests
- Requests over the past hour with the status SUCCESS, WARNING, ERROR or CANCEL.

You can change the metrics displayed on the Service Summary performance page. Click the **Show Metric Palette** button to display the metric palette for Oracle Enterprise Scheduler.

2. You can change the duration for which performance data display.
3. Optionally, you can change the display format of performance data on the page. Choose **View > Grid** to remove the grid display format from the charts. Click the **Table View** link to display the performance data in table format.
4. You can overlay an additional scheduler service over the current service performance metrics.
 - a. Choose **Overlay > Another Scheduler Service**.
In the Search and Select: Targets window, enter the name of the target and host server and click **Go** to search for the relevant service.
 - b. From the search results table, select the relevant service and click **Select** to choose the scheduler service to overlay.
5. Additionally, you can take any of the following actions.
 - Save the chart set, as described in [Section 10.4.2.5.1](#).
 - Reorder the chart set, as described in [Section 10.4.2.5.3](#).
 - Revert to an old chart set, as described in [Section 10.4.2.5.4](#).
 - Delete a chart set, as described in [Section 10.4.2.5.2](#).

10.4.2.2 Viewing Current Oracle Enterprise Scheduler Activity

The Oracle Enterprise Scheduler Current Activity page displays pending and running requests by user and application.

To view current Oracle Enterprise Scheduler activity:

1. From the **Scheduling Service** menu, choose **Performance > Current Activity**.
2. You can view pending and running requests by user or application, or workload by work assignment.

To view pending and running requests by user, choose **View > Requests By User**.

To view pending and running requests by application, choose **View > Requests By Application**.

To view workloads by work assignment, choose **View > Workload By Work Assignment**. For more information about work assignment details, see the "Creating or Editing a Work Assignment" section in the *Oracle Fusion Middleware Administrator's Guide for Oracle Enterprise Scheduling Service*.

10.4.2.3 Viewing Historical Oracle Enterprise Scheduler Reports

The Oracle Enterprise Scheduler Historical Reports page displays completed job requests by user, application and work assignment.

To view historical Oracle Enterprise Scheduler reports:

1. From the **Scheduling Service** menu, choose **Performance > Historical Reports**.
2. From the **View** list, choose one of the following report views:
 - Completed Request Statistics By Job Name

- Completed Request Statistics By User
 - Request Metrics By Work Assignment
 - Completed Request Statistics
3. Adjust the filter criteria as required. The filter criteria vary depending on the report view you choose.

Table 10–2 Completed Request Statistics by Job Name

Filter Criterion	What to Do...
Application	From the list, choose the application whose statistics you want to view.
Minimum Run Time	Enter the shortest run time period for jobs whose statistics you want to view. Use the text field to enter a number, and select a unit of time from the list: Seconds, Minutes, Hours or Days .
Completed between	Use the calendar icons on either side of the text fields to select a start and end date and time.

Table 10–3 Completed Request Statistics by User

Filter Criterion	What to Do...
Application	From the list, choose the application whose statistics you want to view.
Minimum Total Run Time	Enter the shortest total time period for jobs whose statistics you want to view. Use the text field to enter a number, and select a unit of time from the list: Seconds, Minutes, Hours or Days .
Completed between	Use the calendar icons on either side of the text fields to select a start and end date and time.

Table 10–4 Request Metrics By Work Assignment

Filter Criterion	What to Do...
Application	From the list, choose the application whose statistics you want to view.
Completed after	Use the calendar icon to choose the date and time after which the work assignment has completed.

Table 10–5 Completed Request Statistics

Filter Criterion	What to Do...
Application	From the list, choose the application whose statistics you want to view.
Completed between	Use the calendar icons on either side of the text fields to choose a start and end date and time.

4. Click **Go** to search for and display the historical data.

10.4.2.4 Viewing Oracle Enterprise Scheduler Metrics

You can view performance metrics for specific components of Oracle Enterprise Scheduler. Each set of metrics displays performance data that is relevant to the component.

To view Oracle Enterprise Scheduler performance metrics:

1. From the Scheduling Service menu, choose **Performance > Service Summary**.
2. Click **Key Performance Indicators**.

3. Click **Show Metric Palette** and select the metrics you want to display for each component.

10.4.2.5 Managing Chart Sets

You can save or delete the chart set displayed in the Oracle Enterprise Scheduler Performance Summary page. You can also re-organize the order of a chart and revert to an old chart set.

This section contains the following topics:

- [Saving a Chart Set](#)
- [Deleting a Chart Set](#)
- [Reordering Charts](#)
- [Reverting to an Old Chart Set](#)

10.4.2.5.1 Saving a Chart Set You can save the chart set displayed in the Oracle Enterprise Scheduler Performance Summary page. Saving a chart set allows preserving a specific set of metrics selected for display, as opposed to the default metrics.

To save a chart set:

1. Display the Oracle Enterprise Scheduler Performance Summary page.
2. Choose **View > Save Chart Set**.
3. In the Save Chart Set dialog box, enter a name for the chart set and click **OK**.
4. Once you have saved the chart set, you can display it by choosing **Chart Set > My Chart Set**.

10.4.2.5.2 Deleting a Chart Set You can delete saved chart set files created in the Oracle Enterprise Scheduler Performance Summary page.

To delete a chart set:

1. Display the Oracle Enterprise Scheduler Performance Summary page.
2. Choose **View > Delete Chart Sets**.
3. In the Delete Chart Sets dialog box, select the name of the chart set to be deleted and click **OK**.

10.4.2.5.3 Reordering Charts You can organize the display of performance charts in the Oracle Enterprise Scheduler Performance Summary page. Charts can be moved up or down, or removed entirely.

To reorder chart sets:

1. Display the Oracle Enterprise Scheduler Performance Summary page.
2. Choose **View > Reorder Charts**.
3. In the Reorder Charts popup window, select the chart you want to move and click **Move Up** or **Move Down**, respectively.

To delete a chart from the Performance Summary page, select the chart and click **Delete**.

4. Click **OK**.

10.4.2.5.4 Reverting to an Old Chart Set If you have re-organized your chart sets, you can revert to the chart set as you had originally saved it.

To revert chart sets:

1. Display the Oracle Enterprise Scheduler Performance Summary page.
2. From the **Chart Set** list, choose the chart you want to revert to its original display.
3. Choose **View > Revert Chart Set**.

10.4.2.6 Selecting the Performance Time Period

In the Oracle Enterprise Scheduler Performance Summary page, you can change the time period for which the performance summary displays.

To select the performance time period:

1. Display the Oracle Enterprise Scheduler Performance Summary page.
2. Change the duration for which job requests display.
 - **All**: Click to display all job requests within a specified, finite period.
 - **2 hours**: Click to display all job requests within the past two hours.
 - **15 minutes**: Click to display all job requests within the past fifteen minutes.
 - **Slider**: Click to display a slider that enables scrolling horizontally to show job requests from a specified period. The slider defaults to whichever time segment is selected, all, 2 hours, or 15 minutes.
 - **Calendar**: Click the calendar icon to select a time period for which you want to display job requests.

10.4.3 Monitoring Oracle Identity Manager Performance Metrics

To monitor performance metrics for Oracle Identity Manager with Fusion Applications Control:

1. From the navigation pane, expand **Identity and Access** and **OIM**.
2. Select the Oracle Identity Management server.
3. From the Oracle Identity Management home page, view the overall metrics.
4. From the **Oracle Identity Manager** menu, choose **Performance Summary** to view detailed performance metrics.

The Performance Summary page displays. For more information about using this page, see [Section 2.5.7](#).

10.4.4 Locating Key Performance Metrics for Specific Components

For each component in the Oracle Fusion Applications environment, there are performance metrics. This section helps you locate the key performance metrics related to software monitoring and configuration.

This section contains the following topics:

- [Monitoring Key Performance Metrics for Oracle WebLogic Server](#)
- [Monitoring Key Performance Metrics for Oracle SOA Suite](#)
- [Monitoring Key Performance Metrics for Oracle Enterprise Scheduler](#)

10.4.4.1 Monitoring Key Performance Metrics for Oracle WebLogic Server

Use [Table 10–6](#) to locate the key performance metrics for Oracle WebLogic Server.

Table 10–6 Oracle WebLogic Server Key Performance Metrics

Category	Metric Name	Location in Fusion Applications Control
JMS	Messages Pending Messages Received	1. From the navigation pane, expand the farm, then WebLogic Domain , and then the domain.
		2. Expand the Managed Server. The Managed Server home page displays.
		3. From the WebLogic Server menu, choose Performance Summary .
		4. From the Performance Summary page, click Show Metric Palette .
		5. From the Metric Palette, expand the JMS Server Metrics folder.
		6. Select the metrics to view the charts.
JDBC	Available Connections Connection Leaks (per minute) Connection Pool Size Connection Requests (per minute) Connections Requests Waiting Connections in Use	1. From the navigation pane, expand the farm, then WebLogic Domain , and then the domain.
		2. Expand the Managed Server. The Managed Server home page displays.
		3. From the WebLogic Server menu, choose Performance Summary .
		4. From the Performance Summary page, click Show Metric Palette .
		5. From the Metric Palette, expand the Datasource Metrics folder.
		6. Select the metrics view the charts.
JVM Memory	Heap Usage (MB) Non Heap Usage (MB)	1. From the navigation pane, expand the farm, then WebLogic Domain , and then the domain.
		2. Expand the Managed Server. The Managed Server home page displays.
		3. From the WebLogic Server menu, choose JVM Performance .
		4. In the Memory section, view the charts.
JVM Garbage Collections	Collections (per minute)	1. From the navigation pane, expand the farm, then WebLogic Domain , and then the domain.
		2. Expand the Managed Server. The Managed Server home page displays.
		3. From the WebLogic Server menu, choose JVM Performance .
		4. In the Memory section, view the Collections (per minute) in the Garbage Collectors table.
	Heap Usage (%)	1. From the navigation pane, expand the farm, then WebLogic Domain , and then the domain.
		2. Expand the Managed Server. The Managed Server home page displays.
		3. From the WebLogic Server menu, choose Performance Summary .
		4. From the Performance Summary page, click Show Metric Palette .
		5. From the Metric Palette, expand the JVM Metrics folder.
		6. Select the metric to view the chart.

Table 10–6 (Cont.) Oracle WebLogic Server Key Performance Metrics

Category	Metric Name	Location in Fusion Applications Control
Oracle WebLogic Server	CPU Usage (%)	1. From the navigation pane, expand the farm, then WebLogic Domain , and then the domain.
	Active Threads	2. Expand the Managed Server. The Managed Server home page displays.
		3. From the WebLogic Server menu, choose Performance Summary . 4. From the Performance Summary page, click Show Metric Palette . 5. From the Metric Palette, expand the JVM Metrics folder. 6. Select the metric to view the chart.
	Deadlocked Threads	1. From the navigation pane, expand the farm, then WebLogic Domain , and then the domain. 2. Expand the Managed Server. The Managed Server home page displays. 3. From the WebLogic Server menu, choose Performance Summary . 4. From the Performance Summary page, click Show Metric Palette . 5. From the Metric Palette, expand the JVM Threads folder. 6. Select the metric to view the chart.
	Threads	1. From the navigation pane, expand the farm, then WebLogic Domain , and then the domain. 2. Expand the Managed Server. The Managed Server home page displays. 3. From the WebLogic Server menu, choose JVM Performance . 4. View the thread metrics in the Threads section.
Server Work Manager	Work Manager Pending Requests	1. From the navigation pane, expand the farm, then WebLogic Domain , and then the domain.
	Work Manager Requests (per minute)	2. Expand the Managed Server. The Managed Server home page displays.
	Work Manager Stuck Threads	3. From the WebLogic Server menu, choose Performance Summary . 4. From the Performance Summary page, click Show Metric Palette . 5. From the Metric Palette, expand the Work Manager folder. 6. Select the metrics to view the charts.

10.4.4.2 Monitoring Key Performance Metrics for Oracle SOA Suite

Use [Table 10–7](#) to locate the key performance metrics for Oracle SOA Suite.

Table 10–7 Oracle SOA Suite Key Performance Metrics

Metric Name	Location in Fusion Applications Control
Audit Trail Threshold (for BPEL)	<ol style="list-style-type: none"> 1. From the navigation pane, expand SOA and then soa-infra. 2. From the SOA Infrastructure menu, choose SOA Administration > BPEL Properties. 3. From the BPEL Service Engine Properties page, view the Audit Trail Threshold. From the BPEL Service Engine Properties page, adjust the Audit Trail Threshold to bolster BPEL engine performance and avoid performance overheads of capturing audit information. From the Audit Level list, you can also level of information collected by the instance tracking infrastructure.
Composite-level rate metrics	<ol style="list-style-type: none"> 1. From the navigation pane, expand SOA and then soa-infra. 2. Expand the partitions. 3. Select a specific SOA composite application. 4. From the SOA Composite menu, choose Monitoring > Performance Summary. 5. From the BPEL Service Engine Properties page, view the Audit Trail Threshold.
Dispatcher System Threads Dispatcher Invoke Threads Dispatcher Engine Threads (for BPEL)	<ol style="list-style-type: none"> 1. From the navigation pane, expand SOA and then soa-infra. 2. From the SOA Infrastructure menu, choose SOA Administration > BPEL Properties. 3. From the BPEL Service Engine Properties page, view the settings for Dispatcher System Threads, Dispatcher Invoke Threads, and Dispatcher Engine Threads. Adjust these thresholds when the response-time is high. Tune the thread configuration to either increase or decrease the number depending upon the BPEL engine processing requirement.
Average Request Processing Time (synchronous requests) Average Request processing time (asynchronous requests)	<ol style="list-style-type: none"> 1. From the navigation pane, expand SOA and then soa-infra. 2. From the SOA Infrastructure menu, choose Monitoring > Request Processing. 3. From the Request Processing page, in the Service Engines and Service Infrastructure sections, view the Average Request Processing Time - Synchronous Requests and the Average Request Processing Time - Asynchronous Requests columns.

Table 10–7 (Cont.) Oracle SOA Suite Key Performance Metrics

Metric Name	Location in Fusion Applications Control
Audit trail	<ol style="list-style-type: none"> 1. From the navigation pane, expand SOA and then soa-infra. 2. Expand the partitions. 3. Select a specific SOA composite application. 4. In the Oracle SOA Composite page, in the Recent Instances section, select an instance. 5. In the Flow Trace window, select a component instance in the Trace section. The audit trail for the instance displays.
State (for a BPEL instance)	<ol style="list-style-type: none"> 1. From the navigation pane, expand SOA and then soa-infra. 2. Expand the partitions. 3. Select a specific SOA composite application. 4. In the Oracle SOA Composite page, in the Component Metrics section, select a BPEL instance. 5. In the BPEL Component page, in the Recent Instances table, view the State column.
State (for an Oracle Mediator instance)	<ol style="list-style-type: none"> 1. From the navigation pane, expand SOA and then soa-infra. 2. Expand the partitions. 3. Select a specific SOA composite application. 4. In the Oracle SOA Composite page, in the Component Metrics section, select an Oracle Mediator instance. 5. In the Mediator Component page, in the Recent Instances table, view the State column.

10.4.4.3 Monitoring Key Performance Metrics for Oracle Enterprise Scheduler

Use [Table 10–8](#) to locate the key performance metrics for Oracle Enterprise Scheduler.

Table 10–8 Oracle Enterprise Scheduler Key Performance Metrics

Metric Name	Location in Fusion Applications Control
Number of active processor threads for each instance	<ol style="list-style-type: none"> 1. From the navigation pane, expand Scheduling Services. 2. Select the Oracle Enterprise Scheduler component. 3. From the Scheduling Service menu, choose Performance > Current Activity. 4. From the View list, choose Workload by Work Assignment. 5. In the Work Assignments with active Workshifts report, view the Thread Allocation columns.
Processing Time (for jobs)	<ol style="list-style-type: none"> 1. From the navigation pane, expand Scheduling Services.
Wait Time	<ol style="list-style-type: none"> 2. Select the Oracle Enterprise Scheduler component. 3. From the Scheduling Service menu, choose Performance > Historical Reports. 4. In the Historical Reports page, from the View list, choose Request Metrics By Work Assignment. 5. In the Results section, view the Wait Time and Processing Time columns.
Job running time	<ol style="list-style-type: none"> 1. From the navigation pane, expand Scheduling Services. 2. Select the Oracle Enterprise Scheduler component. 3. From the Scheduling Service home page, in the Top 10 Long Running Request tab, select a job. 4. In the Request Detail page, under Execution Trail, view the Run Time metric.

10.5 Performing Proactive Performance Monitoring with Cloud Control

Managing applications and the underlying middleware technology can be difficult and IT organizations often have to rely on a variety of specialized tools. This can lead to inefficiency and may introduce complexities and risks.

Cloud Control enables you to discover middleware targets. To access these pages, see the following topic:

- [Monitoring Middle Tier Targets](#)

In addition to drilling down to specific targets for performance metrics, Oracle recommends the following best practices:

- [Creating Monitoring Templates](#)
- [Creating Blackouts](#)
- [Setting Up Notifications](#)

10.5.1 Monitoring Middle Tier Targets

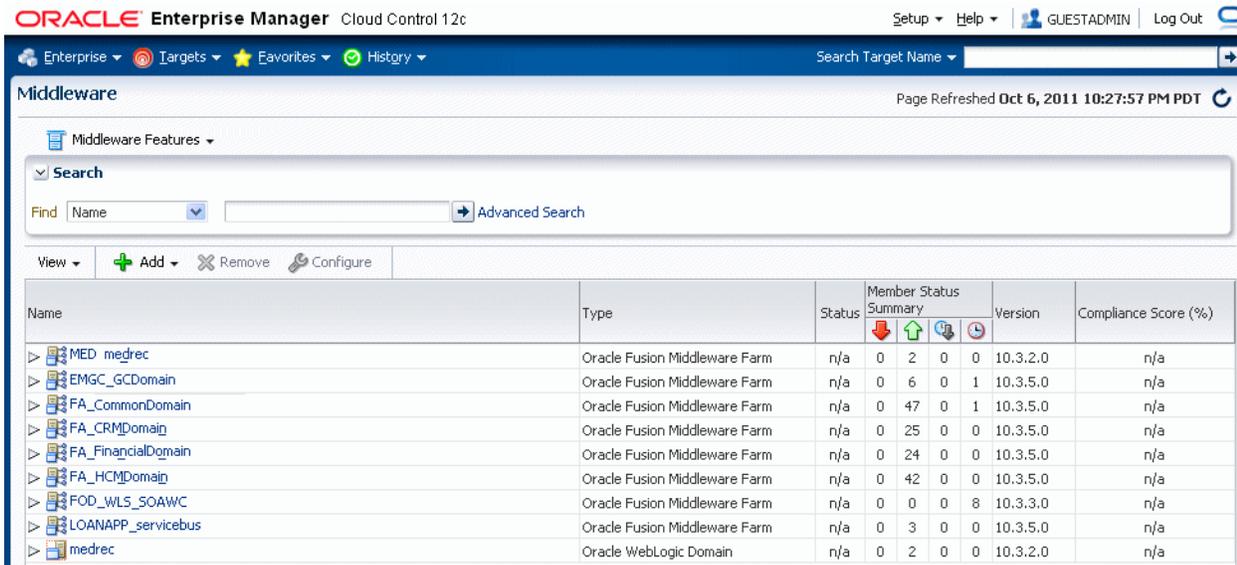
After you have added a middle tier target through discovery, you can view general information about the targets, including their status and availability on the **Middleware** page. You can select a specific type of target, such as Oracle WebLogic Server domain, to narrow the view. You then can drill down into each target to get further details like how the target is performing, where it is deployed, the version, location of its home directory, and so on.

You can also view the number of critical, warning, and error alerts generated for the past 24 hours. These alerts indicate that a particular metric condition has been encountered. For example, an alert is triggered when a metric threshold is reached. Using these details, you can drill down to investigate the target and the problem that triggered the alert.

You can also use Cloud Control to monitor performance metrics as a current real time value (30 seconds, 1 minute, or 5 minutes) or a previous value (past 24 hours, 7 days, or 31 days). The historical information is displayed as graphs and a table. By using graphs, you can easily watch for trends, and by using tables, you can examine details of past metric severity history.

[Figure 10–1](#) shows the **Middleware** page for Oracle WebLogic Server domains in the Oracle Fusion Applications environment.

Figure 10–1 Middleware Target Home Page of Cloud Control



To access the Middleware target home page in Cloud Control:

1. From the **Targets** menu, choose **Middleware**.
2. From the **Search** area, click **Advanced Search**.
3. From the **Type** list, choose a specific target type to monitor, and then click **Search**.
In [Figure 10–1](#), the target selected is **Oracle WebLogic Domain**.

For more information about discovering and monitoring middleware targets, see the chapters in the "Managing Oracle Fusion Middleware" part in the *Oracle Enterprise Manager Getting with Fusion Middleware Management*.

10.5.2 Creating Monitoring Templates

Monitoring templates simplify the task of setting up monitoring for large numbers of targets by allowing you to specify the monitoring and policy settings once and applying them as often as needed. You can save, edit, and apply these templates across one or more targets or groups. A monitoring template is specified for a particular target type and can only be applied to targets of the same type. For example, you can configure a monitoring template for Oracle WebLogic Server domains.

To create a monitoring template from Cloud Control:

1. From the **Enterprise** menu, choose **Monitoring > Monitoring Templates** to access the Monitoring Templates main page.
2. Modify a pre-built monitoring template, or create a new template without any data from another template:

To modify a pre-build monitoring template:

- a. In the **Template Name** field, enter `Fusion Apps` or `FMW`.
- b. Click option **Display Oracle provided templates and Oracle Certified templates**.
- c. Click **Go**.
- d. From the table, select the template, and click **Edit**.

To create a new template, click **Create**.

3. Fill out the tabs accordingly.

In the **Metric Threshold** tab, select the metrics to add to the template. Use the following tables in the *Oracle Fusion Applications Performance and Tuning Guide* for tips on recommended metrics to add to the template:

- Host Metrics
- Oracle Enterprise Scheduler Metrics
- Oracle HTTP Metrics
- Oracle Identity Manager Metrics
- Oracle Internet Directory Metrics
- Oracle SOA Composite Metrics
- Oracle SOA Infrastructure Metrics
- Oracle WebLogic Server Metrics

10.5.3 Creating Blackouts

Blackouts allow you to suspend collection of metrics on a target when performing scheduled maintenance on the target. Blacking out a target suspends monitoring on the target for the duration of the blackout. Blackouts are named entities that mainly consist of a list of targets that participate in the blackout and the schedule that defines when the blackout is in effect. Blackouts may be defined on a single target or on aggregate targets such as groups, hosts, systems, or web applications.

To create a blackout from Cloud Control:

1. From the **Enterprise** menu, choose **Monitoring** > **Blackouts** to access the Blackouts main page.
2. From the table, click **Create**.

Cloud Control displays a wizard page to guide you through the steps required to create a blackout. Click **Help** from any wizard page for more information on specific steps.

10.5.4 Setting Up Notifications

You can associate corrective actions with metrics whose alerts trigger them.

To register corrective actions for these notifications from Cloud Control. From the **Setup** menu, choose **Notifications** > **Notification Methods** access the Notifications Methods page.

See the "Notifications" chapter in the *Oracle Enterprise Manager Cloud Control Administrator's Guide* to complete the configuration.

10.6 Tuning Middle Tiers for Oracle Fusion Applications

One of the most challenging aspects of performance tuning is knowing where to begin. Start with the following documentation resources:

- *Oracle Fusion Applications Performance and Tuning Guide*
- *Oracle Fusion Middleware Fusion Developer's Guide for Oracle Application Development Framework*

- "Managing Log Files and Diagnostic Data" in the *Oracle Fusion Middleware Administrator's Guide*

Diagnosing the Oracle Fusion Applications Middle Tier Performance

This chapter describes how to diagnose Java applications in the middle tier using Oracle Enterprise Manager Cloud Control (Cloud Control):

This chapter includes the following topics:

- [Introduction to Java Diagnostics in the Middle Tier](#)
- [Diagnosing the Oracle Fusion Applications Middle Tier](#)

11.1 Introduction to Java Diagnostics in the Middle Tier

Mission critical Java applications often suffer from availability and performance problems. Developers and IT administrators spend a lot of time diagnosing the root cause of these problems. Many times, the problems occurring in production environments either cannot be reproduced or may take too long to reproduce in other environments. This can cause severe impact on the business.

Oracle Enterprise Manager Cloud Control (Cloud Control) enables you to diagnose performance problems in Java applications in the production environment. By eliminating the need to reproduce problems, it reduces the time required to resolve these problems. This improves application availability and performance. Using Java Virtual Machine (JVM) diagnostics, you can identify the root cause of performance problems in the production environment without having to reproduce them in the test or development environment. It does not require complex instrumentation or restarting of the application to get in-depth application details. Application administrators will be able to identify Java problems or Database issues that are causing application downtime without any detailed application knowledge.

To configure you environment for JVM Diagnostics, see the "Using JVM Diagnostics" chapter in the *Oracle Enterprise Manager Getting Started with Fusion Middleware Management*.

11.2 Diagnosing the Oracle Fusion Applications Middle Tier

This section describes how to diagnose Java applications in the middle tier using Cloud Control:

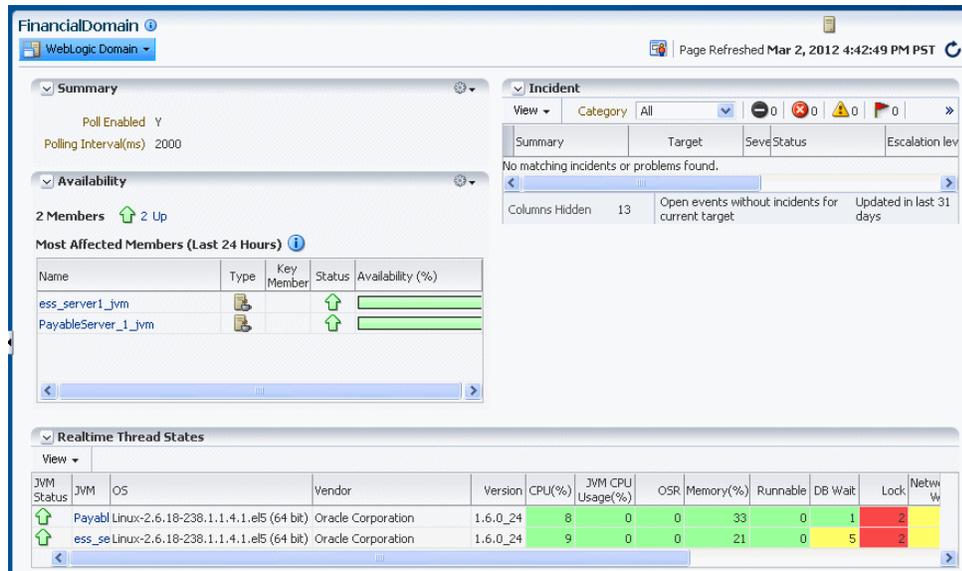
- [Viewing JVM Diagnostics Summary Information](#)
- [Finding the Top Java Methods](#)
- [Finding the Top SQL Queries](#)
- [Analyzing Stuck Threads](#)

- [Drilling Down from JVM Diagnostics to SQL Instances](#)
- [Analyzing Potential Memory Leaks](#)

11.2.1 Viewing JVM Diagnostics Summary Information

To get started with monitoring JVM diagnostics with Cloud Control:

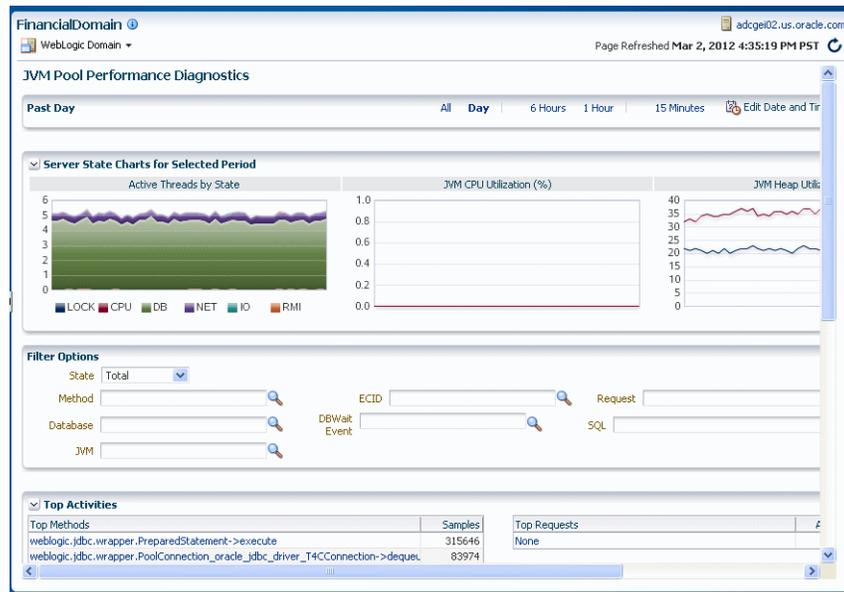
1. From the **Targets** menu, choose **Targets > Middleware**.
The Middleware target home page displays.
2. Search for the Oracle WebLogic Server domains:
 - a. From the **Search** area, click **Advanced Search**.
 - b. From the **Type** list, select Oracle WebLogic Domain or **Oracle WebLogic Server**; deselect the other options.
 - c. Click **Search**.
3. Click on a domain or one of its servers which has JVM Diagnostics agents deployed. For more information about installing JVM Diagnostics, see the *Oracle Enterprise Manager Cloud Control Basic Installation Guide*.
The WebLogic Server Domain Home page or WebLogic Server Home page displays.
4. From the **WebLogic Domain** or **WebLogic Server** menu, choose **Diagnostics > JVM Diagnostics Home**.
The JVM Diagnostics Home page displays.



5. Take note of the following:
 - In the **Summary** section, view the make and version of the JVM, optimization level, and the maximum and minimum heap memory.
 - In the **Active Threads** section, view the number of Java threads that are currently running in the for the domain or server. It is color-coded by thread state.

- In the **Realtime Thread State** section, view the overall thread state, including locks.
- a. From the **WebLogic Domain** or **WebLogic Server** menu, choose **Diagnostics > JVM Diagnostics**.

The JVM Diagnostics Pool Performance Diagnostics page displays for an Oracle WebLogic Server domain and the JVM Performance Diagnostics page displays for a Managed Server. While the JVM Diagnostics Pool Performance Diagnostics page provides information for the pool of JVMs in the domain, the JVM Performance Diagnostics page provides information for a single JVM. The following figure shows the JVM Diagnostics Pool Performance Diagnostics page for the domain:



- b. View the various charts:

Server State Charts For Selected Period section

This section displays the Active Threads, CPU Utilization, and Heap Utilization of IO, CPU, lock, and network resources during the selected time. The number of Java threads (daemon and non-daemon) that are currently running in the virtual machine for this Oracle WebLogic Server

Active Threads by State: This chart displays the number of Java threads that are currently running in the for the domain or server. It is color-coded by thread state.

CPU Utilization: This chart shows the CPU utilization across the JVMs in the pool.

Heap Utilization (%): This chart shows the heap utilization across the JVMs in the pool.

Garbage Collections (Invocations/min): This chart shows the number of times the JVM garbage collector was invoked in the time period. It includes both major and minor garbage collections.

General tab	<p>This section displays data from the JVM itself:</p> <p>Active Threads by State: This chart displays the number of Java threads that are currently running in the for the domain or server. It is color-coded by thread state.</p> <p>Top Requests: This chart shows the top page requests in the selected time period.</p> <p>Top Methods: This chart shows the most expensive Java methods in the selected time period.</p> <p>Top SQLs: This chart displays the list of SQL calls ordered by their cost (the number of samples).</p> <p>Top DBWait Events: This chart shows the cross-tier correlation with the database.</p> <p>Top Databases: This chart shows where the pool of JVM is spending in each area.</p> <p>JVM CPU Utilization: This chart shows the CPU utilization across the JVMs the pool.</p> <p>JVM Heap Utilization (%): This chart shows the heap utilization across the JVMs in the pool.</p>
Threads tab	<p>Threads State Transition: This chart shows how the threads have transitioned from one state to the other in the selected period. You can change the time interval and move it to a different time period by using the quick time selection control at the top of the page. You can hover over the colored bars to see the transition changes from one state to the other, for example from Runnable to Not Active or to Runnable. Click on a bar graph in the State column to view a detailed analysis on the state of the thread. This feature allows you to analyze each sample (JVM snapshot at a specific time) in the monitored data.</p> <p>Metric By Active States: This chart shows how long each of the threads have been in the various states.</p> <p>Top Ecid(s): This chart displays the ECIDs for tracking transactions</p>

Note: The **Compared with** feature enables you to compare the diagnostics across two specified periods of time.

11.2.2 Finding the Top Java Methods

If you have a slow-running application, locate the Java method causing the potential issue.

To find the top Java methods with Cloud Control:

1. From the **Targets** menu, choose **Targets > Middleware**.
The Middleware target home page displays.
2. Search for the Oracle WebLogic Server domains:
 - a. From the **Search** area, click **Advanced Search**.

- b. From the **Type** list, select Oracle WebLogic Domain or **Oracle WebLogic Server**; deselect the other options.
 - c. Click **Search**.
3. Click on a domain or one of its servers which has JVM Diagnostics agents deployed. For more information about installing JVM Diagnostics, see the *Oracle Enterprise Manager Cloud Control Basic Installation Guide*.
The WebLogic Server Domain Home page or WebLogic Server Home page displays.
4. From the **WebLogic Domain** menu or **WebLogic Server** menu, choose **Diagnostics > JVM Performance Diagnostics**.
The JVM Diagnostics Pool Performance Diagnostics page displays for an Oracle WebLogic Server domain and the JVM Performance Diagnostics page displays for a Managed Server.
5. In the **Top Methods** section, review the list of methods ordered by their cost (the number of samples).
6. In the **Top Methods** section, click on a method to add it as a filter.
7. In the Method Stack Tree dialog, click on the plus sign next to the method to investigate further.
8. The Method Stack Tree dialog displays with a call stack that you can traverse.
9. In the Method Stack Tree dialog, select the method, and then click **OK**.
The **Filter Options** section auto-fills the information on the method and the charts update to reflect that method, for example:
 - Heavy use of IO, CPU, lock, and network resources (**Active Threads** chart)
 - Request causing it to be invoked (**Top Requests** chart)
 - SQL or database state it causes (**Top DBStates** and **Top SQLs** charts)
10. After you are done viewing the method, in **Filter Options** section, clear out the **Method** text box filter and click anywhere to remove the filter.

11.2.3 Finding the Top SQL Queries

If you suspect a slow SQL query is causing a network requests or IO issue, find the slowest SQL queries.

To find the top SQL calls using JVM diagnostics with Cloud Control:

1. From the **Targets** menu, choose **Targets > Middleware**.
The Middleware target home page displays.
2. Search for the Oracle WebLogic Server domains:
 - a. From the **Search** area, click **Advanced Search**.
 - b. From the **Type** list, select Oracle WebLogic Domain or **Oracle WebLogic Server**; deselect the other options.
 - c. Click **Search**.
3. Click on a domain or one of its servers which has JVM Diagnostics agents deployed. For more information about installing JVM Diagnostics, see the *Oracle Enterprise Manager Cloud Control Basic Installation Guide*.

The WebLogic Server Domain Home page or WebLogic Server Home page displays.

4. From the **WebLogic Domain** menu or **WebLogic Server** menu, choose **Diagnostics > JVM Performance Diagnostics**.

The JVM Diagnostics Pool Performance Diagnostics page displays for an Oracle WebLogic Server domain and the JVM Performance Diagnostics page displays for a Managed Server.

5. In the **Top SQLs** section, review the list of SQL calls ordered by their cost (the number of samples).
6. In the **Top SQLs** section, click on a SQL call to view the charts for that call.

The **Filter Options** section auto-fills the information on the method and the charts update to reflect that method. Adding the statement as a filter enables you to see everything related to that SQL call, for example:

- Methods that invoke it (**Top Methods** chart)
 - Request causing it to be invoked (**Top Requests** chart)
 - Database state it causes (**Top Databases** chart)
7. After you are done viewing the method, in **Filter Options** section, clear out the **SQL** field and click anywhere to remove the filter.

11.2.4 Analyzing Stuck Threads

If application users report a spinning status indication after clicking in the application, investigate the stuck threads.

To find the top SQL calls using JVM diagnostics with Cloud Control:

1. From the **Targets** menu, choose **Targets > Middleware**.

The Middleware target home page displays.

2. Search for the Oracle WebLogic Server domains:
 - a. From the **Search** area, click **Advanced Search**.
 - b. From the **Type** list, select Oracle WebLogic Domain or **Oracle WebLogic Server**; deselect the other options.
 - c. Click **Search**.

3. Click on a domain or one of its servers which has JVM Diagnostics agents deployed. For more information about installing JVM Diagnostics, see the *Oracle Enterprise Manager Cloud Control Basic Installation Guide*.

The WebLogic Server Domain Home page or WebLogic Server Home page displays.

4. From the **WebLogic Domain** menu, choose **Diagnostics > Live Thread Analysis**.

The Live Thread Analysis page displays.

5. In the **JVMs** section, click on a thread to show details in the **JVM Threads** section.
6. In the **JVM Threads** section, look for a thread having the prefix **[STUCK THREAD]** and click on it.
7. In the **Thread Info** and **Thread Stack** sections, look at the **Current Call**, **File Name**, **Line**, and **State** for the thread.

This information provides you with the key information on how to locate the code that is causing the problem:

- **Current Call:** This field displays the name of the method call where the code is stuck.
 - **File Name:** This column identifies the file with the problem.
 - **Line:** This column identifies the line number in the file where the problematic code is.
 - **State:** This column displays the state of the thread (for example, CPU, IO, Network, DB Wait, Lock, and so on).
8. Look for the **Lock Held** in the **Thread Info** section.
 9. If the stuck thread is in the **DB Wait** state, then click on the link and go directly to the database session to see what that thread is doing in the database, or use the technique described in [Section 11.2.5](#).

11.2.5 Drilling Down from JVM Diagnostics to SQL Instances

If you issue an SQL query and it does not return, then analyze the SQL statement.

To analyze SQL from Cloud Control:

1. From the **Targets** menu, choose **Targets > Middleware**.

The Middleware target home page displays.

2. Search for the Oracle WebLogic Server domains:

- a. From the **Search** area, click **Advanced Search**.
- b. From the **Type** list, select Oracle WebLogic Domain or **Oracle WebLogic Server**; deselect the other options.
- c. Click **Search**.

3. Click on a domain or one of its servers which has JVM Diagnostics agents deployed. For more information about installing JVM Diagnostics, see the *Oracle Enterprise Manager Cloud Control Basic Installation Guide*.

The WebLogic Server Domain Home page or WebLogic Server Home page displays.

4. From the **WebLogic Domain** menu, choose **Diagnostics > JVM Performance Diagnostics**.

JVM Diagnostics Pool Performance Diagnostics page displays for an Oracle WebLogic Server domain and the JVM Performance Diagnostics page displays for a Managed Server.

5. In the **Top SQLs** section, click on a SQL call to view the charts for that call.

6. In **Filter Options** section, in the **SQL** field, copy the ID of the SQL call.

7. Click on **Databases** in secondary navigator.

8. Select the relevant target database.

9. Click the **Performance** tab.

10. In the **Additional Instance Monitoring Links** section, click **Search SQL**.

11. Paste the ID of the SQL call into the relevant field with any other choices you may need and then click **Search**.

12. Analyze the SQL.

11.2.6 Analyzing Potential Memory Leaks

To find and analyze memory leaks, you can use Cloud Control to take and analyze snapshots of the heap.

Analyzing heap requires a large amount of free space in the Oracle Database tablespace being used. As a standard practice, ensure you have five times the size of heap dump file being loaded in the tablespace. Since you know the size of your dump file, make sure that there is adequate space to accommodate the dump file before it is loaded into the database.

To use heap snapshot to analyze memory leaks:

- [Task 1, "Review Metrics for a Potential Memory Leak"](#)
- [Task 2, "Perform a Live Heap Analysis"](#)
- [Task 3, "Create a Heap Snapshot"](#)

Task 1 Review Metrics for a Potential Memory Leak

To create a snapshot of the heap for later loading and examination for leaks:

1. From the **Targets** menu, choose **Targets > Middleware**.
The Middleware target home page displays.
2. Search for the Oracle WebLogic Server domains:
 - a. From the **Search** area, click **Advanced Search**.
 - b. From the **Type** list, select Java Virtual Machine; deselect the other options.
 - c. Click **Search**.
3. Click on a Java Virtual Machine with JVM Diagnostics agents deployed. For more information about installing JVM Diagnostics, see the *Oracle Enterprise Manager Cloud Control Basic Installation Guide*.
The Java Virtual Machine Home page displays.
4. From the **Java Virtual Machine** menu, choose **Performance Summary**.
The Performance Summary page displays.
5. Review the following metrics for any periods of time where the **Warning Thresholds** or **Critical Thresholds** were reached:

JVM GC Overhead This metric shows the percentage of CPU the JVM is using for garbage collections in relation to total CPU usage including servicing application workload (the lower, the better). This metric and its trending can help determine when the garbage collector is making the CPU spin on garbage collection instead of on application workload.

JVM Heap Usage (%) This metric shows the heap utilization for the JVM. This metric provides an indicator of heap size as they fluctuate between garbage collections.

JVM Heap Used After GC This metric shows the percent of heap utilization used after garbage collection. This metric and its trend over time can provide a good indication that there is a leak. For example, if the chart trending up while the application load is stable, then it is possible there is a leak.

If any of the metrics exceed the **Warning Thresholds** or **Critical Thresholds**, it could indicate memory is a factor in the JVM performance and availability. It could mean there is a memory leak or that the JVM heap configuration is too small for the application load. If the heap configuration is correct, assume there is a leak and investigate the cause.

6. If any of the metrics exceed the **Warning Thresholds** or **Critical Thresholds**, proceed to [Task 2](#).

Task 2 Perform a Live Heap Analysis

To create a snapshot of the heap for later loading and examination for leaks:

1. From the **Java Virtual Machine** menu, choose **Live Heap Analysis**.

The Live Heap Analysis page displays.

2. Review the top portion of the page to analyze the heap and the number of objects added to the garbage collector; review the **JVM Class Detail** table to review the largest-size objects in the heap.

For more information about using this page, see the topic "Viewing the Real-Time Heap Data" in the Cloud Control online help.

Task 3 Create a Heap Snapshot

To create a snapshot of the heap for later loading and examination for leaks:

1. On the Live Heap Analysis page, click **Create Heap Snapshot**.

The Heap Snapshot page displays.

2. Provide the settings for your environment.

Notice the **Heap Snapshot Type** under **Heap Snapshot Only** option enables you to pick either Oracle's JVMD format or a HROFF for use with other tools.

For more information about using this page, see the topic "Taking a Heap Snapshot" in the Cloud Control online help.

3. If you selected **Heap Snapshot Only** option, click **Take Snapshot**.

The heap snapshot is generated and the file name in which it is stored is displayed. You can upload the heap snapshot and analyze it using appropriate options from the Heap Snapshots menu.

4. From the **Java Virtual Machine** menu, choose **Heap Snapshots**.

The Available Heap Snapshots page displays.

5. Select the heap you created from the table, and then click **Detail**.

The Heaps > Roots page displays. The **Roots** tab displays the objects reachable by roots, which are objects that are directly reachable from the JVM itself.

Heaps > Roots

Roots Usage Dominator Roots

View ▾ Compare with ... Swap Detach

adc4110056:17727:ssl 05-MAR-12 21:10:02 : 722/1365 MB [Change Heap](#)

Root	Objects	Total Memory(KB)	Adjusted Memory(KB)
System Classes	5339964	409,947	409,947
ThreadOther	4560698	331,726	6
System Other	4550196	330,965	1
JNI	4550200	330,965	0
Thread Handles	4550172	330,964	0
ObjectMonitor	4550175	330,964	0

6. Click a root name to drill down and view the objects that consume a lot of memory.
The Top 40 Objects page displays.
7. In the Heaps > Roots page, click **Compare with** to compare the current heap with another previously taken heap dump.
When comparing heaps, load the bigger one first. Otherwise you may see negative deltas.
8. In the Select a heap record dialog, select the second heap, and then click **OK**.
9. Compare both the heaps. Compare the number of objects (**Objects**) and the occupied memory size (**Adjusted Memory**) in each heap dump. This measure indicates the objects that are growing over the period of time when the snapshots were taken.
10. Drill down into the root which had the largest delta in order to find the biggest memory leak.

Monitoring and Tuning Oracle Database for Oracle Fusion Applications

This chapter describes how to monitor, diagnose, tune, and test changes to Oracle Database using Oracle Enterprise Manager Database Control. Oracle Database provides several tools that enable you to monitor performance, diagnose and tune problems, and test database changes on a test system before implementing them to a production system. Most of these tools are accessible using Oracle Enterprise Manager Database Control.

This chapter includes the following sections:

- [Monitoring Oracle Database](#)
- [Tuning Oracle Database](#)
- [Testing Oracle Database Changes](#)

12.1 Monitoring Oracle Database

Instructions for how to monitor Oracle Database using Database Control is provided in the following topics:

- [Database Statistics](#)
- [Monitoring Database Alerts](#)
- [Monitoring User and Session Activity](#)
- [Monitoring the Database Instance](#)
- [Monitoring the Host System](#)

12.1.1 Database Statistics

Database statistics provide information about the type of load on the database and the internal and external resources used by the database. To accurately diagnose performance problems with the database using ADDM, statistics must be available.

Oracle Database generates many types of cumulative statistics for the system, sessions, and individual SQL statements. Oracle Database also tracks cumulative statistics about segments and services. Automatic Workload Repository (AWR) automates database statistics gathering by collecting, processing, and maintaining performance statistics for database problem detection and self-tuning purposes.

By default, the database gathers statistics every hour and creates an **AWR snapshot**, which is a set of data for a specific time that is used for performance comparisons. The delta values captured by the snapshot represent the changes for each statistic over the

time period. Statistics gathered by AWR are queried from memory. The gathered data can be displayed in both reports and views.

The database statistics collected and processed by AWR include:

- Time model statistics

Time model statistics measure the time spent in the database by operation type. The most important time model statistic is **database time (DB time)**. Database time represents the total time spent in database calls by foreground sessions, and is an indicator of the total instance workload.

- Wait event statistics

Wait events are incremented by a session to indicate that the session had to wait for an event to complete before being able to continue processing. When a session has to wait while processing a user request, the database records the wait by using one of a set of predefined wait events. The events are then grouped into wait classes.

- Session and system statistics

A large number of cumulative database statistics are available on a system and session level. Some of these statistics are collected by AWR.

- Active session history statistics

The Active Session History (ASH) statistics are samples of session activity in the database. The database samples active sessions every second and stores them in a circular buffer in the System Global Area (SGA). Any session that is connected to the database and using CPU, or is waiting for an event that does not belong to the idle wait class, is considered an active session.

- High-load SQL statistics

SQL statements that are consuming the most resources produce the highest load on the system, based on criteria such as elapsed time and CPU time.

For more information about database statistics, see the "Gathering Database Statistics Using the Automatic Workload Repository" section in the *Oracle Database 2 Day + Performance Tuning Guide*.

12.1.2 Monitoring Database Alerts

Oracle Database includes a built-in alerts infrastructure to notify you of impending problems with the database. By default, Oracle Database enables the following alerts:

- Tablespace Usage
- Snapshot Too Old
- Recovery Area Low on Free Space
- Resumable Session Suspended

In addition to these default alerts, you can use performance alerts to detect any unusual changes in database performance.

This section contains the following topics:

- [Configuring Database Alerts](#)
- [Responding to Database Alerts](#)

For more information about database alerts, see the "Monitoring Performance Alerts" chapter in the *Oracle Database 2 Day + Performance Tuning Guide*.

12.1.2.1 Configuring Database Alerts

A **metric** is the rate of change in a cumulative statistic. This rate can be measured against a variety of units, including time, transactions, or database calls. For example, the number of database calls per second is a metric. You can set thresholds on a metric so that an alert is generated when the threshold is passed.

Performance alerts are based on metrics that are performance-related. These alerts are either environment-dependent or application-dependent.

Environment-dependent performance alerts may not be relevant on all systems. For example, the `AVERAGE_FILE_READ_TIME` metric generates an alert when the average time to read a file exceeds the metric threshold.

Application-dependent performance alerts are typically relevant on all systems. For example, the `BLOCKED_USERS` metric generates a performance alert when the number of users blocked by a particular session exceeds the metric threshold.

To obtain the most relevant information from performance alerts, set the threshold values of performance metrics to values that represent desirable boundaries for your system. You can then fine-tune these values over time until your system meets or exceeds your performance goals.

To set thresholds for alerts, use the Metric and Policy Settings page, accessible from Related Links on the Database Home page.

12.1.2.2 Responding to Database Alerts

When an alert is generated by Oracle Database, it appears under Alerts on the Database Home page.

To respond to a database alert, locate the alert that you want to investigate under Alerts on the Database Home page and click the **Message** link. Follow the recommendations provided on the page.

After taking the necessary corrective measures, you can acknowledge an alert by clearing or purging it. Clearing an alert sends the alert to the Alert History, which can be viewed from the Database Home page under Related Links. Purging an alert removes it from the Alert History.

12.1.3 Monitoring User and Session Activity

The Database Performance page displays information in three sections that enable you to assess the overall performance of the database in real time.

The Average Active Sessions chart of the Performance page shows the average load on the database. The **average active sessions** for a time period equals the total DB time of all sessions during this period divided by the elapsed time (wall clock time) for this period. The chart shows which active sessions are running on the CPU or waiting on an event.

The wait classes show how much database activity is consumed by waiting for a resource such as disk I/O. Values that use a larger block of active sessions represent bottlenecks caused by a particular wait class, as indicated by the corresponding color in the legend. To identify each wait class, move your cursor over the block in the Average Active Sessions chart corresponding to the class.

Click the block of color on the chart or its corresponding wait class in the legend to drill down to the wait class. The Active Sessions Working page for the wait class appears. From this page, you can view the details of wait classes in the following dimensions:

- **Top SQL**

On the Active Sessions Working page, the Top Working SQL table shows the database activity for actively running SQL statements that are consuming CPU resources. If one or several SQL statements are consuming most of the activity, then you should investigate them.
- **Top sessions**

On the Active Sessions Working page, the Top Working Sessions table displays the top sessions waiting for the corresponding wait class during the selected time period.

A **session** is a logical entity in the database instance memory that represents the state of a current user login to the database. A session lasts from the time a user logs in to the database until the user disconnects. If a single session is consuming the majority of database activity, then you should investigate it.
- **Top services**

The Top Services table displays the top services waiting for the corresponding wait event during the selected time period.

A **service** is a group of applications with common attributes, service-level thresholds, and priorities. If a service is using the majority of the wait time, then you should investigate it.
- **Top modules**

The Top Modules table displays the top modules waiting for the corresponding wait event during the selected time period.

Modules represent the applications that set the service name as part of the workload definition. If a single module is using the majority of the wait time, then it should be investigated.
- **Top actions**

The Top Actions table displays the top actions waiting for the corresponding wait event during the selected time period.

Actions represent the jobs that are performed by a module. If a single action is using the majority of the wait time, then you should investigate it.
- **Top clients**

The Top Clients table displays the top clients waiting for the corresponding wait event during the selected time period. A client can be a Web browser or any client process that initiates requests for an operation to be performed by the database. If a single client is using the majority of the wait time, then you should investigate it.
- **Top PL/SQL**

The Top PL/SQL table displays the top PL/SQL subprograms waiting for the corresponding wait event during the selected time period. If a single PL/SQL subprogram is using the majority of the wait time, then you should investigate it.
- **Top files**

The Top Files table displays the average wait time for specific files during the selected time period. This data is available from the Active Sessions Waiting: User I/O page.
- **Top objects**

The Top Objects table displays the top database objects waiting for the corresponding wait event during the selected time period. This data is available from the Active Sessions Waiting: User I/O page.

For more information about monitoring user and session activity, see the "Monitoring User Activity" section in the *Oracle Database 2 Day + Performance Tuning Guide*.

12.1.4 Monitoring the Database Instance

In the Average Active Sessions section of the Performance page, you can use the instance charts to monitor database instance activity in the following dimensions:

- Throughput

Database **throughput** measures the amount of work the database performs in a unit of time. The Throughput charts show any contention that appears in the Average Active Sessions chart. The Throughput charts on the Performance page display:

- Number of logons, transactions, physical reads, and redo size per second
- Number of physical reads and redo size per transaction

Compare the peaks on the Throughput charts with the peaks on the Average Active Sessions chart. If the Average Active Sessions chart displays a large number of sessions waiting, indicating internal contention, and throughput is low, then consider tuning the database.

- I/O

The I/O charts show I/O statistics collected from all database clients. The I/O wait time for a database process represents the amount of time that the process could have been doing useful work if a pending I/O had completed. Oracle Database captures the I/O wait times for all important I/O components in a uniform fashion so that every I/O wait by any Oracle process can be derived from the I/O statistics.

The Latency for Synchronous Single Block Reads chart shows the total perceived **I/O latency** for a block read, which is the time difference between when an I/O request is submitted and when the first byte of the transfer arrives. Most systems are performing satisfactorily if latency is fewer than 10 milliseconds.

You can also monitor I/O by function, type, and consumer groups using the various charts provided.

- Parallel executions

The Parallel Execution charts show system metrics related to parallel queries. **Metrics** are statistical counts per unit. The unit could be a time measure, such as seconds, or per transaction, or session.

A parallel query divides the work of executing a SQL statement across multiple processes. The charts show parallel queries that were waiting for a particular wait event that accounted for the highest percentages of sampled session activity.

- Services

The Services charts show services waiting for the corresponding wait event during the time period shown. Services represent groups of applications with common attributes, service-level thresholds, and priorities. Only active services are shown.

For more information about monitoring the database instance, see the "Monitoring Instance Activity" section in the *Oracle Database 2 Day + Performance Tuning Guide*.

12.1.5 Monitoring the Host System

The Host chart on the Performance page displays utilization information about the system hosting the database.

Using the Host chart, you can view CPU, memory, and disk utilization for the host system. To determine if the host system has enough resources available to run the database, establish appropriate expectations for the amount of CPU, memory, and disk resources that your system should be using. You can then verify that the database is not consuming too many of these resources.

For more information about monitoring the host system, see the "Monitoring Host Activity" section in the *Oracle Database 2 Day + Performance Tuning Guide*.

12.2 Tuning Oracle Database

Performance tuning is an iterative process. Removing the first bottleneck (a point where resource contention is highest) may not lead to performance improvement immediately because another bottleneck might be revealed that has an even greater performance impact on the system.

Oracle Database provides several tools that enable you to diagnose and tune performance problems. Automatic Database Diagnostic Monitor (ADDM) analyzes statistics to provide automatic diagnosis of major performance problems. You can also perform your own analysis using statistics from AWR and ASH reports, wait events, and SQL trace files to identify other bottlenecks in the database.

Instructions for diagnosing and tuning performance problems are provided in the following sections:

- [Creating a Baseline of the Database](#)
- [Diagnosing and Tuning Database Performance Problems Using ADDM](#)
- [Diagnosing and Tuning Database Performance Problems Reactively](#)
- [Tuning SQL Statements](#)

12.2.1 Creating a Baseline of the Database

Before you can tune your database, you need to have an established performance baseline that can be used for comparison if a performance problem arises. Oracle Database automatically maintains a system-defined moving window baseline that contains all AWR data within the AWR retention period, which by default is 8 days. Using Database Control, you can also create your own baseline by specifying and preserving a pair or a range of snapshots as a baseline. The snapshots contained in a baseline are excluded from the automatic AWR purging process and are retained indefinitely.

Before creating a baseline, carefully consider the time period you choose as a baseline because it should represent the database operating at an optimal level. In the future, you can compare these baselines with other baselines or snapshots to compare performance. You can create a single baseline captured at a single, fixed time interval, or a repeating baseline that repeats during a time interval over a specific period.

To create a baseline of the database, use the AWR Baselines page in Database Control, accessible from the Statistics Management section of the Server tab.

For more information about creating a baseline of the database, see the "Managing Baselines" section in the *Oracle Database 2 Day + Performance Tuning Guide*.

12.2.2 Diagnosing and Tuning Database Performance Problems Using ADDM

Automatic Database Diagnostic Monitor (ADDM) is self-diagnostic software built into Oracle Database. ADDM examines and analyzes data captured in the AWR to identify possible database performance problems.

An ADDM analysis is performed after each AWR snapshot (every hour by default), and the results are saved in the database. The results of ADDM analysis are displayed as ADDM findings under Diagnostic Summary on the Database Home page in Oracle Database Control.

Each ADDM finding belongs to one of three types: problem, symptom, and information. Each problem finding is quantified with an estimate of the portion of DB time that resulted from the performance problem. When a specific problem has multiple causes, ADDM may report multiple findings. In this case, the impacts of these multiple findings can contain the same portion of DB time.

A problem finding can be associated with a list of recommendations for reducing the impact of a performance problem. Each recommendation has a benefit that is an estimate of the portion of DB time that can be saved if the recommendation is implemented. When multiple recommendations are associated with an ADDM finding, the recommendations may contain alternatives for solving the same problem. If this is the case, then choose the easiest solution to implement that yields the greatest benefit.

Recommendations are composed of actions and rationales. You must apply all the actions of a recommendation to gain its estimated benefit. The rationales explain why the set of actions was recommended, and provide additional information for implementing them.

You should review ADDM findings and implement the recommendations as part of regular database maintenance. Even when the database is operating at an optimal performance level, you should continue to use ADDM to monitor database performance on an ongoing basis.

For more information about using ADDM, see the "Automatic Database Performance Monitoring" chapter in the *Oracle Database 2 Day + Performance Tuning Guide*.

12.2.3 Diagnosing and Tuning Database Performance Problems Reactively

While ADDM enables you to proactively diagnose database performance problems when they happen, there may be times when you want to identify database performance problems reactively. For example, you may want to compare database performance over time, or analyze a very short duration for short-lived performance problems. You may also want to examine wait events to determine if user response time can be improved by reducing the time that is spent waiting by server processes.

Oracle Database provides various tools, aside from ADDM, to enable you to diagnose and tune database performance problems reactively. The instructions are provided in the following topics:

- [Comparing Database Performance Over Time Using AWR Reports](#)
- [Identifying Transient Performance Problems Using ASH Reports](#)
- [Identifying Wait Time Using Wait Events](#)

12.2.3.1 Comparing Database Performance Over Time Using AWR Reports

Performance degradation of the database occurs when your database was performing optimally in the past, such as 6 months ago, but has gradually degraded to a point

where it becomes noticeable to the users. The Automatic Workload Repository (AWR) Compare Periods report enables you to compare database performance between two periods of time to identify any performance degradation that may have occurred over time.

The AWR Compare Periods report compares a new baseline or a pair of snapshots to an existing baseline. Before generating this report, you should have an existing baseline that represents the system operating at an optimal level, as described in ["Creating a Baseline of the Database"](#) on page 12-6. If an existing baseline is unavailable, you can use this report to compare two periods of time using two pairs of snapshots. To generate the AWR Compare Periods report, use the Automatic Workload Repository page in Database Control, accessible from the Database Server page.

For more information about generating AWR Compare Periods reports, see the "Running the AWR Compare Periods Reports" section in the *Oracle Database 2 Day + Performance Tuning Guide*.

12.2.3.2 Identifying Transient Performance Problems Using ASH Reports

ADDM tries to report the most significant performance problems during an analysis period in terms of their impact on DB time. If a particular problem lasts for a very short duration, then its severity might be averaged out or minimized by other performance problems in the analysis period. Therefore, the problem may not appear in the ADDM findings. Whether a performance problem is captured by ADDM depends on its duration compared to the interval between the AWR snapshots.

To capture a detailed history of database activity, Oracle Database samples active sessions each second with the Active Session History (ASH) sampler. ASH gathers sampled data at the session level rather than at the instance level. By capturing statistics for only active sessions, the size of the sampled data is directly related to the work being performed.

To generate the ASH report, under Average Active Sessions on the Database Performance page, click **Run ASH Report**.

For more information about generating ASH reports, see the "Running Active Session History Reports" section in the *Oracle Database 2 Day + Performance Tuning Guide*.

12.2.3.3 Identifying Wait Time Using Wait Events

Wait events are statistics that are incremented by a server process or thread to indicate that it had to wait for an event to complete before being able to continue processing. Whenever an Oracle Database process waits for something, it records the wait using one of a set of predefined wait events. Wait event data reveals various symptoms of problems that might be impacting performance, such as latch contention, buffer contention, and I/O contention.

Wait event statistics include the number of times an event was waited for and the time waited for the event to complete. To minimize user response time, reduce the time spent by server processes waiting for event completion. Not all wait events have the same wait time. Therefore, it is more important to examine events with high wait time rather than wait events with a high number of occurrences.

You can query wait event statistics from various `v$` dynamic performance views. For more information about using wait events, see the "Using Wait Event Statistics to Drill Down to Bottlenecks" section in the *Oracle Database Performance Tuning Guide*.

12.2.4 Tuning SQL Statements

When Oracle Database executes a SQL statement, the query optimizer (also called the optimizer) first determines the best and most efficient way to retrieve the results. It compares the cost of all possible approaches and chooses the approach with the least cost. The access method for physically executing a SQL statement is called an execution plan, which the optimizer is responsible for generating. The determination of an execution plan is an important step in the processing of any SQL statement, and can greatly affect execution time.

The optimizer can also help you tune SQL statements. By using SQL Tuning Advisor and SQL Access Advisor, you can run the optimizer in advisory mode to examine a SQL statement or set of statements and determine how to improve their efficiency. SQL Tuning Advisor and SQL Access Advisor can make various recommendations. SQL Access Advisor is primarily responsible for making schema modification recommendations, such as adding or dropping indexes and materialized views. SQL Tuning Advisor makes other types of recommendations, such as creating SQL profiles and restructuring SQL statements. Using Database Control, you can accept and implement many of these recommendations.

Review the following topics for information on tuning SQL statements:

- [Collecting Optimizer Statistics](#)
- [SQL Tuning Sets](#)
- [SQL Profiles](#)
- [Using SQL Tuning Advisor](#)
- [Using SQL Access Advisor](#)
- [Using SQL Tracing](#)

12.2.4.1 Collecting Optimizer Statistics

Oracle Database can execute an SQL statement in many different ways. The optimizer determines the most efficient way to execute an SQL statement after considering many factors related to the objects referenced and the conditions specified in the query. For the query optimizer to produce an optimal execution plan, the statistics in the data dictionary should accurately reflect the volume and data distribution of the tables and indexes.

Oracle Fusion Applications uses the `DBMS_STATS` package to automatically gather optimizer statistics. In this case, the database automatically collects optimizer statistics for tables with absent or stale statistics. If fresh statistics are required for a table, then the database collects them both for the table and associated indexes. For any on-demand gathering statistics on a particular table, you can use Database Control.

The output from the optimizer is an execution plan that describes an optimal method of execution. The plan shows the combination of the steps Oracle Database uses to execute an SQL statement. Each step either retrieves rows of data physically.

For more information about the optimizer, see "The Query Optimizer" chapter in the *Oracle Database Performance Tuning Guide*.

Note: Oracle Fusion Applications require several database initialization parameters to be set correctly in order to ensure optimal performance. For more information about these database initialization parameters and Fusion Applications best practices, see MOS document 1270340.1.

12.2.4.2 SQL Tuning Sets

A SQL tuning set is a database object that includes one or more SQL statements and their execution statistics and context. You can use the set as an input for advisors such as SQL Tuning Advisor, SQL Access Advisor, and SQL Performance Analyzer. You can load SQL statements into a SQL tuning set from different SQL sources, such as AWR, the cursor cache, or high-load SQL statements that you identify. SQL tuning sets are transportable, enabling SQL workloads to be transferred between databases for remote performance diagnostics and tuning. When high-load SQL statements are identified on a production system, it may not be desirable to perform investigation and tuning activities directly on this system. This feature enables you to transport the high-load SQL statements to a test system, where they can be safely analyzed and tuned.

To create a SQL tuning set, use SQL Tuning Sets page, accessible from Additional Monitoring Links on the Database Performance page.

For more information about SQL tuning sets, see the "Managing SQL Tuning Sets" in the *Oracle Database 2 Day + Performance Tuning Guide*.

12.2.4.3 SQL Profiles

A SQL profile is a set of auxiliary information that is built during automatic tuning of a SQL statement. A SQL profile is to a SQL statement what statistics are to a table.

During SQL profiling, the optimizer uses the execution history of the SQL statement to create appropriate settings for optimizer parameters. After SQL profiling completes, the optimizer uses the information in the SQL profile and regular database statistics to generate execution plans. The additional information enables the database to produce well-tuned plans for corresponding SQL statements.

After running a SQL Tuning Advisor task with a comprehensive scope, a SQL profile may be recommended. If you accept the recommendation, then the database creates the SQL profile and enables it for the SQL statement.

To manage SQL profiles, use the Plan Control tab, accessible from the SQL Details page of the SQL statement that is using a SQL profile.

For more information about SQL profiles, see the "Managing SQL Profiles" section in the *Oracle Database 2 Day + Performance Tuning Guide*.

12.2.4.4 Using SQL Tuning Advisor

Oracle Database can generate SQL tuning reports automatically. Automatic SQL Tuning runs during system maintenance windows as an automated maintenance task, searching for ways to improve the execution plans of high-load SQL statements.

ADDM also automatically identifies high-load SQL statements. If ADDM identifies such statements, then click **Schedule/Run SQL Tuning Advisor** on the Recommendation Detail page to run SQL Tuning Advisor on these statements.

For more information about SQL Tuning Advisor, see the "Tuning SQL Statements Using SQL Tuning Advisor" section in the *Oracle Database 2 Day + Performance Tuning Guide*.

12.2.4.5 Using SQL Access Advisor

To achieve optimal performance for data-intensive queries, materialized views and indexes are essential for SQL statements. However, implementing these objects does not come without cost. Creation and maintenance of these objects can be time-consuming. Space requirements can be significant. SQL Access Advisor enables

you to optimize query access paths by recommending materialized views and view logs, indexes, SQL profiles, and partitions for a specific workload.

To run SQL Access Advisor, use the **SQL Access Advisor** link on the SQL Advisors page, accessible from the Advisor Central page under Related Links on the Database Home page.

For more information about SQL Access Advisor, see the "Optimizing Data Access Paths" chapter in the *Oracle Database 2 Day + Performance Tuning Guide*.

12.2.4.6 Using SQL Tracing

You can use the SQL Trace facility to monitor Oracle Fusion Applications running against Oracle Database by assessing the efficiency of the SQL statements that Oracle Fusion applications are running. The SQL Trace facility provides performance information for individual SQL statements and generates detailed statistics for each statement.

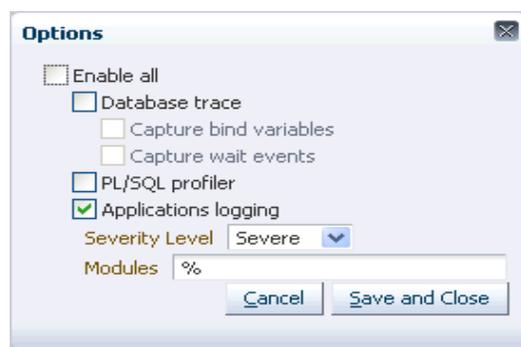
You can then run the TKPROF program to format the contents of the SQL trace file and place the output into a readable output file. TKPROF reports each statement executed with the resources it has consumed, the number of times it was called, and the number of rows which it processed. This information lets you easily locate SQL statements that are using the greatest resource.

For more information about using SQL tracing, see the "Using the SQL Trace Facility and TKPROF" section in the *Oracle Database Performance Tuning Guide*.

If a performance problem is reported for a particular business transaction for a given session in Oracle Fusion Applications, you can instruct the user of the session to enable a SQL trace for that transaction.

To enable SQL trace for a specific business transaction that is reported to be causing performance problem:

1. Sign in to Oracle Fusion Applications with a user account that is provisioned with the necessary role. Contact your security administrator for details.
2. From the **Help** menu, choose **Troubleshooting**, then **Troubleshooting Options**.
3. In the Options page, select **Database Trace**.
4. Select **Capture bind variables**, **Capture wait events** as appropriate.



5. Click **Save and Close**.

This enables SQL trace for the selected transaction and does not affect transactions that belong to other sessions.

12.3 Testing Oracle Database Changes

Oracle Real Application Testing enables you to perform real-world testing of Oracle Database. By capturing production workloads and assessing the impact of system changes on a test system before production deployment, Oracle Real Application Testing minimizes the risk of instabilities associated with changes.

See the following sections for more information about Oracle Real Application Testing:

- [Testing SQL Changes](#)
- [Testing Database Changes](#)

12.3.1 Testing SQL Changes

System changes—such as a upgrading a database or adding an index—may cause changes to execution plans of SQL statements, resulting in a significant impact on SQL performance. In some cases, the system changes may cause SQL statements to regress, resulting in performance degradation. In other cases, the system changes may improve SQL performance. Being able to accurately forecast the potential impact of system changes on SQL's performance by identifying regressed and improved SQLs, enables you to tune the system beforehand.

Also we can validate and measure the performance gain of the SQL and the system.

SQL Performance Analyzer automates the process of assessing the overall effect of a change on the full SQL workload by identifying performance divergence for each SQL statement. A report that shows the net impact on the workload performance due to the change is provided. For regressed SQL statements, SQL Performance Analyzer also provides appropriate executions plan details along with tuning recommendations. As a result, you can remedy any negative changes before the end users are affected.

To use SQL Performance Analyzer, under Real Application Testing on the Software and Support page, click **SQL Performance Analyzer**.

For more information about SQL Performance Analyzer, see the "SQL Performance Analyzer" part in the *Oracle Database Real Application Testing User's Guide*.

12.3.2 Testing Database Changes

Before system changes are made, such as hardware and software upgrades, extensive testing is usually performed in a test environment to validate the changes. However, despite the testing, the new system often experiences unexpected behavior when it enters production because the testing was not performed using a realistic workload. The inability to simulate a realistic workload during testing is one of the biggest challenges when validating system changes.

Database Replay enables realistic testing of system changes by essentially re-creating the production workload environment on a test system. Using Database Replay, you can capture a workload on the production system and replay it on a test system with the exact timing, concurrency, and transaction characteristics of the original workload. This enables you to fully assess the impact of the change, including undesired results, new contention points, or plan regressions. Extensive analysis and reporting is provided to help identify any potential problems, such as new errors encountered and performance divergence.

To use Database Replay, under Real Application Testing on the Software and Support page, click **Database Replay**.

For more information about Database Replay, see the "Database Replay" part in the *Oracle Database Real Application Testing User's Guide*

Managing Oracle Fusion Applications Log Files and Diagnostic Tests

This chapter discusses how to use incidents, log files, QuickTrace, and diagnostics tests to support normal operations for Oracle Fusion Applications and to prepare for future troubleshooting.

This chapter contains the following topics:

- [Introduction to Incidents, Log File Management, QuickTrace, and Diagnostic Tests](#)
- [Relationships Between Log Files and Diagnostic Tests](#)
- [Standard Log File Administration Tasks and Tools](#)
- [Log Files and Settings](#)
- [Log Message Structure](#)
- [Viewing and Searching Log Files During Normal Operation](#)
- [Configuring Settings for Log Files During Normal Operation](#)
- [Standard Diagnostic Testing Administration Tasks and Tools](#)
- [Configuring the Diagnostic Testing Framework for Normal Operation](#)
- [Using Diagnostic Tests to Monitor Normal System Health](#)

For information about troubleshooting using log settings, log files, the QuickTrace feature (in-memory logging), diagnostic tests, and incidents, see [Chapter 14](#).

13.1 Introduction to Incidents, Log File Management, QuickTrace, and Diagnostic Tests

Incidents, log files, QuickTrace, and diagnostic tests can all help you administer Oracle Fusion Applications. If your Oracle Fusion Applications environment includes Oracle Enterprise Manager Cloud Control (Cloud Control), then you can use Support Workbench to help you investigate, report, and resolve problems.

Incidents are collections of information about problematic error conditions. It is strongly recommended that you follow Information Technology Infrastructure Library (ITIL) best practices by establishing a help desk or service desk within your organization to support Oracle Fusion Applications, and have the help desk personnel use incidents to track the troubleshooting and resolution of all problems. Some incidents are created and gather information automatically when problems occur. For example, the information associated with an automatically created incident may include detailed operational info collected by the QuickTrace (in-memory logging) feature for Oracle Fusion Applications. For problems that do not automatically create

incidents, administrators or help desk personnel can manually create incidents and manually gather and add related system information.

Log files contain information about both normal and problematic events. Log files can help you both to monitor normal operation diagnose and to address some problems, yourself. For example, log messages that state that a service cannot be reached might indicate a hardware failure. If you discover a more complex issue, Oracle Support personnel may use log files to trace the execution code paths of relevant requests, as part of diagnosing the problem. And log files are particularly helpful if your Oracle implementation contains custom code that needs debugging, especially when using a debugger is not feasible, such as on a production system

The **QuickTrace** (in-memory logging) feature continuously records a specified level of log detail in an area of memory. The memory is recycled on an ongoing basis, with the oldest information being deleted or overwritten first. Because QuickTrace writes to memory instead of to a log file, it can gather operational information continuously without significantly affecting system performance. The information that QuickTrace stores in memory is written to disk only when an incident occurs or when an administrator manually dumps the contents of a QuickTrace buffer.

Diagnostic tests are executables that are designed to exercise particular aspects of Oracle Fusion applications, to determine whether they are operating correctly and to help identify and resolve any problems. The Diagnostic Testing Framework for Oracle Fusion Applications lets you execute diagnostic tests and collects the results into detailed diagnostic reports. Oracle provides diagnostics tests that are installed along with Oracle Fusion Applications releases and patches.

In Cloud Control, **Support Workbench** helps you investigate, report, and resolve problems (critical errors). You can use Support Workbench to perform the following kinds of operations:

- View summary information about recent problems and incidents
- View detailed diagnostic data that was gathered automatically
- Manually trigger additional dumps of diagnostic data
- Connect to the Diagnostic Dashboard for Oracle Fusion Applications, which provides you with additional diagnostic tests that you can run
- Create or update Service Requests with Oracle Support, including uploading diagnostic data to Oracle

For information about using Support Workbench to administer Oracle Fusion Applications targets, specifically, see [Chapter 14](#). For general information about using Support Workbench, see the topic "Support Workbench Page" in the Cloud Control online help and the chapter on investigating, reporting, and resolving problems in *Oracle Database 2 Day DBA*.

13.2 Relationships Between Log Files and Diagnostic Tests

The following features are designed to work together to help you administer and support Oracle Fusion Applications over time:

- **Logs and Error-Handling:** Oracle developers use mechanisms such as application programming interface (API) calls in Oracle Fusion Applications code to record application operations in log files and to provide error messages as appropriate. You can set system log levels to determine how much information is logged.

For more information about setting log levels, see [Section 13.7](#). For more information about using logs to help diagnose a problem, see [Chapter 14](#). For more information about monitoring log files, see [Section 13.6](#) and [Section 13.7.4](#).

- **Diagnostic Tests:** Oracle developers create tests that you can use to help diagnose and resolve Oracle Fusion application problems. A diagnostic test may or may not be associated with a particular error message. If an Oracle Fusion application handles a particular error in a way that triggers the creation of an incident, then any diagnostic tests that are associated with the error message for the incident run automatically. The test results are associated with the incident and the identity of the user who received the error message is recorded.

For more information about diagnostic tests, see [Section 13.8](#), [Section 13.9](#), [Section 13.10](#), and [Chapter 14](#).

It is important to be familiar with the following additional concepts that are related logs and diagnostic tests:

- **Seed Data:** Information that Oracle provides to you in the form of database records. Both error messages and diagnostic tests are included in seed data.
- **Profiles:** Settings that you can select in order to determine details of how the application operates. Oracle Fusion applications include profile options that affect how much information to log either for an entire site or a specific user.

For more information about profiles, see [Section 13.7](#) and the chapter on maintaining common reference objects in the *Oracle Fusion Applications Common Implementation Guide*.

13.3 Standard Log File Administration Tasks and Tools

Under ordinary circumstances, the following administrative tasks are part of necessary setup and maintenance of Oracle Fusion Applications log files:

- Tuning system performance by adjusting configuration settings for logs
- Reviewing log files for general monitoring of system health
- Searching for specific information in log files
- Managing available disk space for log files

If your Oracle Fusion Applications environment includes Cloud Control, then you may find it convenient to use Cloud Control to monitor log files in multiple domains, simultaneously.

You can use the Oracle Enterprise Manager Fusion Applications Control (Fusion Applications Control) user interface to complete all of the tasks listed in this section, but only for one WebLogic domain at a time. In particular, you can use Fusion Applications Control to configure log profile option values for Oracle Fusion Applications, although you can also configure profile options by using the Manage Administrator Profile Values screen in the Setup and Maintenance work area. You can use either Cloud Control or Fusion Applications Control for viewing and searching Oracle Fusion Applications log files.

Note: Fusion Applications Control and Cloud Control are compliant with Information Technology Infrastructure Library (ITIL) best practices.

During troubleshooting activities, additional log administration tasks may include downloading log files from servers and packaging incidents for transmittal to Oracle Support. For more information about these activities and the tools for accomplishing them, see [Chapter 14](#).

13.4 Log Files and Settings

Oracle provides settings that determine the amount of information that is gathered into the log files for your Oracle Fusion applications. You can either use the default setting values or change one or more values to adjust how much information is gathered. When your system is operating correctly, you may need log entries only for particularly important kinds of events. If your system experiences a problem, you can temporarily increase the amount of information that is logged to get more detailed information while you attempt to reproduce and resolve the problem.

Most Oracle Fusion applications write log output in Oracle Diagnostic Logging (ODL) format. For information about the attributes that appear in standard log files for Oracle Fusion Applications, see the "Understanding ODL Messages and ODL Log Files" section in the *Oracle Fusion Middleware Administrator's Guide*.

Some Oracle Fusion applications use nonstandard logging mechanisms that are disabled by default. For these applications, you must turn on the logging facility when you need it and specify the kind of information that you want to record in the log file. For more information about these special logging mechanisms, see [Section 13.7.6](#).

13.5 Log Message Structure

A typical log message consists of three parts:

- Attributes that are logged by the Oracle Fusion Middleware layer or the Oracle Database layer
- Attributes that are logged by logging APIs for Oracle Fusion Applications, including AppsLogger
- Message text

This section contains the following topics:

- [Log Message Attributes Supplied by the Oracle Fusion Middleware and Oracle Database Layers](#)
- [Log Message Attributes Supplied by Logging APIs for Oracle Fusion Applications](#)
- [Log Message Attributes Supplied by the Application Session](#)
- [Log Message Attributes Supplied by Oracle Enterprise Scheduler Job Requests](#)
- [Log Message Attributes Supplied by SOA](#)

13.5.1 Log Message Attributes Supplied by the Oracle Fusion Middleware and Oracle Database Layers

The log message attributes logged by the Oracle Fusion Middleware layer or the Oracle Database layer may include the following:

- **Date/Time:** The date and time when the message was recorded in the log.
- **Component ID:** The component or Oracle WebLogic Server instance from which the message originated. For the Oracle Fusion Middleware layer, a typical value is the name of the Oracle WebLogic Server that was executing the Oracle Fusion

application when the message was generated. For the database layer, a typical value is `rdbms`.

- **Message Type:Level:** Shows the level being logged for the particular message
- **Message ID:** A unique identifier for a seeded message, composed of the product code and a message number. A typical value for a Message ID might be `FND-12343`.
- **Module ID:** The system or application module that generated the message. This is usually the name of the logger object that generated the message. In some cases, the logger name may reflect the full Java class name of the application code module that was executing when the message was logged. In other cases, the logger object may generate messages for multiple Java classes.
- **Execution Context ID (ECID):** A global unique identifier and a sequence number of the thread of execution that the originating component participates in. The identifier can be used to correlate messages from several components that may be involved in the same thread of execution.
- **Host:** The host name where the message originates. For Java, this should be the value returned by `java.net.InetAddress.getLocalHost().getHostName()`.
- **Thread ID (TID):** A unique identifier for the thread within the Java process where the message was generated.
- **Java EE Application Name:** Name of a Java EE application that was executing when the message was logged.
- **User Name:** A unique identifier that the user enters when signing in to an Oracle Fusion application. This attribute is logged both by Oracle Fusion Middleware and by the logging APIs for Oracle Fusion Applications.
- **Selective Trace ID (labeled `ODL_TRACE_ID` in the log file):** An identifier for operations in the Oracle Fusion Middleware layer that match criteria an administrator supplied as part of a request to log additional information.

Note: Selective trace operations do not affect the information that is logged from the Oracle Fusion Applications layer, but the Oracle Fusion Middleware log messages that selective tracing operations obtain are stored in the same log files as log messages from the Oracle Fusion Applications layer. Similarly, log messages from the Oracle Fusion Applications layer are listed along with Oracle Fusion Middleware log messages in selective trace output if those log messages match the trace criteria that the administrator specified. For more information about selective tracing, see the "Configuring and Using Selective Tracing" section in the *Oracle Fusion Middleware Administrator's Guide*.

13.5.2 Log Message Attributes Supplied by Logging APIs for Oracle Fusion Applications

The log message attributes logged by logging APIs for Oracle Fusion Applications include the following:

- **User Name:** A unique identifier that the user enters when signing in to an Oracle Fusion application. This attribute is logged both by Oracle Fusion Middleware and by the logging APIs for Oracle Fusion Applications.
- **User GUID:** A global unique identifier representing the user.

- Role IDs: A list of IDs representing the job roles granted to the user.
- Session ID: A unique identifier for the application user session.
- Thread Name: A name that identifies the thread generating the log in JVM. Applies only to Oracle Fusion applications that are written in Java. Logging APIs for Oracle Fusion applications that are written in C and PL/SQL do not populate this attribute.
- Apps Source: The portion of the Oracle Fusion Applications code that is executing when the message is logged.
- Apps Auto Log: Indicates whether the message being logged was logged implicitly.
- DB Connection URL: The URL connection string for the application database data source.

Depending on the value of the `AFLOG_EXTENDED_ENABLED` profile option, the following supplemental attributes may also be logged:

- Message Cause: The reason that the message is being logged.
- Admin Action: Recommended follow-up action for Oracle Fusion Applications administrators.
- Admin Details: Additional information for administrators about the condition being logged.
- User Action: Recommended follow-up action for Oracle Fusion Applications users.
- User Details: Additional information for users about the condition being logged.

For more information about how the `AFLOG_EXTENDED_ENABLED` profile option affects the logging of these supplemental attributes, see [Section 13.7.1](#).

13.5.3 Log Message Attributes Supplied by the Application Session

The log message attributes logged by the application session include the following:

- Product Family: The name of the product family that is executing when the message is logged.
- Product: The name of the product that is executing when the message is logged.

13.5.4 Log Message Attributes Supplied by Oracle Enterprise Scheduler Job Requests

The following attributes appear in the log if the log entry is written from an Oracle Enterprise Scheduler job request.

- Job Request ID: The identifier of the Oracle Enterprise Scheduler job request that is being executed when this message is logged.
- Job Definition Name: The name of the Oracle Enterprise Scheduler job definition. A job definition is the smallest unit of work that is performed in the context of the client application.
- Job Package Name: The package location of the Oracle Enterprise Scheduler job definition that is executing when this message is logged.
- Job Definition Application: The application that owns the Oracle Enterprise Scheduler job.

13.5.5 Log Message Attributes Supplied by SOA

The following attributes appear in the log if the log entry is written from SOA:

- SOA Composite Name: The name of the SOA Composite that is executing when the message is logged
- SOA Component Name: The name of the SOA Component that is executing when the message is logged
- SOA Composite Instance ID: The instance ID of the SOA Composite that is executing when the message is logged
- SOA Component Instance ID: The instance ID of the SOA Component that is executing when the message is logged
- Log Token Map: one or more primary key values for the logical entity that a BPEL process is handling when the message is logged. Oracle Support may require this information for troubleshooting purposes.

13.6 Viewing and Searching Log Files During Normal Operation

To view and search Oracle Fusion application log files effectively, it is important to be familiar with the software that is available for working with log files. From time to time you may also find it helpful to know where log files are stored for various application modules.

This section contains the following topics:

- [Viewing Logs from Fusion Applications Control](#)
- [Searching for Specific Information in Log Files within a Single Domain](#)
- [Searching for Specific Information in Log Files Across Multiple Domains](#)
- [Typical Log File Locations](#)
- [Location of Logged Information When Using Multiple Managed Servers](#)

Note: Each Oracle Fusion Applications module is written in one of the following code languages: Java, SOA, PL/SQL, or C. Details of logging vary depending on the coding language of the application module.

In general, you can administer logging for Oracle Fusion Applications without knowing which programming language implements particular modules. However, it is important to become familiar with the configuration settings and log file locations used for each language, so you can monitor or adjust all relevant parts of the system when necessary.

13.6.1 Viewing Logs from Fusion Applications Control

When you want to view log files for one Oracle Fusion Applications product family or product that uses standard logging functionality, one typical way that you can view the log files is to use Fusion Applications Control.

By default, you can use Fusion Applications Control to view standard log files for the Oracle WebLogic Servers that host Oracle Fusion applications and standard log files generated by Java code within Oracle Fusion applications. In addition, you can configure your system to allow Fusion Applications Control to view log files

generated by PL/SQL or C code within Oracle Fusion applications. For more information about how to configure this capability, see [Section 13.7.7](#).

From Fusion Applications Control, you can view log messages for different system components by selecting different targets. In the product family part of the navigation tree, the following target types allow viewing of log messages for Oracle Fusion Applications:

- An entire product family. When you select this target type and view log files, you view the aggregated logs for all of the servers and clusters in the selected product family's domain.
- A particular Oracle Fusion application, such as Procurement, sometimes called an Oracle Fusion Applications cluster application. When you select this target type and view log files, you view the aggregated logs for all of the servers in the selected cluster.
- A particular instance of an Oracle Fusion application deployed to a Managed Server. When you select this target type and view log files, you view the log for the selected individual Managed Server.

The Fusion Applications Control page also gives you access to log files for the following target types in the farm part of the navigation tree, but the functionality used to view those logs is part of Fusion Middleware Control. For more information, see the "Viewing Log Files and Their Messages Using Fusion Middleware Control" section in the *Oracle Fusion Middleware Administrator's Guide*.

- A particular Oracle WebLogic Server domain
- A particular server cluster
- A particular server within a particular server cluster

Note: For information about viewing Oracle Fusion Middleware log files, including log files for Oracle Enterprise Scheduler and Oracle Enterprise Crawl and Search Framework, see the "Viewing and Searching Log Files" section in the *Oracle Fusion Middleware Administrator's Guide*.

To view the contents of Oracle Fusion Applications logs from Fusion Applications Control:

1. From the navigation pane, select the target for which you want to display log contents.

In the part of the navigation tree for the product family, you can display logs for the following kinds of targets:

- A product family target lets you display the aggregated log entries for all of the Oracle Fusion applications in the selected product family. For example, **Financials** is a product family target that appears as a top-level folder in the navigation tree.
- An Oracle Fusion Applications cluster application target lets you display the aggregated logs for all of the Managed Servers that run the selected application, including the logs for any other Oracle Fusion applications that run on those Managed Servers. For example, if you expand the product family, **Financials**, and then expand **Fusion Applications**, then the **PayablesApp** listing is a cluster application target.

- An Oracle Fusion application instance target lets you display just the log entries for the selected application on the selected sever. For example, if you expand the product family, **Financials**, and then expand **Fusion Applications**, and then expand **PayablesApp**, then the **PayablesApp (PayablesServer_1)** listing is an application instance target.

Note: When you select a target, the header of the context pane displays a dropdown menu name that depends on the target type:

- For a product family target, the menu name is **Product Family**.
 - For an Oracle Fusion Applications cluster application target, the menu name is **Fusion Cluster Application**.
 - For an Oracle Fusion application instance target, the menu name is **Fusion J2EE Application**.
-

Alternatively, if you want to view a log file for an Oracle Fusion Middleware target, you can select such a target from the farm part of the navigation tree. For more information about viewing log files for Oracle Fusion Middleware components, see the "Viewing Log Files and Their Messages Using Fusion Middleware Control" section in the *Oracle Fusion Middleware Administrator's Guide*.

2. In the context pane, from the dynamic target menu, choose **Logs > View Log Messages** to display the log entries for the target you selected.

13.6.2 Searching for Specific Information in Log Files within a Single Domain

For most Oracle Fusion applications, you can use Fusion Applications Control to search for log messages that have specific characteristics. Searches that use Fusion Applications Control are very similar to searches that use Fusion Middleware Control. For more information about these searches, see the "Viewing and Searching Log Files" section in the *Oracle Fusion Middleware Administrator's Guide*.

By default, you can use Fusion Applications Control to search standard log files for Oracle Fusion applications that are written in Java and the Oracle WebLogic Servers that host those applications. You can configure your system to allow Fusion Applications Control to search log files for Oracle Fusion applications that are written in PL/SQL or C, as well. For more information about the necessary configuration steps for PL/SQL and C logs, see [Section 13.7.7](#).

To search standard Oracle Fusion Applications log files using Fusion Applications Control:

1. From the navigation pane for Fusion Applications Control, select the target for which you want to search log contents.

For your target, you can select a product family such as **Financials**, an Oracle Fusion Applications cluster application such as **LedgerApp**, or an Oracle Fusion application instance such as **LedgerApp (LedgerServer_1)**. For information about searching log files in multiple domains, using a Fusion Instance target, see [Section 13.6.3](#).

Note: When you select a target, the header of the context pane displays a dropdown menu name that depends on the target type:

- For a product family target, the menu name is **Product Family**.
 - For an Oracle Fusion Applications cluster application target, the menu name is **Fusion Cluster Application**.
 - For an Oracle Fusion application instance target, the menu name is **Fusion J2EE Application**.
-

For example, if you want to search the aggregated log entries for all of the Oracle Fusion applications in the **Financials** product family, then you would select the **Financials** entry in the navigation tree. The **Product Family** dropdown menu appears under **Financials** in the context pane header.

If you want to search the aggregated logs for the Ledger cluster application, you would expand the **Financials** entry in the navigation tree, then expand the **Fusion Applications** entry, and then select **LedgerApp**. The **Fusion Cluster Application** dropdown menu appears under **LedgerApp** in the context pane header.

If you would rather search the Ledger application log entries for only a particular server, you would expand the **Financials** entry in the navigation tree, then expand the **Fusion Applications** entry, then expand the **LedgerApp** entry, and then select the individual **LedgerApp** Managed Server for the log you want to search. The **Fusion J2EE Application** dropdown menu appears under **LedgerApp** in the context pane header.

2. In the context pane, from the dynamic target menu, choose **Logs > View Log Messages**.

If you selected an Oracle Fusion application as your target, the content pane displays the aggregated log messages from all of the servers for the Oracle Fusion application that you selected.

If you selected an individual Oracle Fusion application instance running on a particular Managed Server as your target, the content pane displays just the log messages for the selected Oracle Fusion application operations that were handled by the selected server.

If you selected a product family such as **Financials** as your target, the content pane displays the aggregated log messages for all of the Managed Servers that handle Oracle Fusion applications in the selected product family.

3. If necessary, expand the **Search** area of the content pane.
4. If you selected an Oracle Fusion application instance in Step 1, skip to Step 5.

If you selected a product family or Oracle Fusion Applications cluster application in Step 1, complete the following substeps to determine which members of the selected target will be included in the log search:

- a. Expand the **Selected Targets** section of the **Search** area.
- b. If you want to search the log files for only some of the listed target members, select the ones you want to omit from the search and click **Remove**.
- c. If you want the log search operation to search log files that are not currently listed, click **Add**, select the additional targets that you want to search, and click **Select**.
- d. Proceed to Step 5.

5. Enter or select applicable search criteria in any visible search fields, such as the following fields that are visible by default:
 - In the **Date Range** fields, specify the time period for which you want to display log messages.

To specify a specific number of the most recent minutes, hours, or days, select **Most Recent** from the dropdown list, enter your preferred number, and select the appropriate time units.

To specify the time period using starting and ending dates and times, select **Time Interval** from the dropdown list and enter the starting and ending dates and times.
 - In the **Message Types** field, select one or more checkboxes to specify the kinds of log messages that you want to display.
 - In the **Message** fields specify some or all of the message text in the log messages that you want to display.

If you select either the **matches** search operator or the **does not match** search operator, you must either specify the full text of the message, or you must use a standard Java language expression to specify the characteristics of the missing text. For example, in standard Java expressions, the period character (.) acts as a wildcard that matches most single characters, and the combination of period and asterisk characters (.*) acts as a wildcard to match any string of characters.

In contrast, the following search operators contain implicit wildcard semantics, so you do not need wildcard characters or other Java expressions when you use any of these operators to search for a part of the text of a message:

- **contains**
- **does not contain**
- **starts with**
- **does not start with**
- **ends with**
- **does not end with**

Note: If you search for multiple values within a single search field, using commas to separate the values, then the search results display log entries that match any of the criteria specified for that field (logical OR). However, if you specify criteria in multiple search fields, the search results display only log entries that match specified criteria in all of those search fields (logical AND).

6. If you want to use search criteria in fields that are not currently visible, complete the following substeps; otherwise, skip this step.
 - a. Click **Add Fields**.
 - b. Select one or more checkboxes for the additional fields you want to use in your search, and then click **Add**.

For example, you might select **ECID** to search for all log messages that are associated with a particular execution context identifier.

- c. In each field that you added, specify the value that you want the search to match.

For example, you might enter the **ECID** value of
004bYSyedEi3v1A5Jb8Dyf0002kx003Xff.

7. Click **Search**.

13.6.3 Searching for Specific Information in Log Files Across Multiple Domains

If your Oracle Fusion Applications environment includes Cloud Control, you can use Cloud Control to search for information in the log files of multiple domains within a single Fusion instance.

To search standard Oracle Fusion Applications log files in multiple domains using Cloud Control:

1. In Cloud Control, from the **Target** menu, choose **Fusion Applications**.
2. In the list of Oracle Fusion instances and Oracle Fusion Applications product families, click the name of the instance that interests you.
3. From the dynamic Fusion Instance target menu, choose **Logs > View Log Messages**
4. If it is collapsed, expand the **Search** area of the **Log Messages** pane.
5. Expand the **Selected Targets** part of the **Search** area, and make sure that it lists the targets for all of the log files that you want to search:
 - If you want to search the log files for only some of the listed target members, select the ones you want to omit from the search and click **Remove**.
 - If you want the log search operation to search log files that are not currently listed, click **Add**, select the additional targets that you want to search, and click **Select**.
6. Enter or select applicable search criteria in any visible search fields, such as the following fields that are visible by default:
 - In the **Date Range** fields, specify the time period for which you want to display log messages.

To specify a specific number of the most recent minutes, hours, or days, select **Most Recent** from the dropdown list, enter your preferred number, and select the appropriate time units.

To specify the time period using starting and ending dates and times, select **Time Interval** from the dropdown list and enter the starting and ending dates and times.
 - In the **Message Types** field, select one or more checkboxes to specify the kinds of log messages that you want to display.
 - In the **Message** fields specify some or all of the message text in the log messages that you want to display.

If you select either the **matches** search operator or the **does not match** search operator, you must either specify the full text of the message, or you must use a standard Java language expression to specify the characteristics of the missing text. For example, in standard Java expressions, the period character (.) acts as a wildcard that matches most single characters, and the combination of period and asterisk characters (.*) acts as a wildcard to match any string of characters.

In contrast, the following search operators contain implicit wildcard semantics, so you do not need wildcard characters or other Java expressions when you use any of these operators to search for a part of the text of a message:

- **contains**
- **does not contain**
- **starts with**
- **does not start with**
- **ends with**
- **does not end with**

Note: If you search for multiple values within a single search field, using commas to separate the values, then the search results display log entries that match any of the criteria specified for that field (logical OR). However, if you specify criteria in multiple search fields, the search results display only log entries that match specified criteria in all of those search fields (logical AND).

7. If you want to use search criteria in fields that are not currently visible, complete the following substeps; otherwise, skip this step.

a. Click **Add Fields**.

b. Select one or more checkboxes for the additional fields you want to use in your search, and then click **Add**.

For example, you might select **ECID** to search for all log messages that are associated with a particular execution context identifier.

c. In each field that you added, specify the value that you want the search to match.

For example, you might enter the **ECID** value of
004bYSyedEi3v1A5Jb8Dyf0002kx003XFf.

8. Click **Search**.

13.6.4 Typical Log File Locations

Note: The following information is provided for reference purposes. For maximum clarity and convenience, it is recommended that you use Fusion Applications Control to work with Oracle Fusion Applications log files, rather than working directly with the log files.

Each Oracle Fusion application runs in one or more logical Managed Servers. Each Managed Server is dedicated to a single Oracle Fusion application. Different Managed Servers do not normally share log files with each other, even if they are running on the same physical server computer.

Most Oracle Fusion Applications code modules use standard logging code. On each Managed Server, all Java and SOA application code modules that use standard logging code write log entries to a single file. The location of this file is specified by a `log_`

handler entry in a logging.xml file. The default location for the logging.xml file is as follows, where *DOMAIN_HOME* is the path to your Oracle WebLogic Server domain, and *WebLogic_Server_Name* is the name of the WebLogic Server that uses the logging.xml file:

```
DOMAIN_HOME/config/fmwconfig/servers/WebLogic_Server_Name/logging.xml
```

[Example 13–1](#) shows a typical log_handler entry in the logging.xml file, which includes the path and file name of the log file.

Example 13–1 Typical log_handler Entry Showing Log File Name and Location

```
<log_handlers>
  <log_handler name='apps-handler'
class='oracle.core.ojdl.logging.ODLHandlerFactory'>
  <property name='path'
value='${domain.home}/servers/${weblogic.Name}/logs/apps/${weblogic.Name}-
diagnostic.log' />
  <property name='maxFileSize' value='10485760' />
  <property name='maxLogSize' value='104857600' />
  <property name='encoding' value='UTF-8' />
  <property name='supplementalAttributes' value='APPS_USER_NAME,
APPS_SESSION_ID, APPS_THREAD_NAME, APPS_SOURCE, APPS_USER_ID, APPS_AUTO_LOG,
APPS_JOB_REQUEST_ID, APPS_JOB_DEFINITION_NAME, APPS_JOB_PACKAGE_NAME,
APPS_JOB_DEFINITION_APP, APPS_COMPOSITE_NAME, APPS_COMPONENT_NAME,
APPS_COMPOSITE_INSTANCE_ID, APPS_COMPONENT_INSTANCE_ID, APPS_PRODUCT_FAMILY,
APPS_PRODUCT, APPS_INDUSTRY, APPS_TERRITORY, APPS_DB_CONNECTION_URL,
APPS_MESSAGE_CAUSE, APPS_USER_ACTION, APPS_ADMIN_ACTION, APPS_USER_DETAILS,
APPS_ADMIN_DETAILS, APPS_LOG_TOKEN_MAP' />
  </log_handler>
</log_handlers>
```

[Example 13–2](#) shows the default location for the standard log file is as follows, for Oracle Fusion applications that are implemented in Java and SOA, where *DOMAIN_HOME* is the path to your Oracle WebLogic Server domain:

Example 13–2 Default Log File Location for Oracle Fusion Applications Implemented in Java and SOA

```
DOMAIN_HOME/servers/domain_name/logs/apps/server_name-diagnostic.log
```

For Oracle Fusion Applications code modules that are implemented in PL/SQL, the location of the standard log file is set by the *AFLOG_PLSQL_FILENAME* profile option. By default, the value of this profile option is *APPLLOG_DIR*/diagnostic.log, where *APPLLOG_DIR* is a directory object that was defined as a custom variable through use of the Oracle Fusion Applications Repository Creation Utility during installation. For more information about directory objects, see the section about the *Create Directory* command in the *Oracle Database SQL Language Reference*

For Oracle Fusion Applications code modules that are implemented in C, the location of the standard log file is set by the *AFLOG_FILENAME* profile option. By default, the value of this profile option is *diagnostic.log*, but it is recommended that you set the value to *directory_path*/Cdiagnostic.log, where *directory_path* is a location that can be written to by all system users and by the Managed Server where Oracle Enterprise Scheduler is deployed for your Oracle Fusion applications. For example, you might set the value of *directory_path* to */tmp*, or to an explicitly specified directory path that corresponds to the value of the *APPLCP_DIR* custom variable that was defined through use of the Oracle Fusion Applications Repository Creation Utility during installation.

By default, the following Oracle Fusion Applications code modules do not write log entries. If you configure these modules to write log entries, then the entries for those modules are written to nonstandard log file locations. These modules do not specify any configuration settings using the standard `logging.xml` file:

- The following kinds of batch jobs for Oracle Fusion Incentive Compensation:
 - Calculation
 - Classification
 - Collection
 - Crediting
 - Rollup

Note: For optimum performance, it is recommended that you use the logging functionality that is available for these kinds of Oracle Fusion Incentive Compensation batch jobs only when troubleshooting an existing problem. For more information about using this feature and viewing the results, see [Section 14.2.9](#) and [Section 14.5.2](#).

- The following kinds of batch jobs for Oracle Fusion General Ledger:
 - OpenPeriod
 - Posting
 - Translation
 - Close Process - Create Income Statement Closing Journals
 - Close Process - Create Balance Sheet Closing Journals

Note: For optimum performance and log file sizes, it is recommended that you use the logging functionality for these kinds of Oracle Fusion General Ledger batch jobs only when troubleshooting an existing problem. For more information about using this feature and viewing the results, see [Section 14.2.10](#).

- The AutoInvoice portion of the Oracle Fusion Receivables application

Note: The amount of information that is logged for AutoInvoice depends on the value of the Log File Message Level system option setting for each business unit. Any AutoInvoice log messages that meet the level requirements are written to the standard log file for Oracle Enterprise Scheduler.

For more information about configuring the Log File Message Level system option setting, see [Section 13.7.6.3](#). For more information about adjusting this setting for troubleshooting and viewing the results, see [Section 14.2.11](#).

13.6.5 Location of Logged Information When Using Multiple Managed Servers

Note: The following information is provided for reference purposes. For maximum clarity and convenience, it is recommended that you use Fusion Applications Control to work with Oracle Fusion Applications log files, rather than working directly with the log files.

If an Oracle Fusion application is deployed to multiple Managed Servers in a cluster domain, each Managed Server records log entries in its own log file, for just the transactions or other operations handled by that server.

For your convenience in monitoring Oracle Fusion applications that run on multiple Managed Servers, Fusion Applications Control lets you view either individual standard log files for individual Managed Servers or the aggregated contents of the standard log files of all of the servers in the cluster domain. The aggregated log is especially useful if you need to find a particular log entry but do not know which Managed Server recorded it. For more information about viewing log files, see [Section 13.6.1](#).

However, Fusion Applications Control does not currently support viewing the contents of any nonstandard log files generated by Oracle Fusion Incentive Compensation, Oracle Fusion General Ledger, or the AutoInvoice portion of the Oracle Fusion Receivables application. For information about viewing nonstandard log files for Oracle Fusion Incentive Compensation, see [Section 14.5.2](#). For information about viewing nonstandard log files for Oracle Fusion General Ledger, see [Section 14.2.10](#). Log entries for the AutoInvoice portion of Oracle Fusion Receivables are placed in the standard log file for Oracle Enterprise Scheduler. For information about viewing Oracle Fusion Middleware log files, including log files for Oracle Enterprise Scheduler, see the "Viewing and Searching Log Files" section in the *Oracle Fusion Middleware Administrator's Guide*.

Note: When an Oracle Fusion application is deployed to multiple Managed Servers in a cluster domain, it is generally recommended that you set the log level to the same value for all of the Managed Servers, so that all of the servers will log comparable amounts and kinds of information. However, you can configure each Managed Server's log level independently, if you have a reason to do so.

13.7 Configuring Settings for Log Files During Normal Operation

Although critical business logic sections of Oracle Fusion applications may write more information to log files than less critical areas of the application code, the amount of information that Oracle Fusion applications log depends primarily on how the environment is configured. Oracle supplies default values for log settings, but you can specify different setting values if you want to adjust the amount of information to be logged. Most Oracle Fusion Applications components use a standard set of log configuration settings. For information about the selected components that use nonstandard log settings, see [Section 13.7.6](#).

This section contains the following topics:

- [Default System Log Settings](#)
- [When Changes to Log File Configuration Settings Take Effect](#)
- [Standard Logging Levels](#)

- [Managing Log File Size and Disk Space Usage](#)
- [Using Profile Options to Configure Standard Log Settings](#)
- [Using Additional Settings to Configure Logs for Selected Components](#)
- [Configuring Access to Logs for Fusion Applications Control](#)

The concepts of logging levels and profile options apply to most Oracle Fusion applications. Logging levels are thresholds that can be set to control how much information to log.

To set logging levels either for a whole site or for a particular user, administrators change profile option values from either the Log Configuration dialog box in Fusion Applications Control or the Manage Administrator Profile Values screen in the Setup and Maintenance work area. For more information, see the chapter on maintaining common reference objects in the *Oracle Fusion Applications Common Implementation Guide*.

Under certain specific conditions, Oracle Fusion Applications users can set their own profile option for logging levels, from the Troubleshooting Options dialog box under the Oracle Fusion Applications **Help > Troubleshooting** menu. Setting a logging level for a particular user is useful if you want only that user to gather more detailed log information while attempting to reproduce a problem. However, the **Help > Troubleshooting** menu displays the **Troubleshooting Options** command only for Oracle Fusion Applications users who have a job role that is mapped to the following three duty roles:

- Supportability Level Management Duty (CRM) (FND_SET_SUPPORTABILITY_LEVEL_DUTY_CRM)
- Supportability Level Management Duty (FSCM) (FND_SUPPORTABILITY_LEVEL_MANAGEMENT_DUTY_FSCM)
- Supportability Level Management Duty (HCM) (FND_SUPPORTABILITY_LEVEL_MANAGEMENT_DUTY_HCM)

Note: Oracle Fusion Applications seed data does not include a preconfigured job role that maps to these three duty roles. It is recommended that you define such a job role when you first need to give an Oracle Fusion Applications user access to the Troubleshooting Options dialog box. You can reuse the same job role for other users who subsequently need the same access.

For more information about making the **Troubleshooting Options** command and dialog box available to selected Oracle Fusion Applications users, and about working with those users to gather troubleshooting data, see [Section 14.2.1](#).

The following types of settings affect how logging is done, including the effective logging level:

- `AFLOG_SettingName` profile option values. For information about the `AFLOG_SettingName` profile options, see [Section 13.7.1](#) through [Section 13.7.5](#). For general information about working with profile options at both the `SITE` and `USER` levels, see the "Setting and Accessing Profile Values" and "Managing Profile Definitions" sections in the *Oracle Fusion Applications Developer's Guide*.

Note: For `AFLOG_SettingName` profile options, the `PRODUCT` level is reserved for future use.

- **Oracle Diagnostics Logging Level:** The log level for Oracle Fusion applications that are written in Java or SOA. The value for this setting is specified in each Managed Server's `logging.xml` file, in the `oracle.apps` entry. Log levels for most entries are adjusted using the Managed Server's **Log Configuration** page in Fusion Middleware Control, rather than by directly editing the `logging.xml` file. However, the log level for the `oracle.apps` entry should remain set to the default value of `All` unless Oracle Support gives you specific instructions to change it. For more information about this setting, see "Setting the Level of Information Written to Log Files" in the *Oracle Fusion Middleware Administrator's Guide*.

Caution: The logging framework for Oracle Fusion Applications is designed to work with the Log Level for the `oracle.apps` logger set to the default value of `All`. If you use other values for this setting, the diagnostic framework may not appear to work correctly, because setting values other than `All` can interfere with normal logging operations.

13.7.1 Default System Log Settings

The seed data for Oracle Fusion Applications contains default values for many profile option settings. Normal operation settings for log files are described in [Table 13–1](#). For information about default system settings for incidents and QuickTrace, see [Section 14.2](#).

Table 13–1 Profile Options for Oracle Fusion Applications Logging

Profile Option Name (and Display Name)	Environment	Description	Possible Values or Example	Applicable Profile Hierarchy Levels	Default Value
AFLOG_BUFFER_MODE (FND: Buffer Mode for PL/SQL)	PL/SQL only	Asynchronous buffer mode for PL/SQL logging. A value of 0 disables buffering; a value greater than 0 enables buffering for messages that have levels lower than WARNING. PL/SQL log messages that have SEVERE level or WARNING level are not buffered. You may want to use asynchronous buffering if you have set your log levels to collect very detailed information and find that your system is running slowly. When asynchronous buffering is enabled, log messages from PL/SQL code that are at the INFO log level or below are written to a buffer rather than being immediately written to a log file. This reduces disk I/O activity. The buffer is written to the log file only when it contains the number of records specified in the AFLOG_BUFFER_SIZE setting.	0, 1	Site, User	0
AFLOG_BUFFER_SIZE (FND: Buffer Size for PL/SQL)	PL/SQL only	Number of PL/SQL log records that are buffered in memory before they are written to the log file.	1024	Site, User	1000
AFLOG_ENABLED (FND: Log Enabled)	Java, PL/SQL, C, and SOA	Enables or disables standard logging functionality for Oracle Fusion Applications. If the value of this profile option is N, and if the profile option sets the effective logging level, then standard Oracle Fusion Applications logging does not occur at runtime for application modules written in Java, PL/SQL, SOA, or C. This profile option does not affect the configuration for Oracle Fusion Middleware.	Y, N	Site, User	Y

Table 13–1 (Cont.) Profile Options for Oracle Fusion Applications Logging

Profile Option Name (and Display Name)	Environment	Description	Possible Values or Example	Applicable Profile Hierarchy Levels	Default Value
AFLOG_EXTENDED_ENABLED (FND: Log Extension Enabled)	Java and SOA	<p>Determines whether or not to log the following extended attributes:</p> <ul style="list-style-type: none"> ▪ Message Cause (APPS_MESSAGE_CAUSE) ▪ User Action (APPS_USER_ACTION) ▪ User Details (APPS_USER_DETAILS) ▪ Admin Action (APPS_ADMIN_ACTION) ▪ Admin Details (APPS_ADMIN_DETAILS) <p>When the value of this profile option is Y, all of these extended attributes are logged. When the value of this profile option is N, none of these extended attributes are logged.</p>	Y, N	Site	Y
AFLOG_FILENAME (FND: Log File for C)	C only	<p>Full path and file name of the log file for all of the Oracle Fusion applications that are written in C and that use standard logging functionality.</p> <p>It is recommended that you set the value of this profile option to <i>directory_path/Cdiagnostic.log</i>, where <i>directory_path</i> is a location that can be written to by all system users and by the Managed Server where Oracle Enterprise Scheduler is deployed for your Oracle Fusion applications. For example, you might set the value of <i>directory_path</i> to <i>/tmp</i>, or to an explicitly specified directory path that corresponds to the value of the <i>APPLCP_DIR</i> custom variable that was defined through use of the Oracle Fusion Applications Repository Creation Utility during installation.</p> <p>If you do not set a directory path for this log file, the file is written to the default location specified by Oracle Enterprise Scheduler.</p>	/temp/Cdiagnostic.log	Site, User	diagnostic.log

Table 13–1 (Cont.) Profile Options for Oracle Fusion Applications Logging

Profile Option Name (and Display Name)	Environment	Description	Possible Values or Example	Applicable Profile Hierarchy Levels	Default Value
AFLOG_LEVEL (FND: Log Level)	Java, PL/SQL, C, and SOA	Minimum level of information detail to be logged for Oracle Fusion applications that use standard logging functionality. If no value is set for this profile option, the default value is 1000 (SEVERE). For more information about log levels, see Section 13.7.3 .	1000, 900, 800, 700, 500, 400, 300	Site, User	1000
AFLOG_MAX_FILE_SIZE (FND: Maximum size for log file in MB.)	PL/SQL	Specifies the size in megabytes beyond which the current standard log file for Oracle Fusion applications that are written in PL/SQL is automatically renamed and a new log file is started.	10	Site	10
AFLOG_MODULE (FND: Log Module Filter)	Java, PL/SQL, C, and SOA	Specifies the Oracle Fusion applications for which logging takes place. Use a comma-separated list of modules for value of this setting, and use % as a wild card character. If no value is specified, all modules are logged.	%, %financial%	Site, User	%

Table 13–1 (Cont.) Profile Options for Oracle Fusion Applications Logging

Profile Option Name (and Display Name)	Environment	Description	Possible Values or Example	Applicable Profile Hierarchy Levels	Default Value
AFLOG_NUMBER_OF_LOG_FILES (FND: Number of old log files)	PL/SQL	The maximum number of PL/SQL log files the system keeps at any one time.	Any integer greater than zero.)	Site	10
AFLOG_PLSQL_FILENAME (FND: Log File for PL/SQL)	PL/SQL only	<p>The location and name of the log file for standard Oracle Fusion Applications log messages that are generated from PL/SQL. The location must be expressed as a directory object.</p> <p>By default, the <code>SITE</code> level value for this profile option is set to <code>APPLLOG_DIR/ diagnostic.log</code>, where <code>APPLLOG_DIR</code> is a directory object that was defined using the Oracle Fusion Applications Repository Creation Utility during installation.</p> <p>If you would like the log file name to indicate that the log messages are generated from PL/SQL, you can change the <code>AFLOG_PLSQL_FILENAME</code> profile option value to a value such as <code>APPLLOG_DIR/plsqldiagnostic.log</code>.</p> <p>For more information about directory objects, see the section about the <code>Create Directory</code> command in the <i>Oracle Database SQL Language Reference</i>.</p>	<code>APPLLOG_DIR/mylog.log</code>	Site, User	<code>APPLLOG_DIR/diagnostic.log</code>
FND_MESSAGE_MODE (Message Mode)	Java, PL/SQL, C, and SOA	For error conditions that use messages from the message dictionary, this setting determines whether administrator-level or user-level message details and suggested actions are displayed and logged.	ADMIN, USER	Site, Product, User	User

13.7.2 When Changes to Log File Configuration Settings Take Effect

For PL/SQL and C processes such as scheduled jobs, changes to log file profile options take effect in the same ways as changes to any Oracle Fusion Applications profile option values. For more information, see the chapter on maintaining common reference objects in the *Oracle Fusion Applications Common Implementation Guide*.

For user sessions, users may need to log out from an Oracle Fusion application and log in again in order to have changes to log profile options take effect.

13.7.3 Standard Logging Levels

Seven severity levels are used for log messages in Oracle Fusion Applications, and each level is associated with a number. In circumstances where the effective logging level depends on the `AFLOG_LEVEL` logging profile option (rather than a lower `odlLevel` setting value), then, once the profile option level is set, only application messages that have a predefined severity level greater than or equal to the value of the `AFLOG_LEVEL` profile option are logged.

You can set the value of the `AFLOG_LEVEL` profile option to one value for the site as a whole, and to another value for any user whose operations you want to log at a different level of detail.

For example, setting the level to the lowest severity, 300 (`FINEST`), for a particular user means that messages of all seven severities are logged for that user's operations. Setting the level to 900 (`WARNING`) or the site means that logging occurs for 900 (`WARNING`) and 1000 (`SEVERE`) messages for all site operations initiated by other users.

Note: The default severity level for Oracle Fusion Applications is 1000 (`SEVERE`). For optimum performance, it is recommended that you use this logging level for your sites unless you need to investigate a problem that specifically requires a change to a site's severity level. Gathering detailed log information for an entire site (rather than a single user) can decrease system performance and make it difficult to find relevant information in a log file.

Table 13–2 describes the seven severity levels that are used for log messages in Oracle Fusion Applications. The messages in the log file identify the severity of errors using the ODL Message Type/Level value.

Note: Oracle Fusion applications that are written in PL/SQL have logging level names that start with `LEVEL_`. Oracle Fusion applications that are written in Java, SOA, and C do not use this `LEVEL_` prefix for level names. For example, the log level that is called `LEVEL_SEVERE` in PL/SQL application code is comparable to the log level that is called `SEVERE` in Java, SOA, or C code.

Most log levels provide information that is useful to Oracle Fusion Applications administrators. However, the information that is logged at the `FINER` and `FINEST` log levels is primarily intended for use by Oracle.

Table 13–2 Severity Levels for Logging Messages

ODL Message Type/Level (in Log Files)	Log Level Name	AFLOG_LEVEL Profile Option Value (internal values stored in database tables)	Usage and Examples
<i>Log Levels for Reporting Failures and Normal Events:</i>			
ERROR:1	SEVERE	1000	<p>Highest severity level. Unexpected errors that occur during normal execution. Fatal exceptions or any other serious problems that require immediate attention from the System Administrator. An error at this level may also create an incident. For example, the following error messages might be associated with this log level:</p> <p>The process cannot be started now because another process is using the data.</p> <p>The participant's rate could not be found. Contact your benefits administrator to validate the standard rate for this compensation object.</p>
WARNING:1	WARNING	900	<p>Internal software failures. Non-fatal exceptions or errors that allow processing to continue without requiring immediate attention from an administrator. Any potential problem that should be reviewed by the System Administrator.</p> <p>For example, the following error message might be associated with this log level:</p> <p>You must change your password before date.</p>
NOTIFICATION:1	INFO	800	<p>Errors, warnings and other kinds of information that allow processing to continue. May include key flow steps, high-level functional progress messages, and major life cycle events such as the activation or deactivation of a primary sub-component or feature.</p>
<i>Log Levels for Tracing and Reporting Progress:</i>			
NOTIFICATION:16	CONFIG	700	<p>Configuration properties and environment settings. This is a finer level of granularity for reporting normal events.</p>

Table 13–2 (Cont.) Severity Levels for Logging Messages

ODL Message Type/Level (in Log Files)	Log Level Name (in Oracle Enterprise Manager Fusion Applications Control and Help > Troubleshooting > Troubleshooting Options)	AFLOG_LEVEL Profile Option Value (internal values stored in database tables)	Usage and Examples
TRACE:1	FINE	500	<p>High-level technical progress messages; more detailed than INFO level. Trace or debug information for events that are meaningful to users of the product, such as public application programming interface (API) entry/exit points.</p> <p>For example, the following error message might be associated with this log level:</p> <p>User has chosen to Cancel the Purge operation.</p>
TRACE:16	FINER	400	<p>Logging messages that are called at the beginning and end of a routine. Detailed trace or debug information that can help Oracle Support diagnose problems with a particular subsystem.</p> <p>For example, the following error message might be associated with this log level:</p> <p>Entering validateLogin.</p>
TRACE:32	FINEST	300	<p>Very detailed trace or debug information that would be useful to an Oracle developer who is working on the product and who is familiar with implementation details of the sub-system that generates the message.</p> <p>For example, the following error message might be associated with this log level:</p> <p>Copying string from buffer xyz to buffer zzz.</p>

Note: Whenever an incident is created, a different log level, `INCIDENT_ERROR`, automatically appears in the corresponding log entry. This is the only context in which the `INCIDENT_ERROR` log level is used. For more information about working with incidents, see [Chapter 14](#).

13.7.4 Managing Log File Size and Disk Space Usage

In busy computing environments, the amount of disk space used by log files can become a concern. Very large log files can also affect system performance. Oracle Fusion applications that are written in Java, SOA, or PL/SQL address this concern using automatic log file rotation.

This section contains the following topics:

- [Managing Rotating Log File Space Usage for Java and SOA Applications](#)
- [Managing Rotating Log File Space Usage for PL/SQL Applications](#)
- [Managing Log File Space Usage for C Applications](#)

Note: Oracle Fusion applications that are written in C require you to monitor the space used and manually discard log files that are no longer needed.

13.7.4.1 Managing Rotating Log File Space Usage for Java and SOA Applications

For Oracle Fusion Applications modules that are implemented using Java or SOA and that use standard logging, when a the log file reaches a specific size or when a specific time period has passed, the file is automatically renamed, and a new log file is created.

You can use any of the following methods to adjust the settings that determine the maximum log file size and the maximum length of time that a log file covers:

- Use Fusion Middleware Control commands to change the log rotation policies for the `oracle.apps` logger. For more information, see "Specifying Log File Rotation Using Fusion Middleware Control" in the *Oracle Fusion Middleware Administrator's Guide*.
- Edit the settings in the `log_handler` section of the ODL configuration file, at the following location:

`DOMAIN_HOME/config/fmwconfig/servers/WebLogic_Server_Name/logging.xml`

where `DOMAIN_HOME` is the domain home directory for the Oracle WebLogic Server domain, and `WebLogic_Server_Name` is the name of the WebLogic Server that uses the `logging.xml` file.

- Use WebLogic Scripting Tool (WLST) Log Configuration commands to change the log rotation policies for the `apps-handler` log handler. For more information, see the section on the `configureLogHandler` command in the chapter on logging commands in the *Oracle Fusion Middleware WebLogic Scripting Tool Command Reference* and "Specifying Log File Rotation Using WLST" in the *Oracle Fusion Middleware Administrator's Guide*.

The `maxFileSize` or Maximum Log File Size setting determines the maximum size that a standard Java or SOA log file can reach before being renamed. The default value is 10485760 bytes.

The `rotationFrequency` or Frequency setting determines the maximum amount of time that can pass before a standard Java or SOA log file is renamed. There is no default value for this setting. If no value is specified, then the log file is not renamed after any particular time period, but only when it reaches its maximum allowed size. Valid values for the `rotationFrequency` setting are numbers representing the length of the time period in minutes, as well as the following case-insensitive values:

- hourly
- daily
- weekly

When a standard Java or SOA log file is renamed, the new name is of the format `log_file_name-n.log`, where `n` is a positive integer and `log_file_name` is the file name specified for the path property of the `apps-handler` log handler in the `logging.xml` file. (By default, `log_file_name` is set to `server_name-diagnostic.log`).

In `log_file_name-n.log`, the value of `n` depends on the names of the log files that are already present in the directory. If the directory contains no previously renamed log files, then the first renamed log file is called `log_file_name-1.log`. If other log files do exist, then `n` is set to the next higher integer after the highest integer that is already in

use. For example, if the directory contains *log_file_name-1.log* through *log_file_name-8.log* at the time when *log_file_name.log* reaches a time or size limit, then the *log_file_name.log* file is renamed to *log_file_name-9.log*

When the aggregated size of current and older log files reaches a specific value, older log files are deleted automatically, to keep the disk space usage of the log file directory below that specific value.

The `maxLogSize` or Maximum Size of All Log Files setting determines the directory size at which older log files begin to be deleted. The default value is 104857600 bytes.

13.7.4.2 Managing Rotating Log File Space Usage for PL/SQL Applications

For Oracle Fusion Applications modules that are implemented using PL/SQL, when a `diagnostic.log` file reaches a specific size, the `diagnostic.log` file is automatically renamed, and a new `diagnostic.log` file is created. If the `AFLOG_PLSQL_FILENAME` profile option is set so that the logging framework uses a log file name other than `diagnostic.log`, then the file name that the profile option specifies is used, instead of `diagnostic.log`.

Use the following profile options settings to specify the maximum log file size:

- `AFLOG_MAX_FILE_SIZE`: This setting specifies the size in megabytes beyond which a PL/SQL log file name is automatically renamed and a new log file is started. The default value is 10 megabytes.

Note: If the `AFLOG_BUFFER_MODE` profile option is set to a value larger than 0, enabling asynchronous buffering of PL/SQL log entries, then the actual maximum size of any single PL/SQL log file is the value of `AFLOG_MAX_FILE_SIZE` plus the number of megabytes that are flushed from the buffer. This value is always approximate, because the amount of information that can accumulate in the buffer is set using the `AFLOG_BUFFER_SIZE` setting, which specifies a specific number of log records, rather than a specific number of megabytes.

- `AFLOG_NUMBER_OF_LOG_FILES`: This setting specifies the maximum number of PL/SQL log files the system keeps at any one time. The default value is 10 files.

PL/SQL log rotation is currently done only on the basis of file size, not on the basis of the passage of a specified amount of time.

When a PL/SQL log file is renamed, the new name depends on whether or not the `AFLOG_PLSQL_FILENAME` profile option is set:

- If the profile option is set, then the new log file name is of the format `AFLOG_PLSQL_FILENAME_value-n.log`, where *n* is a positive integer.
- If the profile option is not set, then the new log file name is of the format `diagnostic-n.log`, where *n* is a positive integer.

The value of *n* depends on the names of the log files that are already present in the directory. If the directory contains no previously renamed log files, then the first renamed log file is called `diagnostic-1.log` or `AFLOG_PLSQL_FILENAME_value-1.log`. If other log files do exist, then *n* is set to the next higher integer after the highest integer that is already in use. For example, if the directory contains `diagnostic-1.log` through `diagnostic-8.log` at the time when the `diagnostic.log` file surpasses the size limit set in the `AFLOG_MAX_FILE_SIZE` profile option, then the `diagnostic.log` file is renamed to `diagnostic-9.log`.

When the number of log files reaches the value specified using the `AFLOG_NUMBER_OF_LOG_FILES` profile option, then older log files are deleted automatically, to prevent the disk space usage of the log file directory from growing too large.

Over time, the value of *n* in `diagnostic-n.log` or `AFLOG_PLSQL_FILENAME_value-n.log` can grow large enough to cause usability challenges or exceed the number of characters that the operating system allows in a file name. If you want to have the value of *n* start over at 1, you can move all existing log files except the currently active `diagnostic.log` file or `AFLOG_PLSQL_FILENAME_value.log` file into another directory. When the active file surpasses the size limit and the log rotation code finds no previously renamed log files in the directory, the active file is renamed using a value of 1 for *n*.

Note: If your Oracle Fusion Applications environment includes multiple database nodes such as RAC systems, each database node corresponds to a server instance that has its own location for log files.

If an incident is created, the server instance that creates the incident handles all subsequent jobs related to that incident. Identifiers for incidents are unique within a specific instance, but not across instances. For more information about working with incidents, see [Chapter 14](#).

13.7.4.3 Managing Log File Space Usage for C Applications

Oracle Fusion Applications modules that are implemented in C currently produce log files that continually increase in size.

To manage log file space usage for log files created by Oracle Fusion Applications modules that are written in C:

1. Navigate to the directory that contains the log files:
 - If the `AFLOG_FILENAME` profile option is set, navigate to the location designated by the profile option value.
 - If the `AFLOG_FILENAME` profile option is not set, navigate to the location set by Oracle Enterprise Scheduler Service.

You can use Fusion Applications Control to determine the location of Oracle Enterprise Scheduler log files, as follows:

- a. In the navigation pane, expand the **Farm** listing, then expand the **Scheduling Services** listing, and then select an Oracle Enterprise Scheduler server as your target.
- b. In the context pane, from the **Scheduling Service** dropdown menu, choose **Logs > View Log Messages**.
- c. Click **Target Log Files** to view a list of log files associated with the selected server.

For example, a typical path and file name might be the following, where *DOMAIN_HOME* is the domain home directory, *SERVER_HOME* is the server home directory, and *serverName* is the name of the Oracle Enterprise Scheduler server:

```
DOMAIN_HOME/servers/SERVER_HOME/logs/serverName-diagnostic.log
```

2. Rename the log file that is currently in use.

For example, if the current log file is called `Cdiagnostic.log`, you might rename it to `Cdiagnostic_MMDDYYYY.log`, where `MMDDYYYY` is the current date.

3. Delete any previously renamed log files that you no longer need.

13.7.5 Using Profile Options to Configure Standard Log Settings

Oracle Fusion applications that use standard logging functionality use profile option values to configure how much information to log. Ordinarily, these profile options are set at the `Site` level, but some can also be set at the `User` level, to gather more information into the log file for a particular user.

Note: For information about configuring Oracle Fusion Middleware log settings, including settings for Oracle Enterprise Scheduler and Oracle Enterprise Crawl and Search Framework, see "Configuring Settings for Log Files" in the *Oracle Fusion Middleware Administrator's Guide*. For information about configuring profile option values for incident and QuickTrace settings, see [Section 14.2](#).

You can use either Fusion Applications Control or the Setup and Maintenance work area to configure profile options for standard Oracle Fusion Applications logging functionality. For information about how to use the Setup and Maintenance work area for this purpose, see the chapter on maintaining common reference objects in the *Oracle Fusion Applications Common Implementation Guide*.

Note: In order to set logging profile option values that affect other users, you must log in as a user who has the Manage All Application Profile Values function security privilege. By default, this privilege is carried by the Applications Common Application Profile Value Administration Duty role, which the predefined Application Administrator job role inherits. You can use Oracle Identity Manager to determine whether you have the Application Administrator job role.

If you want to change profile option values for others without being provisioned with the Application Administrator job role or the Applications Common Application Profile Value Administration Duty role, you can use Oracle Authorization Policy Manager to determine which other duty roles have the Manage All Application Profile Values function security privilege and which job roles inherit those duty roles. You can use Oracle Identity Manager to make sure that you have a job role that inherits one of those duty roles.

For more information about working with roles and privileges, see the *Oracle Fusion Applications Security Guide*, the *Oracle Fusion Middleware User's Guide for Oracle Identity Manager*, and the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)*.

To use Fusion Applications Control to configure profile options for standard Oracle Fusion Applications logging functionality:

1. From the navigation pane, select one of the following target types:

- A product family target that lets you configure the log profile options for all of the Oracle Fusion applications in the selected product family. For example, you could select the **Financials** product family target.
 - An Oracle Fusion Applications cluster application target that lets you configure the log profile options for the selected application on all of the servers in the selected cluster. For example, you could expand the **Financials** listing and the **Fusion Applications** listing and then select the **PayablesApp** cluster application target.
 - An Oracle Fusion application instance target that lets you configure the log profile options for the selected application on the selected sever. For example, you could expand the **Financials** listing and the **Fusion Applications** listing and the **PayablesApp** listing and then select the **PayablesApp (PayablesServer_1)** application instance target.
2. In the context pane, from the target type dropdown menu, choose **Logs > Log Configuration**.
 3. In the Logging Profile Configuration dialog box, click one of the following tabs to determine whether you set the logging level for the whole site or for a single user:
 - **Site-Level**
 - **User-Level**
 4. If you clicked **Site-Level**, skip to Step 5.

If you clicked **User-Level**, complete the following substeps:

- a. Inspect the table on the **User-Level** tab for the user name for which you want to change logging levels.

If the user name is listed, select it and click **Edit** to display the Edit User-Level Configuration dialog box.

If the user name is not listed, click **Add** and enter it in the **Name** field of the Add Logging Profile Configuration dialog box.
- b. From the **Log Level** dropdown list, select the log level that corresponds to the kinds of information you want to gather.

For more information, see [Table 13-2](#).
- c. In the **Enabled** field, make sure that the checkbox is selected.
- d. In the **Module Filters** field, indicate if you want Oracle Fusion Applications to log only for a particular code module for the selected user—by default, log entries are written for all Oracle Fusion Applications code modules that use standard logging implementation for Oracle Fusion Applications.

Note: If you specify that logging for the selected user should be done only for a particular code module, that setting will not affect the information that is logged for other users.

- e. If you want to change whether PL/SQL modules buffer log entries of levels lower than Warning, select **Enabled** or **Disabled** for the **Buffer Mode** setting.
- f. If **Buffer Mode** is set to **Enabled** and you want to change the number of PL/SQL log records that will be buffered in memory before they are written to the log file, enter the number of records to buffer in the **Buffer Size** field.

13.7.6.1 Configuring Additional Log Settings for Oracle Fusion Incentive Compensation Batch Jobs

The logging functionality for certain Oracle Fusion Incentive Compensation batch jobs is separate from the standard logging functionality for Oracle Fusion Applications. By default, the following kinds of Oracle Fusion Incentive Compensation batch jobs do not write log entries:

- Calculation
- Classification
- Collection
- Crediting
- Rollup

For optimum performance and log file sizes, it is recommended that you use the logging functionality for these areas only when troubleshooting an existing problem. For more information, see [Section 14.2.9](#).

13.7.6.2 Configuring Additional Log Settings for Oracle Fusion General Ledger

In the Oracle Fusion Financials product family, some logging functionality for the Oracle Fusion General Ledger application is separate from the standard logging functionality for Oracle Fusion Applications. By default, the following kinds of Oracle Fusion General Ledger batch jobs do not write log entries:

- OpenPeriod
- Posting
- Translation
- Close Process - Create Income Statement Closing Journals
- Close Process - Create Balance Sheet Closing Journals

For optimum performance and log file sizes, it is recommended that you use the logging functionality for these areas only when troubleshooting an existing problem. For more information, see [Section 14.2.10](#).

13.7.6.3 Configuring Additional Log Settings for Oracle Fusion Receivables AutoInvoice

In the Oracle Fusion Financials product family, logging functionality for the AutoInvoice portion of the Oracle Fusion Receivables application is separate from the standard logging functionality for Oracle Fusion Applications. The amount of information that is logged for AutoInvoice depends on the value of the Log File Message Level system option setting.

To set the amount of information that the Oracle Fusion Receivables application logs for the AutoInvoice functionality area:

1. In the Oracle Fusion Receivables application, choose **Setup and Maintenance** from the **Navigator** menu.
2. Complete the following substeps to navigate to the **Edit System Options** screen for the **Manage Receivables System Options** task.
 - a. Expand the **Search: Tasks** pane.
 - b. Enter `Manage Receivables System Options` in the **Name** field and click **Search**.

- c. In the **Manage Receivables System Options** row of the **Search Results** table, click **Go to Task**.
 - d. In the **Search** area of the **Manage Receivables System Options** screen, select **Business Unit** from the drop-down list, enter the name of a business unit for which you want to set up the amount of AutoInvoice information to be logged, and then click **Search**.
Alternately, you can click **Search** without specifying a business unit name to display a list of the available business units.
 - e. In the **Search Results** table, click the name of the business unit for which you want to set up the amount of AutoInvoice information to be logged.
3. In the **Edit System Options** screen, scroll down to display the **AutoInvoice** area of the screen, and then set **Log File Message Level** to the value of 0, which is the recommended value for normal operations.

Note: For optimum performance and log file sizes, it is recommended that you keep the **Log File Message Level** system option set to the lowest value that meets your everyday needs, and increase the value of the setting only when troubleshooting an existing problem. For more information about using this setting when troubleshooting, see [Section 14.2.11](#).

4. Click **Save** to put your changes into effect.

Repeat this entire procedure for each additional business unit for which you want to configure AutoInvoice system options.

For information about configuring the **Maximum Memory in Bytes** system option for AutoInvoice, see product-specific documentation in Oracle Fusion Applications Help.

13.7.7 Configuring Access to Logs for Fusion Applications Control

By default, you can use Fusion Applications Control to view log file entries that are generated by Java code in Oracle Fusion applications. To make it possible to view log file entries that are generated by PL/SQL or C code, you must perform the following configuration steps.

To make PL/SQL and C log file entries visible in Fusion Applications Control:

1. Determine the locations of the PL/SQL and C log files written by your Oracle Fusion applications.

The values of the `AFLOG_PLSQL_FILENAME` and `AFLOG_FILENAME` profile options normally determine log file locations for PL/SQL and C log files, respectively. For more information about these profile options and the log file locations that are used when these profile options are not set, see [Section 13.6.4](#) and [Section 13.7.1](#).

2. Decide which Managed Server you want to use to view PL/SQL log files, C log files, or both, and make sure that the server has read access to all log file locations that you determined in Step 1.

It is recommended that you use an Oracle WebLogic Server that runs Oracle Enterprise Scheduler for this purpose.

Note: If you want to use more than one Managed Server to view PL/SQL log files, C log files, or both, you can do so, provided that each such server has read access to the log files that you want to view through that server. However, it is not especially valuable to use multiple servers to view a particular log file, since each server is providing access to the same data.

3. On a Managed Server that you picked in Step 2, navigate to the following location, where *DOMAIN_HOME* is the domain home for the Managed Server and *server_name* is the name of the Managed Server:

DOMAIN_HOME/config/fmwconfig/servers/*server_name*/diagnostics-registration

If this location does not currently exist, create it now.

4. In the *diagnostics-registration* subdirectory, for each log file that you determined in Step 1 and want to view using this server, use a text editor to create and save a file that contains the following lines, using any file name that has the format *yourChosenFileName.xml*:

```
<?xml version='1.0' encoding='UTF-8'?>
<logs xmlns='http://www.oracle.com/ias/EMComponent/ojdl'>
  <log path="path_to_log_file">
    <logreader class="oracle.core.ojdl.reader.ODLLogReaderFactory">
      </logreader>
    </log>
  </logs>
```

Substitute the full path and name of the log file for *path_to_log_file*, using the form of the path that is correct for the current server to use when accessing the file:

- For a PL/SQL log, make sure that the value for *path_to_log_file* corresponds to the same location and file name as the value that you have set for the *AFLOG_PLSQL_FILENAME* profile option.

Note: You must resolve and enter the full path to the *path_to_log_file* location, explicitly, rather than entering the directory object that is used to set the value of the *AFLOG_PLSQL_FILENAME* profile option.

- For a C log, make sure that the value for *path_to_log_file* indicates the same location and file name as the value that you have set for the *AFLOG_FILENAME* profile option.

Repeat this step as needed to create an individual *.xml* file for each log file that you determined in Step 1 and want to view using this server. You can save each file using any unique legal file name that ends in *.xml*. For example, the file name might be *plsqlifindlog.xml* for a PL/SQL log, and *cfndlog.xml* for a C log.

5. In Fusion Applications Control, verify that each log file from Step 1 is now listed in an appropriate **Target Log Files** list:
 - a. From the navigation pane, select a target for which you created a *.xml* file.

For example, you might expand the **Financials** product family entry, and then expand the **Fusion Applications** and **PayablesApp** entries, and then select the **PayablesApp (PayablesServer_1)** application instance target.

- b. In the context pane, from the dynamic target menu, choose **Logs > View Log Messages** to display the log messages for the selected target.
 - c. Click **Target Log Files**.
 - d. Verify that the resulting list includes the all of the log files that are applicable for the selected target, including any applicable log files for which you created .xml files.
The View Log Message page includes information from all log files that appear on this list.
6. Repeat Step 5 for any other targets for which you created a .xml file.

13.8 Standard Diagnostic Testing Administration Tasks and Tools

Under ordinary circumstances, the following administrative tasks are associated with Oracle Fusion Applications diagnostic tests:

- Configuring security to provide appropriate access to diagnostic tests. You can assign job roles to particular users to grant those users the ability to perform various diagnostic operations. For more information, see [Section 13.9.1](#).

Caution: In the current release, a job role for diagnostic operations grants the user the ability to perform the specified operations for all diagnostic tests that are provided with Oracle Fusion Applications. When choosing whether to grant a diagnostic job role to specific users, be aware that some diagnostic tests may include sensitive information in their results.

- Running diagnostic tests. You can use diagnostic tests for the following purposes:
 - Routinely checking the health of your Oracle Fusion applications
 - Troubleshooting a problem with an Oracle Fusion application
 - Collecting detailed data that may help Oracle Support to resolve a problem for you

Some diagnostic tests require a specific Oracle Fusion application to be running while the test is performed—these diagnostic tests are called **internal** diagnostic tests. Other diagnostic tests can perform their functions even if the Oracle Fusion application to be tested is not running—these tests are called **external** diagnostic tests.

The distinction between internal and external tests is important because it affects both when you can run the tests and which interfaces you can use to run the tests. The Diagnostic Testing Framework provides two interfaces:

- The Diagnostic Dashboard application provides a graphical user interface that lets you perform the following tasks:
 - Execute and monitor both internal and external diagnostic tests for Oracle Fusion applications
 - Purge diagnostic test results
 - Register any special-purpose diagnostic tests that Oracle Support may provide to you
- The `diagctl` command line interface lets you perform the following tasks:

- Execute external diagnostic tests (tests that do not require a specific Oracle Fusion application to be running)

Note: Technical constraints prevent the `diagctl` command line interface from returning useful results for internal diagnostic tests (tests that require a specific Oracle Fusion application to be running when the test is performed). You must use the Diagnostic Dashboard to run any internal diagnostic test.

You must also use the Diagnostic Dashboard to determine whether a particular test is an internal or an external test. For more information about making this determination, see Step 9 through Step 11 in [Section 13.10.1.1](#).

- View diagnostic test results
- Register any special-purpose diagnostic tests that Oracle Support may provide to you

13.9 Configuring the Diagnostic Testing Framework for Normal Operation

You can use diagnostic tests to check normal system health and to troubleshoot system problems. You can configure your Oracle Fusion Applications environment to run all Oracle Fusion Applications diagnostic tests using the Diagnostic Dashboard application, and to run external diagnostic tests using the `diagctl` command line interface.

This section contains the following topics:

- [Controlling Access to Diagnostic Testing Functionality](#)
- [Navigating to the Diagnostic Dashboard](#)

Note: Technical constraints prevent the `diagctl` command line interface from returning useful results for internal diagnostic tests (tests that require a specific Oracle Fusion application to be running when the test is performed). You must use the Diagnostic Dashboard to run any internal diagnostic test.

You must also use the Diagnostic Dashboard to determine whether a particular test is an internal or an external test. For more information about making this determination, see Step 9 through Step 11 in [Section 13.10.1.1](#).

Both the Diagnostic Dashboard application and the `diagctl` command line interface are automatically installed and configured as part of the Oracle Fusion Applications installation. However, you must assign appropriate job roles to specific users to give them the ability to display and perform operations using the Diagnostic Dashboard application. Access to the `diagctl` command line interface is controlled at the level of the server operating system. For more information about granting appropriate access, see [Section 13.9.1](#).

To help you locate diagnostic tests for specific purposes, the diagnostic tests that you receive with Oracle Fusion applications are all assigned to predefined categories.

Note: You cannot change the tag name and tag value assignments that Oracle uses to categorize diagnostic tests, and you cannot remove those tag names or tag values from the database. The following related links in the Task pane of the Diagnostic Dashboard application are intended for use by Oracle personnel, only:

- Add New Tag
 - Add New Tag Value
 - Assign Tags to Tests
 - Unassign Tags from Tests
 - Remove Tag
 - Remove Tag Value
-
-

Caution: Do not attempt to modify the diagnostic test seed data provided to you by Oracle. Unauthorized modification of this seed data may prevent diagnostic tests from functioning correctly, lengthening the amount of time required to resolve both current and future problems.

13.9.1 Controlling Access to Diagnostic Testing Functionality

Access to diagnostic testing functionality is controlled separately for the Diagnostic Dashboard and the `diagctl` command line interface.

For the `diagctl` command line interface, access is controlled at the level of the server operating system. If a user can log in to the server where `diagctl` is stored, and if that user has operating system permissions to read and execute `diagctl`, then that user can use `diagctl` to perform all diagnostic operations that the command line interface supports.

For the Diagnostic Dashboard, you can use Oracle Identity Manager to assign specific users to any of the four preconfigured job roles that grant users access to the Diagnostic Dashboard. Each of these four job roles provides access to a different amount of the functionality of the dashboard.

Note: Oracle Fusion applications display the **Troubleshooting > Run Diagnostic Tests** command in the **Help** menu only for users who are associated with the preconfigured job roles that grant access to Diagnostic Dashboard operations.

- The `Diagnostic Viewer` job role can view and analyze diagnostic test results for Oracle Fusion applications.
- The `Diagnostic Regular User` job role can execute diagnostic test runs and view diagnostic test results for Oracle Fusion applications, and cancel diagnostic test runs that were started by the current user.
- The `Diagnostic Advanced User` job role can schedule and execute diagnostic test runs, view diagnostic test results, attach test results to application incidents for Oracle Fusion applications, and cancel diagnostic test runs that were started by the current user. In general, this job role is recommended for running Oracle Fusion

Applications diagnostic tests, since its added capabilities allow users to work with administrators more flexibly during troubleshooting.

- The `Diagnostic Administrator` job role can use all diagnostic testing functionality provided for Oracle Fusion applications, including purging test results from the database and canceling test runs started by other users.

Caution: In the current release, any job role for diagnostic operations grants the user the ability to perform the role's specified operations for all diagnostic tests that are provided with Oracle Fusion applications. When choosing whether to grant any diagnostic job role to specific users, be aware that some diagnostic tests may include sensitive information in their results.

To grant specific users permission to use the Diagnostic Dashboard:

1. Decide which users need the capabilities of each of the four preconfigured job roles for diagnostic operations.
2. Use Oracle Identity Manager to assign the appropriate job role to each user.

13.9.2 Navigating to the Diagnostic Dashboard

The Diagnostic Dashboard application for Oracle Fusion Applications provides a graphical user interface that lets you execute and monitor diagnostic tests, display and purge test results, and register any special-purpose diagnostic tests that Oracle Support may provide to you. Each product family within Oracle Fusion Applications has its own instance of the Diagnostic Dashboard. Provided that you are assigned to an appropriate job role, you can navigate to the Diagnostic Dashboard from any Oracle Fusion application, or from Cloud Control.

13.9.2.1 Navigating to the Diagnostic Dashboard from an Oracle Fusion Application

If you want to use the Diagnostic Dashboard to execute or monitor diagnostic tests or display or purge test results while you are using an Oracle Fusion application, you can navigate to the Diagnostic Dashboard directly from the application.

To display to the Diagnostic Dashboard from an Oracle Fusion application:

1. Sign in to the relevant Oracle Fusion application as a user who has access to the specific Diagnostic Dashboard operations that you need.

For more information about the job roles that grant access to Diagnostic Dashboard operations, see [Section 13.9.1](#).

2. From the **Help** menu in the application, choose **Troubleshooting > Run Diagnostic Tests** to display the Diagnostic Dashboard.

Note: Oracle Fusion applications display the **Troubleshooting > Run Diagnostic Tests** command in the **Help** menu only for users who are assigned to the preconfigured jobs roles that grant access to Diagnostic Dashboard operations. For more information about these job roles, see [Section 13.9.1](#).

13.9.2.2 Navigating to the Diagnostic Dashboard from Cloud Control:

If you want to use the Diagnostic Dashboard to execute or monitor diagnostic tests or display or purge test results while you are using Cloud Control, such as while you are using Support Workbench to gather additional information about an existing incident, you can navigate to the Diagnostic Dashboard directly from Cloud Control.

To display to the Diagnostic Dashboard from Cloud Control:

1. In Oracle Enterprise Manager, select the product family or cluster application for which you want to run diagnostic tests or view diagnostic test results.
2. From the dynamic drop-down menu, choose **Diagnostics > Fusion Applications Diagnostic Dashboard**.

A login screen for the Diagnostic Dashboard appears in a new window.

3. Log in using an account for the Oracle Fusion Applications product family that you intend to test.

The account that you use must also be assigned to a job role that provides access to the Diagnostic Dashboard. For more information, see [Section 13.9.1](#).

13.10 Using Diagnostic Tests to Monitor Normal System Health

Oracle Fusion Applications diagnostic tests are designed to help you to monitor the health of your system and to help you troubleshoot when necessary.

This section contains the following topics:

- [Running Diagnostic Tests](#)
- [Searching for Diagnostic Tests by Name, Categorization Tag, or Module](#)
- [Checking Diagnostic Test Availability](#)
- [Providing Input Parameters for Diagnostic Tests](#)
- [Running Diagnostic Tests Immediately](#)
- [Scheduling Diagnostic Tests to Run Later](#)
- [Checking the Status of a Diagnostic Test](#)
- [Canceling Diagnostic Test Runs](#)
- [Viewing the Results of Diagnostic Tests](#)
- [Identifying Diagnostic Test Launch Methods from Test Run Names](#)
- [Purging the Results of Selected Diagnostic Test Runs](#)

Note: The user name that you use to sign in to an Oracle Fusion application affects which diagnostic operations are available to you. Be sure that you sign in using an account that has a job role for the diagnostic operations that you need. For more information, see [Section 13.9.1](#).

Some diagnostic tests can be used with all Oracle Fusion applications. Other tests apply to specific product families within Oracle Fusion Applications. For information about the individual diagnostic tests that are provided with this release, see the *Oracle Fusion Applications Common User Guide* in Oracle Fusion Applications Help.

13.10.1 Running Diagnostic Tests

You can use either the Diagnostic Dashboard or the `diagctl` command line interface to run external diagnostic tests—tests that do not depend on the availability of any specific Oracle Fusion application. However, for technical reasons, you must use the Diagnostic Dashboard to run internal diagnostic tests—tests that require a specific Oracle Fusion application to be running at the time when the test is run. For information about determining whether a particular test is an internal or an external test, see Step 9 through Step 11 in [Section 13.10.1.1](#).

This section contains the following topics:

- [Using the Diagnostic Dashboard to Run Diagnostic Tests](#)
- [Using the `diagctl` Command Line Interface to Run Diagnostic Tests](#)

13.10.1.1 Using the Diagnostic Dashboard to Run Diagnostic Tests

The Diagnostic Dashboard application provides a graphical user interface that lets you execute and monitor diagnostic tests, display and purge test results, and register any special-purpose diagnostic tests that Oracle Support may provide to you.

To run a diagnostic test from the Diagnostic Dashboard:

1. Navigate to the Diagnostic Dashboard for the Oracle Fusion applications you are administering, and log in using an account for the application you intend to test that has one of the following job roles:
 - Diagnostic Regular User
 - Diagnostic Advanced User
 - Diagnostic AdministratorFor more information about these job roles, see [Section 13.9.1](#). For more information about navigating to the Diagnostic Dashboard, see [Section 13.9.2](#).
2. In the Regional area on the left side of the Diagnostic Dashboard, locate, and, if necessary, expand one of the following panels:
 - **Search by Tests**
 - **Search by Tags**
3. In the search panel, use standard Oracle query techniques to specify the test characteristics that you want to search for, and then click **Search**.

For information about the kinds of test characteristics that you can specify, see [Section 13.10.2](#).

The results of the search appear in a table below the **Search** button.

Note: In the current release, you can search for and display information about all diagnostic tests that are associated with a tag name that you specify. You cannot currently limit those searches to particular pairs of tag names and tag values. If you need to locate diagnostic tests that are associated with a particular tag name and tag value, you must search for the tag name and then scan the results for the tag value you require.

4. In the search results table, select the checkbox for each test that you want to run, then click **Add To Run**.

The **Choose Tests to Run and Supply Inputs** table appears in the upper Local area of the screen, listing characteristics of the tests you selected.

5. From the **View** menu, use standard techniques to adjust how the table displays its data.

If you want to display nested tests and test steps, you can also expand the tree structure in the **Choose Tests to Run and Supply Inputs** table.

6. In the **Choose Tests to Run and Supply Inputs** table, inspect the **Input Status** column and perform the appropriate action for the value you find:
 - For any root level test that displays the message *Required Input Values Missing*, either clear the checkbox to omit that test (and its nested tests and test steps) from the test run and skip to Step 8, or click the alert icon in the **Input Status** column to display the Input Parameters dialog box and proceed to Step 7.
 - For any root level test that displays the message *Required Input Values Validated* or the message *Input Values Validated*, consider whether you want to inspect the parameter values the test is currently set to use. If so, or if you know that you want to change an existing input parameter value, click the check mark icon in the **Input Status** column to display the Input Parameters dialog box and proceed to Step 7.

Note: You can use the Input Parameters dialog box to override current input parameter values, including in tests that have a valid input status.

- For any root level test that displays the message *No Inputs Required*, skip to Step 8. A test that displays this message does not use input parameters.
7. If you clicked the alert icon in the **Input Status** column to display the Input Parameters dialog box, specify new parameter values as needed, according to the parameter type, as follows; otherwise, skip this step.
 - For Boolean parameters, select the appropriate button in the **New Value** field.
 - For numerical parameters and general text parameters, enter the appropriate value in the **New Value** field.
 - For date parameters, click the icon in the **New Value** field to display a calendar pop-up, and then select an appropriate year, month, and day.
 - For parameters that must be supplied from a list of values (LOV), click the magnifying glass icon in the **New Value** field to display the Search and Select dialog box. In the Search and Select dialog box, select the appropriate value for the parameter and click **OK**.

Note: The values that are available in the list of values are determined in the metadata for the diagnostic test.

- If you want to save your current input values for convenient future use, click **Save** to display the Save As Input Set dialog box. Supply a name for the set of values that you are saving, plus any additional information about the input set that you want to store, and then click **OK**.

- If you want to use a set of previously stored input values, select the appropriate set from the **Input Set** dropdown list, and click **Load**.
- If you want to revert to the default values for all parameters in the current test, click **Defaults**. The Diagnostic Test Framework removes all values for text parameters and resets other parameter types to default values.

When you are finished setting input parameters for the current test, click **OK**.

8. Repeat Step 6 through Step 7 for any other tests that are missing input values, or that have parameter values that you want to override in your test run.
9. From the **View** menu for the **Choose Tests to Run and Supply Inputs** table, choose **Availability** to display the Diagnostic Test Availability dialog box.
10. In the **Select a Diagnostic Test for Details** table, select each test listing and inspect the icon displayed in its **Availability** column:
 - If the **Availability** column shows a check mark icon, then the selected test is currently available to be run.
 - If the **Availability** column shows a triangular warning icon, or if you want to determine whether the test is internal or external, click **Detach** in the **Available Details** table header to display the **Available Details** table in a larger window, and then proceed to Step 11.

If all listed tests are available to be run and if you do not need to know whether they are internal or external, skip to Step 13.

11. Inspect the expanded **Detached Table** for messages about whether the selected test depends on particular Oracle Fusion applications or about why the test is not available, and take the appropriate action:
 - If you want to know whether the selected test is internal or external, inspect the **Details of Required Test Components** column of the **Web Applications Accessible** row.
 - If that cell of the **Detached Table** contains the name of an Oracle Fusion application, then the test is an **internal** diagnostic test that you can run only by using the Diagnostic Dashboard when the specified application is available. You cannot run such a test by using the `diagctl` command line interface or when the application is not present.
 - If that cell of the **Detached Table** does not list an application, then the test is an **external** diagnostic test.
 - If the **Error** column says, "The following Java classes were not loadable:" This message indicates that the diagnostic testing framework cannot locate the JAR file that contains the selected test. Contact your help desk for assistance in searching for a solution in the My Oracle Support Knowledge Base. If the Knowledge Base does not yield a solution, ask your help desk to open an Oracle Support service request.
 - If the **Error** column says, "The following PL/SQL procedures were not located in the database:" This message indicates that the diagnostic testing framework cannot locate the test code for the selected PL/SQL diagnostic test in your database. Contact your help desk for assistance in searching for a solution in the My Oracle Support Knowledge Base. If the Knowledge Base does not yield a solution, ask your help desk to open an Oracle Support service request.
 - If the **Error** column says, "The following Web Applications were inaccessible:" Use the Oracle WebLogic Server console or Fusion Applications Control to check whether the listed applications are running correctly. This message

indicates that the Diagnostic Testing Framework must have access to running instances of the listed Oracle Fusion applications in order to run the selected diagnostic test—the test is an internal test.

For more information about installing and deploying Oracle Fusion Applications, see the "Provisioning a New Applications Environment" chapter of the *Oracle Fusion Applications Installation Guide*.

If the problem persists when the listed Oracle Fusion applications and the relevant database instance are all running, contact Oracle Support for assistance.

- If the **Error** column says, "The current user does not have execution privileges for the following tests:" You must log in as a user who has appropriate privileges to execute the selected test. For information about the privileges required, see [Section 13.9.1](#).
- If the **Error** column says, "The current user does not have privileges to view reports for the following tests:" You must log in as a user who has appropriate privileges to view the results of the selected test. For information about the privileges required, see [Section 13.9.1](#).

Note: It is possible to have the necessary privileges to view diagnostic test results without having the necessary privileges to run those tests. Use an appropriate user account for the action you want to perform.

12. If you have not already done so, close the Detached Table window and the Diagnostic Test Availability dialog box, repeat Step 9 and Step 10 to verify that all listed tests are now available to be run, and then proceed to Step 13.
13. If you wish, enter a name for your test run in the **Run Name** field in the control bar.

Note: Do not use the word `error` in your test run name. If you use the word `error`, or if you leave the **Run Name** field blank, the Diagnostic Testing Framework automatically assigns the test run a name. For information about the formats used in automatically assigned test run names, see [Section 13.10.10](#).

14. When the **Input Status** column of the **Choose Tests to Run and Supply Inputs** table displays `Required Input Values Validated` in all of the selected rows, choose one of the following from the **Run Options** menu:
 - **Run Now:** Run the selected test or tests immediately after you click **Run**.

Note: Do not click **Run** until you have completed Step 15 and Step 16.

- **Run Later:** Schedule when the test or tests will be run. This option is integrated with Oracle Enterprise Scheduler Service. When you select this option, the **Run** button on the toolbar changes to a **Schedule Run** button. Complete the following sub-steps to schedule when the test or tests will be run:

- a. Click the **Schedule Run** button to display a Schedule Tests dialog.
 - b. Click the **Schedule** tab and then select **Use a schedule**.
 - c. From the **Frequency** dropdown list, select how often to run the selected test or tests.
 - d. In the **Start** field, specify the date and time to start the testing.
 - e. Click **Submit**.
15. Adjust option settings to determine whether or not to run the prerequisite tests for the diagnostic tests you have selected:
- If you want to run prerequisite tests, make sure that the **No Prerequisites** option in the **Run Options** menu is *not* selected.
 - If you do *not* want to run prerequisite tests, make sure that the **No Prerequisites** option in the **Run Options** menu is *is* selected.
16. Adjust option settings to determine how many threads to use when running the selected diagnostic tests:
- To use multiple threads, choose **Run in Parallel** from the **Run Options** menu. Then choose **Number of Threads** from the **Run Options** menu and select a value from 2 through 5. The default number of threads for running in parallel is 3.
 - To use a single thread, choose **Run Synchronously** from the **Run Options** menu.
17. If you chose **Run Now** in Step 14, click **Run** to start executing the test run.
- If you chose **Run Later** in Step 14, the test run will start executing at the time you set in the Schedule Tests dialog.

13.10.1.2 Using the `diagctl` Command Line Interface to Run Diagnostic Tests

The Diagnostic Test Framework command line utility, `diagctl`, lets you specify which tests to run in several different ways: by test name, by associated product codes, by associated tag names and tag values, and by associated module IDs or module keys.

You can run one or more diagnostic tests using a single `diagctl` command. It is particularly appropriate to use `diagctl` when you do not have access to a WebLogic Server.

Note: Technical constraints prevent the `diagctl` command line interface from returning useful results for internal diagnostic tests (tests that require a specific Oracle Fusion application to be running when the test is performed). You must use the Diagnostic Dashboard to run any internal diagnostic test.

You must also use the Diagnostic Dashboard to determine whether a particular test is an internal or an external test. For more information about making this determination, see Step 9 through Step 11 in [Section 13.10.1.1](#).

To run Diagnostic Tests from the `diagctl` command line interface:

1. Obtain the user name and password for the Oracle Fusion Applications account that will run the diagnostic test or tests, and the password for that account.

A user name and password is required for any diagnostic test that you run using the `diagctl` command line interface. The command line syntax for specifying the user name and password is `un=user_name` and `pwd=password`.

2. Decide which of the following methods to use to specify the diagnostic test or tests that you want to run:

- Specify a single test name: To run a single specific test without specifying input parameters, the command line syntax is `test=test_name`
- Specify a test name and parameters: To run a single specific test with one or more input parameters, the command line syntax is `test=test_name input:parameter_name1=parameter_value1 input:parameter_name2=parameter_value2`
- Specify multiple test names: To run several specific tests, the command line syntax is `test=test_name1, test_name2, test_name3`

Note: If you are specifying multiple tests on a single command line, then you cannot specify input parameters on that command line.

- Specify by product codes: To run all of the tests that are associated with one or more specific product codes in the Applications taxonomy, the command line syntax is `app=product_code1, product_code2, product_code3`
- Specify by module ID: To run all of the tests that are associated with specific module IDs in the Applications taxonomy, and all of the tests that are associated with child modules of the module that you specify, the command line syntax is `modid=moduleID1, moduleID2, moduleID3`
- Specify by module key: To run all of the tests that are associated with specific module keys in the Applications taxonomy, the command line syntax is `modkey=module_key1, module_key2, module_key3`
- Specify by tag name and tag value: To run all of the tests that are associated with a specific tag and tag value in the diagnostic test repository, and to run any tests that are associated with any child tag values of the tag name and tag value that you specify, the command line syntax is `tag: tagname1=tagvalue1 tag: tagname2=tagvalue2 tag: tagname3=tagvalue3`

You must use at least one these options in each command to run a diagnostic test from `diagctl`. You can include more than one of these options in a single command, if you prefer.

3. Decide whether to use any, some, or all of the following additional options for the test run:

- Specify a test run name: To specify a particular name for the test run, use the command line syntax `runname=run_name`.

Note: Do not include the word "error" in your test run name. If you include the word "error," or if you do not specify a test run name, the command line utility automatically generates a name for the test run. Automatically generated test run names start with the test name, product code, module ID, module key, or tag name and value that you specified, followed by a colon, a timestamp, another colon, and a sequence number.

- Specify whether to test recursively: To run all of the specified tests recursively, use the command line syntax `recurse=Y`. The default value is N.
 - Specify whether to run prerequisite tests: To identify and run any tests that are prerequisites before running the specified tests, use the command line syntax `prereq=Y`. The default value is N.
 - Specify monitoring interval: To specify how often the status of the test run is uploaded to the test repository, use the command line syntax `moninterval=time_in_seconds`. The default value is 30 seconds.
 - Specify number of threads: To specify the number of parallel threads to spawn for processing this test run, use the command line syntax `nthreads=number_of_threads`. The default value is 5. A value of 1 directs the utility to run the tests serially.
4. At a command prompt for your operating system, enter the appropriate following commands, where `FA_MW_HOME` is the fusionapps Oracle Fusion Middleware home directory and `jdk6_folder_location` is the location of the jdk6 folder labeled **Oracle JRockit Oracle Home** in [Figure 1-6](#) (the Java Development Kit that was included in the Oracle Fusion Applications media pack):

```
(UNIX) export JAVA_HOME=jdk6_folder_location
(UNIX) export MW_HOME=FA_MW_HOME
```

```
(Windows) set JAVA_HOME=jdk6_folder_location
(Windows) set MW_HOME=FA_MW_HOME
```

5. Still at an operating system command prompt, navigate to the location of the `diagctl` executable:

```
(UNIX) ATGPF_ORACLE_HOME/bin/diagctl.sh
(Windows) ATGPF_ORACLE_HOME\bin\diagctl.cmd
```

`ATGPF_ORACLE_HOME` is located under the fusionapps Oracle Fusion Middleware home directory (`FA_MW_HOME`).

6. Enter `diagctl.sh run` (for UNIX) or `diagctl run` (for Windows) followed by the user name and password from Step 1 and the options that you decided upon in Step 2 and Step 3, using the syntax described in those steps.

Note: You can list command arguments that appear after the word `run` in any order. If you do not specify the password on the command line, the utility will prompt you to supply it. For detailed help about running diagnostic tests, enter `diagctl.sh run help` (for UNIX) or `diagctl run help` (for Windows)

For example, to run a single test with two input parameter values specified, you would enter a command such as the following:

```
(UNIX) diagctl.sh run test=oracle.apps.fnd.appltest.sampleTest
input:param1=value1 input:param2=value2 un=sysadmin
```

```
(WINDOWS) diagctl run test=oracle.apps.fnd.appltest.sampleTest
input:param1=value1 input:param2=value2 un=sysadmin
```

Note: If you specify an invalid parameter value, the command line interface returns an error message and does not run the test.

Or, to run all tests that belong to the Application Object Library (FND) and General Ledger products, and to run them recursively and with prerequisite analysis you would enter a command such as:

```
(UNIX) diagctl.sh run app=FND,GL recurse=Y prereq=Y un=sysadmin  
(WINDOWS) diagctl run app=FND,GL recurse=Y prereq=Y un=sysadmin
```

Or, to run all tests that are associated with the given module id, you would enter:

```
(UNIX) diagctl.sh run modid=module1,module2 un=sysadmin  
(WINDOWS) diagctl run modid=module1,module2 un=sysadmin
```

13.10.2 Searching for Diagnostic Tests by Name, Categorization Tag, or Module

You can use the Diagnostic Dashboard to search for available diagnostic tests.

One way you can search is to use **Search by Tags** to search for diagnostic tests that are associated with a particular categorization tag name that you specify.

Another way you can search is to use **Search by Tests** to search for tests with values that you specify for one or more of the following characteristics:

- Test name
- Test description
- Name of the Oracle Fusion Applications code module that the test exercises

The Diagnostic Testing Framework treats any search value that you enter as a case-insensitive matching pattern, returning the result if the pattern appears in anywhere in the specified field. However, if you specify more than one search value to match, the framework returns results only for diagnostic tests that match all of the values that you specify (logical AND).

In general, searching for diagnostic tests is done as a portion of the process of running diagnostic tests from the Diagnostic Dashboard. For more information, see [Section 13.10.1.1](#).

13.10.3 Checking Diagnostic Test Availability

Whether or not you can run a diagnostic test at any given time depends on both the specific requirements of the test and the current state of your Oracle Fusion Applications system. Any of the following factors can prevent a test from being available:

- Java class availability
- PL/SQL procedure availability
- Oracle Fusion Applications accessibility
- Execution privileges for the test
- Report viewing privileges for the test

In general, checking the availability of diagnostic tests is done as a portion of the process of running diagnostic tests. For more information, see [Section 13.10.1](#).

Note: Some diagnostic tests require a specific Oracle Fusion application to be running while the test is performed—these diagnostic tests are called **internal** diagnostic tests. Other diagnostic tests can perform their functions even if the Oracle Fusion application to be tested is not running—these tests are called **external** diagnostic tests.

13.10.4 Providing Input Parameters for Diagnostic Tests

Diagnostic tests often have input parameters. Oracle supplies default values for some input parameters. When you are in the process of preparing to run one or more diagnostic tests, you can change the values for input parameters that have default values and enter values for input parameters that do not have default values. If you know that you will use the same parameter values more than once, you can save those values into an input set that you can reuse for later test runs.

All required input parameters must have values assigned before you can run a diagnostic test. If there are required parameter values missing, the Diagnostic Dashboard application displays *Required Input Values Missing* in the **Input Status** column of the **Choose Tests to Run** and **Supply Inputs** table.

To specify input parameter values, click the icon in the **Input Status** column to display the Input Parameters dialog box. Then you can either enter parameter values individually in the **New Value** column of the **Edit Input Set** table, or you can select a previously saved set of values from the **Input Set** dropdown list. After you have specified all required input parameters, the Diagnostic Dashboard application displays a check mark icon with the message *Inputs Edited: Required Input Values Validated*. For more information, see [Section 13.10.1](#).

13.10.5 Running Diagnostic Tests Immediately

You can use either the Diagnostic Dashboard or the `diagctl` command line interface to run diagnostic tests immediately.

In the Diagnostic Dashboard, you specify that you want to run a diagnostic test immediately by choosing **Run Now** from the **Run Options** menu.

For the `diagctl` command line interface, you specify that you want to run a diagnostic test immediately by entering the command to run the test from an interactive session prompt.

For more information, see [Section 13.10.1](#).

13.10.6 Scheduling Diagnostic Tests to Run Later

In the Diagnostic Dashboard, you can specify a particular time to run diagnostic tests by choosing **Run Later** from the **Run Options** menu. For more information, see [Section 13.10.1.1](#).

To run a delayed diagnostic test using the `diagctl` command line interface, create a script that calls `diagctl.sh` (for UNIX) or `diagctl.cmd` (for Windows) and then use standard Oracle Enterprise Scheduler techniques to schedule when the script runs. For information about submitting and monitoring Oracle Enterprise Scheduler jobs, see [Chapter 5](#).

13.10.7 Checking the Status of a Diagnostic Test

You can check the status of a diagnostic test from the Diagnostic Dashboard or from the `diagctl` command line interface. The command line interface is primarily intended for use if the Diagnostic Dashboard is temporarily unavailable.

This section contains the following topics:

- [Using the Diagnostic Dashboard to Check the Status of a Diagnostic Test](#)
- [Using the `diagctl` Command Line Interface to Check the Status of a Diagnostic Test](#)

13.10.7.1 Using the Diagnostic Dashboard to Check the Status of a Diagnostic Test

In the Diagnostic Dashboard, the Diagnostic Test Run Status table displays two types of status information:

- **Execution Status:** This column displays status information about whether or not a test run request can be executed successfully.
- **Diagnostic Status:** This column displays status information about whether or not individual diagnostic tests detect problems.

Note: After running a diagnostic test using the Diagnostic Dashboard, you may need to click **Refresh** to display the latest status information, including rows for the following kinds of test runs:

- Test runs that were run immediately from the Diagnostic Dashboard
- Test runs that were scheduled to be run later from the Diagnostic Dashboard
- Test runs that were submitted using the `diagctl` command line interface.

When you click **Refresh**, the Diagnostic Dashboard displays listings for any additional test runs that you or other users have submitted in your current Oracle WebLogic Server domain. If your Oracle Fusion Applications deployment uses Global Single Instance (GSI), the Diagnostic Dashboard also displays listings for test runs that were submitted in other domains.

To check the status of a diagnostic test using the Diagnostic Dashboard:

1. If you started the diagnostic test from your current Diagnostic Dashboard application session, the **Diagnostic Test Run Status** table automatically displays in the lower right portion of the screen after you click **Run**. Skip to Step 4.
2. If you are not already displaying the Diagnostic Dashboard for the Oracle Fusion applications you are administering, navigate to the dashboard and log in using an account for the application you are testing. For more information, see [Section 13.9.2](#).
3. In the Regional area of your screen, expand the **Tasks** panel and click **Run Status**.
4. If you want the **Diagnostic Test Run Status** table to display only certain types of rows, select one of the following options from the **Find** dropdown list:
 - **All Runs Submitted in Last Hour**
 - **All Runs Submitted in Last 24 Hours**

- All Running
 - All Running Submitted in Last Hour
 - All Runs with Diagnostic Failures in Last 24 Hours
 - All Runs with Diagnostic Failures
 - All Runs with Diagnostic Warnings
 - All Runs with Execution Errors
 - All Completed
 - All Completed with No Issues
 - All Runs
 - All Runs Run By the Current User in the Last Hour
 - All Runs Run By the Current User
5. If you want to search for specific rows in the **Diagnostic Test Run Status** table, click the **Search Test Runs** icon, enter search criteria, and click **OK**.
- Available search criteria include:
- Test run dates
 - Test run names
 - Test display names
 - User names of those who launched test runs
 - Execution status
6. If you want to display additional columns in the **Diagnostic Test Run Status** table, choose the additional columns from the **View** menu.
7. Expand test run nodes as needed to view the list of test executions for each test, and then inspect the **Execution Status** column for information about whether tests and test runs have completed or encountered execution errors.
8. Inspect the **Diagnostic Status** column for information about whether completed tests and test runs encountered any issues before completing.
9. For additional information about any test execution, click the icon in the **Report** column of the appropriate row of the table.

13.10.7.2 Using the `diagctl` Command Line Interface to Check the Status of a Diagnostic Test

The Diagnostic Test Framework command line utility, `diagctl`, provides three different ways that you can specify the diagnostic test for which you want status information: by test run name, by test run ID, and by test execution ID.

To check the status of a diagnostic test using the `diagctl` command line interface:

1. Obtain the user name and password for the Oracle Fusion Applications account that will run the diagnostic test or tests, and the password for that account.

A user name and password is required whenever you use the `diagctl` command line interface to check the status of a diagnostic test. The command line syntax for specifying the user name and password is `un=user_name` and `pwd=password`.
2. Decide which of the following methods to use to specify the diagnostic test run for which you want status information:

- Specify a test run name: To check the status of a diagnostic test run for which you have the run name, the command line syntax is `runName=run_name`.
- Specify a test run ID: To check the status of a diagnostic test for which you have the test run ID, the command line syntax is `runid=run_ID`.
- Specify a test execution ID: To check the status of a diagnostic test for which you have the execution ID, the command line syntax is `execid=execution_ID`.

You must use at least one these options in each command to check the status of a diagnostic test using `diagctl`.

3. Decide whether you want to view the status of nested test runs.

To check the status of all diagnostic tests that are nested within the specified test, the command line syntax is `printtree=Y`. This setting defaults to a value of `N`, meaning that the status is reported only for the specified test.

4. At a command prompt for your operating system, navigate to the location of the `diagctl` executable under the `fusionapps` Oracle Fusion Middleware home directory:

```
(UNIX) FA_MW_HOME/atgpf/bin/diagctl.sh
(Windows) FA_MW_HOME\atgpf\bin\diagctl.cmd
```

5. Enter `diagctl.sh status` (for UNIX) or `diagctl status` (for Windows) followed by the user name and password from Step 1 and the options that you decided upon in Step 2 and Step 3, using the syntax described in those steps.

Note: You can list command arguments that appear after the word `status` in any order. If you do not specify the password as part of the command, the utility will prompt you to supply it. For detailed help about getting status information, enter `diagctl.sh status help` (for UNIX) or `diagctl status help` (for Windows).

For example, to check the status of a test by using a test run name, you would enter a command such as:

```
(UNIX) diagctl.sh status runName=TrialRun1 un=sysadmin
(Windows) diagctl status runName=TrialRun1 un=sysadmin
```

Or, to check the status of a test and its nested test runs by using a run ID, you would enter a command such as:

```
(UNIX) diagctl.sh status runid=RunID1 printtree=Y un=sysadmin
(Windows) diagctl status runid=RunID1 printtree=Y un=sysadmin
```

Or, to check the status of a test by using an execution key, you would enter:

```
(UNIX) diagctl.sh status execid=TestExecID1 un=sysadmin
(Windows) diagctl status execid=TestExecID1 un=sysadmin
```

13.10.8 Canceling Diagnostic Test Runs

From time to time, you may want to stop a diagnostic test or test run that is currently running. Several constraints affect your ability to cancel a diagnostic test or test run:

- In the current release, you must use the Diagnostic Dashboard to cancel a diagnostic test or test run that is currently running. The `diagctl` command line interface does not provide this capability.
- To cancel a test or test run that you started, you must use an account that has been assigned the Diagnostic Regular User, Diagnostic Advanced User, or Diagnostic Administrator job role.
- To cancel a test or test run that another user started, you must use an account that has been assigned the Diagnostic Administrator job role.
- When a diagnostic test is scheduled to run at a later time, it is immediately submitted to the Oracle Enterprise Scheduler Service. The procedure for canceling such a test depends on whether or not the test has already started to execute at the time when you want to cancel it:
 - To cancel a scheduled diagnostic test that has not yet started to execute, use standard Oracle Enterprise Scheduler techniques to cancel the job. For more information about cancelling scheduled jobs, see the "Cancelling Oracle Enterprise Scheduling Service Job Requests" section in the *Oracle Fusion Middleware Administrator's Guide for Oracle Enterprise Scheduling Service*.
 - To cancel a scheduled diagnostic test that the Diagnostic Test Run Status dashboard panel indicates is already running, use the Diagnostic Dashboard to cancel it in the same way that you would cancel a test that ran immediately.

When you cancel a diagnostic test or test run, the consequences may vary slightly depending on how much of the test run has executed and the language in which the test code is written:

- If you cancel a diagnostic test run while a test step from that test run is in progress, the test step that is currently running is canceled. No additional tests or test steps in the remainder of the cancelled test run are executed. The Diagnostic Test Run Status panel displays an Execution Status of `Canceled` for the run, and a pop-up window displays a message such as the following, but no log messages are recorded to indicate that a diagnostic test run has been canceled:

```
Test Run "test_name" has been canceled. Please check the test run report for further details.
```

- When you cancel a diagnostic test step that has been implemented in Java, the diagnostic framework automatically closes the test step's database connection, using an asynchronous command to kill the thread. However, when you cancel a diagnostic test step that has been implemented using PL/SQL, the diagnostic framework cannot interrupt the test step and use the existing database connection to close that connection. To reclaim the resources allocated to a canceled PL/SQL diagnostic test step, you must establish a separate connection to the database and use an `alter system kill session` command to close the connection that the canceled test step was using, as described later in this section.

To cancel a diagnostic test using the Diagnostic Dashboard:

1. If you started the diagnostic test from your current Diagnostic Dashboard application session, the **Diagnostic Test Run Status** table automatically displays in the lower right portion of the screen after you click **Run**. Skip to Step 4.
2. If you are not already displaying the Diagnostic Dashboard for the Oracle Fusion applications you are administering, navigate to the dashboard and log in using an account for the application you are testing. For more information, see [Section 13.9.2](#).

3. In the Regional area of your screen, expand the **Tasks** panel and click **Run Status**.
4. In the **Diagnostic Test Run Status** panel, locate and select the test run that you want to cancel, verify that its **Execution Status** is **Running**, and then click **Cancel**.
5. In the **Diagnostic Test Run Status** panel, click the **Report** icon for the canceled test run.
6. Inspect the test run report to determine whether the canceled test step was implemented using Java or PL/SQL.
 - If the canceled test step was implemented using Java, skip all of the remaining steps in this procedure.
 - If the canceled test step was implemented using PL/SQL, proceed to Step 7.
7. If the canceled test step was implemented using PL/SQL, make a note of the session ID and serial number for the database connection that the step was using.

For example, the report might display information similar to the following:

Step Report - Diagnostics_Engine_Log

Session Information

The test *test_name* is using a database connection with Session Id 944 and Serial Number 817

8. Using your preferred database client or database monitoring application and a separate connection to the database, determine whether the database connection for the canceled test step is still open.

For example, you could use a SQL client to execute the following query, substituting the session identifier and serial number values that you obtained in Step 7 for *session_Id* and *serial_number*:

```
select * from v$session where sid = session_Id and serial# =
serial_number
```

If this query returns a row that contains the session ID and serial number that you specified, then the database connection from the canceled test is still open and consuming resources. Proceed to Step 9.

If the query does not return a row that contains the session ID and serial number that you specified, then the database connection from the canceled test has been closed. In this case, skip Step 9.

9. If the database connection for the cancelled test step is still open, use a command such as the following to close that connection, substituting the correct session identifier and serial number values for *session_Id* and *serial_number*.

```
alter system kill session 'session_Id, serial_number';
```

13.10.9 Viewing the Results of Diagnostic Tests

You can use either the Diagnostic Dashboard or the `diagctl` command line interface to view reports that show the results of diagnostic tests, whichever you prefer.

This section contains the following topics:

- [Using the Diagnostic Dashboard to View the Results of Diagnostic Tests](#)
- [Using the diagctl Command Line Interface to View the Results of Diagnostic Tests](#)

13.10.9.1 Using the Diagnostic Dashboard to View the Results of Diagnostic Tests

You can view the results of a diagnostic test in the dashboard by checking the status of the test and then clicking the icon in the **Report** column of the selected table row. For more information, see [Section 13.10.7.1](#).

13.10.9.2 Using the `diagctl` Command Line Interface to View the Results of Diagnostic Tests

The `diagctl` command line utility provides three different ways of requesting the results of diagnostic test results: by test run name, by test run ID, and by test execution ID.

To view diagnostic test result reports using the `diagctl` command line interface:

1. Obtain the user name and password for the Oracle Fusion Applications account that you will use to view the test results, and the password for that account.

A user name and password is required whenever you use the `diagctl` command line interface to view the results of a diagnostic test. The command line syntax for specifying the user name and password is `un=user_name` and `pwd=password`.
2. Decide which of the following methods to use to specify the diagnostic test run for which you want to view results:
 - Specify a test run name: To view the results of a diagnostic test run for which you have the run name, the command line syntax is `runName=run_name`. This option includes results for all the executions in the test run.
 - Specify a test run ID: To view the results of a diagnostic test for which you have the test run ID, the command line syntax is `runid=run_ID`. This option includes results for all of the executions in the run.
 - Specify a test execution ID: To view the results of a diagnostic test for which you have the execution ID, the command line syntax is `execid=execution_ID`. This option includes results for the specified execution and any nested executions.

You must use at least one these options in each command to view the results of a diagnostic test using `diagctl`.

3. Decide whether you want to use one or more of the following additional options:
 - Specify a destination directory for results: To write the test results to a specific directory, the command line syntax is `destdir=destination_directory`. If you do not specify a directory, reports are placed in the `java.io.tmpdir/user.name/diagfwk` directory where `java.io.tmpdir` and `user.name` are Java system properties.
 - Specify a format for the result report: Valid values are XML and HTML. The default value is HTML. XML report files are created as a step toward creating HTML report files. These XML report files remain in the same directory as the HTML report files.
 - Specify if the report should be translated: Valid values are Y and N. If the value is Y, any NLS keys that are specified in the report are translated to your session language. If the value is N, no translation is performed. The default value is Y.
4. At a command prompt for your operating system, navigate to the location of the `diagctl` executable under the `fusionapps` Oracle Fusion Middleware home directory:

```
(UNIX) FA_MW_HOME/atgpf/bin/diagctl.sh
(Windows) FA_MW_HOME\atgpf\bin\diagctl.cmd
```

5. Enter `diagctl.sh report` (for UNIX) or `diagctl report` (for Windows), followed by the user name and password from Step 1 and the options that you decided upon in Step 2 and Step 3, using the syntax described in those steps.

Note: You can list command arguments that appear after the word `report` in any order. If you do not specify the password as part of the command, the utility will prompt you to supply it. For detailed help about viewing reports, enter `diagctl.sh report help` (for UNIX) or `diagctl report help` (for Windows).

For example, to view the results of a test by using a run name, and to place the results in a particular directory, you would enter a command such as:

```
(UNIX) diagctl.sh report runName=TrialRun1 destdir=/d1/testreport un=sysadmin
(Windows) diagctl report runName=TrialRun1 destdir=/d1/testreport un=sysadmin
```

Or, to view the results of a test run by using a run ID, with the results placed in the default location, you would enter a command such as:

```
(UNIX) diagctl.sh report runid=RunID1 un=sysadmin
(Windows) diagctl report runid=RunID1 un=sysadmin
```

Or, to check the status of a test by using an execution key, with the results placed in the default location, you would enter:

```
(UNIX) diagctl.sh report execid=TestExecID1 un=sysadmin
(Windows) diagctl report execid=TestExecID1 un=sysadmin
```

6. Navigate to the location of the results file, and use a browser or text editor of your choice to view it.

13.10.10 Identifying Diagnostic Test Launch Methods from Test Run Names

Any test run name that the Diagnostic Testing Framework supplies follows naming conventions that reflect how the test was launched. If you find that an unfamiliar diagnostic test run occurred at an unexpected time, then knowing the test run naming conventions can help you to understand the circumstances in which the test run was launched:

- When you submit a test run to run immediately, without specifying a run name, the name that is automatically assigned to the test run has the format `TestRun_runID`, where `runID` is a unique string of alphanumeric characters generated by the Diagnostic Testing Framework, such as `TestRun_91D818BA54BB29C8E040578C495D6956`. This naming convention applies both when tests are submitted from the Diagnostic Dashboard and when tests are submitted using `diagctl`.
- When you schedule a diagnostic test run to run later, using Oracle Enterprise Scheduler Service, the name that is automatically assigned to the test run has the format `[TestRunName_]ESS_requestID_timestamp`, where `TestRunName` is an optional name that the you may specify when submitting the test, `requestID` is a unique identifier supplied by the Oracle Enterprise Scheduler when it schedules the job, and `timestamp` indicates the time when the job is scheduled to run. This naming convention applies to tests that are submitted from the Diagnostic

Dashboard to run later. (Tests that are submitted using `diagctl` always run immediately.)

For example, if you enter a test run name of `routine`, the full test run name might be:

```
routine_ESS-417-2010-09-09T17:07:09.115-0700
```

- When a test run is submitted automatically because an incident occurred, the name that is automatically assigned to the test run has the format `AppsLogger-DiagnosticTestingFrameworkIntegration_id_timestamp`, where `id` is a number that uniquely identifies the incident, and `timestamp` indicates when the diagnostic test run starts.

13.10.11 Purging the Results of Selected Diagnostic Test Runs

From time to time, you may want to remove diagnostic test run results from your database, to keep the Run Status table from becoming too large.

Note: To remove the results of one or more diagnostic test runs from your database, you must use the Diagnostic Dashboard application with an account that has been assigned the `Diagnostic Administrator` job role.

The `diagctl` command line interface does not currently provide a way to purge test run results.

To purge selected diagnostic test run results from the database:

1. If you started the diagnostic test from your current application session, the **Diagnostic Test Run Status** table automatically displays in the lower right portion of the screen after you click **Run**. Skip to Step 4.
2. If you are not already displaying the Diagnostic Dashboard for the Oracle Fusion applications you are administering, navigate to the dashboard and log in using an account for the application that needs its test results purged. For more information see [Section 13.9.2](#).
3. In the Regional area of your screen, expand the **Tasks** panel and click **Run Status**.
4. In the **Diagnostic Test Run Status** table header, use either of the following methods to locate the test run status records that represent the test results that you want to remove from the database:
 - Select any appropriate filter from the **Find** dropdown list.
 - Click the Search Test Runs icon, enter search criteria, and click **OK**.
5. After searching or filtering, inspect the listings displayed in the **Diagnostic Test Run Status** table, and decide whether you want to remove listed test results from the database individually or as a group:
 - To remove listed test results from the database individually, proceed to Step 6.
 - To remove all of the listed test results from the database, skip to Step 7.
6. To remove a single test run status result record from the database, complete the following substeps in the **Diagnostic Test Run Status** table:

- a. Select the test run status record that represents the results that you want to remove from the database, and click the delete button in the **Diagnostic Test Run Status** table header.
- b. In the Delete Test Run dialog box, select `Delete test run "TestRunName"` and click **OK**.

The selected record is removed from the database immediately.

Skip Step 7 and Step 8.

7. To remove all of the listed test run status result records from the database, click the delete button in the **Diagnostic Test Run Status** table header.

Depending on whether or not you selected a record in the **Diagnostic Test Run Status** table, either the Delete Test Run dialog box or the Confirm Test Run Delete dialog box appears.

8. Complete the purge process using the appropriate instructions for the dialog box displayed on your screen:
 - If your screen displays the Delete Test Run dialog box, select `Delete all test runs in the list` and click **OK** to remove all of the listed test run results from the database immediately.
 - If your screen displays the Confirm Test Run Delete dialog box, click **Yes** to remove all of the listed test run results from the database immediately.

Troubleshooting Oracle Fusion Applications Using Incidents, Logs, QuickTrace, and Diagnostic Tests

This chapter discusses how to use incidents, log settings, log files, QuickTrace, and diagnostic tests to help diagnose and resolve problems with Oracle Fusion Applications.

This chapter also discusses how incidents are related to log files, QuickTrace output, and diagnostic tests.

This chapter contains the following topics:

- [Introduction to Troubleshooting Using Incidents, Logs, QuickTrace, and Diagnostic Tests](#)
- [Configuring Log, QuickTrace, and Incident Settings for Troubleshooting](#)
- [Configuring the Diagnostic Testing Framework for Troubleshooting](#)
- [Investigating, Reporting, and Solving a Problem](#)
- [Troubleshooting Using Log Settings](#)
- [Troubleshooting Using the Diagnostic Testing Framework](#)

For information about using log settings, log files, and diagnostic tests to monitor normal operations for Oracle Fusion Applications, and for information about how log files and diagnostic tests are related to each other, see [Chapter 13](#). For more information about the individual diagnostic tests that are provided with this release, see the *Oracle Fusion Applications Common User Guide* in Oracle Fusion Applications Help.

14.1 Introduction to Troubleshooting Using Incidents, Logs, QuickTrace, and Diagnostic Tests

Incidents are collections of information about problematic system events. Incidents can be created automatically or manually. An automatically created incident normally contains information about the state of a particular Oracle WebLogic Server domain at the time when a specific problem occurred. Monitoring for and responding to incidents is a standard activity for help desk staff who support Oracle Fusion Applications in organizations that use Information Technology Infrastructure Library (ITIL) best practices.

This section contains the following topics:

- [Relationships Between Incidents, QuickTrace Dumps, Log Files, and Diagnostic Tests](#)
- [Introduction to the Oracle Fusion Middleware Diagnostic Framework](#)
- [Standard Incident Administration Tasks and Tools](#)

You can use incidents to diagnose and resolve problems, or to supply your help desk personnel or Oracle Support personnel with information about more complex problems. Oracle developers set up error messages in such a way that incidents are created automatically when certain high-priority error conditions occur. You can also create incidents manually.

Automatically created Oracle Fusion Applications incidents normally include diagnostic dump files. Dump files are provided in order to assist you in resolving incidents for your organization and to assist Oracle representatives in providing you with effective support when needed. The types of dump files provided depend on the nature of the incident. Some incidents include Oracle Fusion Middleware dump files such as Oracle WebLogic Server diagnostic images, Java Virtual Machine thread dumps, and Dynamic Monitoring Service metric dumps. For more information about Oracle Fusion Middleware dumps, see the "Diagnosing Problems" chapter in the *Oracle Fusion Middleware Administrator's Guide*.

Note: Each Oracle Fusion Applications module is written in one of the following code languages: Java, SOA, PL/SQL, or C. The amount and type of information automatically included with an incident may vary depending on the coding language of the relevant application module.

In general, you can administer Oracle Fusion Applications without knowing which programming language implements particular modules. However, in order to monitor and diagnose all types of incidents successfully, it is important to become familiar with all of the incident configuration settings and the types of diagnostic dump files that are automatically provided for different kinds of incidents.

For more information about incident configuration settings for Oracle Fusion Applications, see [Section 14.2.4](#).

Oracle Fusion Applications incidents may also contain diagnostic log excerpts and diagnostic test results that are specific to Oracle Fusion Applications. Before you transmit incident information to Oracle Support personnel, you can add more information (such as screen shots or more log file information) to incidents that were created either automatically or manually.

14.1.1 Relationships Between Incidents, QuickTrace Dumps, Log Files, and Diagnostic Tests

Incidents are designed to work along with other diagnostic features in your Oracle Fusion applications:

- **Error-Handling:** The Oracle Fusion Applications code that handles errors is designed so that errors that need quick attention automatically create incidents.
- **Diagnostic Tests:** Oracle developers create tests that you can use to help diagnose and resolve Oracle Fusion application problems. A diagnostic test may or may not be associated with a particular error message. If Java or SOA code in an Oracle Fusion application handles a particular error in a way that triggers the creation of

an incident, then any diagnostic tests that are associated with the error message for the incident run automatically. The test results are associated with the incident and the identity of the user who received the error message is recorded.

- **Log files:** By default, whenever an Oracle Fusion Applications incident is created automatically, an `INCIDENT_ERROR` level entry is recorded in the standard Oracle Fusion Applications log file and a copy of that log file is included with the incident. If the incident occurs in Oracle Fusion Middleware code two different kinds of log information are automatically included with the incident: a dump file containing all log entries for the relevant execution context ID (ECID) and either a complete log file or a log file excerpt starting five minutes prior to the incident.

For some incidents, you may find all the information that you need to resolve the problem in the log or dump files automatically included with the incident. In other cases, you may need to configure your system to log more detailed information, attempt to reproduce the problem, and then use the more detailed information in the log to diagnose the problem. If you work with Oracle Support to resolve a problem, log file excerpts are among the kinds of information that you can add to incidents before sending them to Oracle.

- **QuickTrace:** If you need to troubleshoot a problem that has generated an incident, and if the information in the relevant Oracle Fusion application log file is insufficient to resolve the problem, then a QuickTrace dump file included with the incident information may supply you with the additional information you need. If so, you can avoid increasing the amount of information to be logged and then attempting to reproduce the problem. (In general, you should increase the level of detail that you collect in your Oracle Fusion Applications logs only when absolutely necessary, as collecting additional information may decrease system performance.)

By default, QuickTrace continuously records a specified level of log detail in an area of memory, generally at a more granular level of detail than is recorded in Oracle Fusion application log files. The memory that QuickTrace uses is recycled on an ongoing basis, with the oldest information being overwritten first. Because QuickTrace does not format information and because it writes to memory instead of to a log file, it can gather operational information continuously without significantly affecting system performance. The information that QuickTrace stores in memory is written to disk only when an incident occurs or when an administrator manually dumps the contents of a QuickTrace buffer.

Note: The same commands in Oracle Fusion Applications code cause messages to be logged both in Oracle Fusion application diagnostic log files and in QuickTrace dumps. However, your system may gather different amounts of detail for the two kinds of output. At default setting levels, QuickTrace stores much more detailed information than diagnostic log files do.

- **Version Information:** If you request assistance from Oracle Support to resolve an incident, you may be asked to provide information about the Oracle product versions you are running along with the incident information. In your Oracle Fusion applications, you can obtain some product version information by selecting *About Applications* from the **Help** menu. Additional version information can be obtained from Remote Diagnostic Agent (RDA) reports. For more information about RDA, see the "Generating an RDA Report" section in the *Oracle Fusion Middleware Administrator's Guide*.

You may also find it helpful to view the performance metrics provided by the Dynamic Monitoring Service (DMS). For more information, see the chapter about the Oracle Dynamic Monitoring Service in the *Oracle Fusion Middleware Performance and Tuning Guide*.

14.1.2 Introduction to the Oracle Fusion Middleware Diagnostic Framework

The Oracle Fusion Middleware Diagnostic Framework provides most of the technologies for creating and handling incidents for Oracle Fusion Applications. For information about the Oracle Fusion Middleware Diagnostic Framework, see the "Understanding the Diagnostic Framework" section in the *Oracle Fusion Middleware Administrator's Guide*.

14.1.3 Standard Incident Administration Tasks and Tools

The following administrative tasks are part of administering incidents:

- Adjusting configuration settings for incidents
- Using Oracle Enterprise Manager Fusion Applications Control (Fusion Applications Control) to monitor for incidents that are created automatically
- Creating incidents manually when needed
- Using incident information to resolve problems
- Adding additional diagnostic information (such as log file excerpts or screen shots) to existing incidents
- Packaging incident information and sending it to Oracle Support as part of obtaining assistance in resolving problems

The standard tools that you use to complete these tasks include:

- Fusion Applications Control user interface (for configuring incident settings)
- Log Viewer (part of Application Server Control, for viewing log files that you may add to incidents)
- `adrci` (Automatic Diagnostic Repository Command Interpreter, for viewing and packaging incidents, provided with each instance of the Automatic Diagnostic Repository (ADR) in Oracle Fusion Applications, Oracle Fusion Middleware, and Oracle Database)

For more information about using `adrci`, see [Section 14.4.2](#), the "ADCRI Command-Line Utility" section in the *Oracle Fusion Middleware Administrator's Guide*, and the "ADRCI: ADR Command Interpreter" chapter of *Oracle Database Utilities*.

- Oracle WebLogic Scripting Tool (WLST) commands related to incidents provided by the Oracle Fusion Middleware Diagnostic Framework.

For more information about these commands, see the "WLST Commands for Diagnostic Framework" section in the *Oracle Fusion Middleware Administrator's Guide*.

- SQL Trace (for assessing the efficiency of the SQL statements that Oracle Fusion applications are running)

For more information about SQL Trace, see [Section 12.2.4.6](#).

14.2 Configuring Log, QuickTrace, and Incident Settings for Troubleshooting

You may find it helpful to change log profile options or QuickTrace properties to gather more information during troubleshooting.

Ordinarily, the default configuration settings for incidents are appropriate for both normal operations and troubleshooting. It is possible to disable incident creation by using configuration settings, but it is recommended that you use the default incident configuration values unless Oracle Support specifically recommends changes to these settings.

This section contains the following topics:

- [Assisting Users in Gathering Data Using Troubleshooting Options](#)
- [Configuring the Oracle Fusion Middleware Diagnostic Framework](#)
- [Precedence and Log Settings for Troubleshooting](#)
- [Default System Settings for Incident Creation and QuickTrace](#)
- [Adjusting Incident Settings for Troubleshooting](#)
- [When Changes to Incident Configuration Settings Take Effect](#)
- [Adjusting QuickTrace Configuration Settings for Troubleshooting](#)
- [Adjusting Standard Log Levels for Troubleshooting](#)
- [Adjusting Log Settings for Oracle Fusion Incentive Compensation Batch Jobs for Troubleshooting](#)
- [Adjusting Oracle Fusion General Ledger Log Settings for Troubleshooting](#)
- [Adjusting Oracle Fusion Receivables AutoInvoice Log Settings for Troubleshooting](#)
- [Disabling Logging of a Particular Message](#)
- [Logging More Detailed Information for a Particular Code Module](#)
- [Configuring and Using Profile Options for Troubleshooting](#)

14.2.1 Assisting Users in Gathering Data Using Troubleshooting Options

Oracle Fusion Applications provides a Troubleshooting Options dialog box that you can display by choosing the **Troubleshooting Options** command in the Oracle Fusion Applications **Help > Troubleshooting** menu. The Troubleshooting Options dialog box provides ways to adjust the following kinds of settings for the current Oracle Fusion Applications user:

- Database trace (optionally capturing bind variables or wait events)
For more information about Database Trace, see the chapter about using application tracing tools in the *Oracle Database Performance Tuning Guide*.
- PL/SQL profiler
For more information about PL/SQL profiler, see the chapter about the PL/SQL hierarchical profiler in the *Oracle Database Advanced Application Developer's Guide*.
- Applications logging Severity Level
- Applications logging Modules

However, the **Help > Troubleshooting** menu displays the **Troubleshooting Options** command only for end users who have a job role that is mapped to the following three duty roles:

- Supportability Level Management Duty (CRM) (FND_SET_SUPPORTABILITY_LEVEL_DUTY_CRM)
- Supportability Level Management Duty (FSCM) (FND_SUPPORTABILITY_LEVEL_MANAGEMENT_DUTY_FSCM)
- Supportability Level Management Duty (HCM) (FND_SUPPORTABILITY_LEVEL_MANAGEMENT_DUTY_HCM)

By default, only the following job roles are matched to the Supportability Level Management Duty duty roles:

- Application Administrator (FND_APPLICATION_ADMINISTRATOR_JOB)
- Application Diagnostics Administrator (FND_DIAG_ADMINISTRATOR_JOB)

These two job roles have broad administrative privileges, so it is recommended that you assign them only to administrative personnel. To give other users access to the Troubleshooting Options dialog box, it is recommended that you find or create a job role that is specifically intended for troubleshooting and that is mapped to the three Supportability Level Management Duty duty roles. Once the troubleshooting job role is available, you can assign users to that job role as needed.

To grant a user access to the **Troubleshooting Options** command and dialog box and assist the user to gather data for troubleshooting:

1. Use Oracle Identity Manager to find or create a job role (also called an external role) that you can assign to users who need access to the **Troubleshooting Options** command and dialog box.

For more information, see the "Creating Roles" section in the *Oracle Fusion Middleware User's Guide for Oracle Identity Manager*.

2. If you created a job role in Step 1, use Oracle Authorization Policy Manager to map the job role to the following three duty roles:
 - Supportability Level Management Duty (CRM) (FND_SET_SUPPORTABILITY_LEVEL_DUTY_CRM)
 - Supportability Level Management Duty (FSCM) (FND_SUPPORTABILITY_LEVEL_MANAGEMENT_DUTY_FSCM)
 - Supportability Level Management Duty (HCM) (FND_SUPPORTABILITY_LEVEL_MANAGEMENT_DUTY_HCM)

Note: To make sure that the job role you create can be used for troubleshooting any Oracle Fusion application, be sure to map the job role to all three of the Supportability Level Management Duty duty roles. Each of these duty roles provides the needed functionality for a different application stripe, where an application stripe is a subset of policies in the policy store that is used by a particular application or group of applications.

For more information about mapping job roles to duty roles, see the "Mapping External Roles to an Application Role" section in the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications)*

Edition). A job role corresponds to an external role, and a duty role corresponds to an application role.

3. For any user who currently needs access to the Troubleshooting Options dialog box, use Oracle Identity Manager to assign the job role that you found or created in Step 1 to the user.

For more information, see the "Assigning Members to a Role" section or the "Adding and Removing Roles" section in the *Oracle Fusion Middleware User's Guide for Oracle Identity Manager*.

4. To make sure that the user has access to the new job role assignment, direct the user to log out from their Oracle Fusion application and log back in.
5. Direct the user to choose **Troubleshooting Options** from the **Help > Troubleshooting** menu and to modify **Troubleshooting Options** settings as needed for the specific problem you are troubleshooting:

Note: If the user does not see **Troubleshooting Options** in the **Help > Troubleshooting** menu and if you performed Step 2 of this procedure within the last few minutes, direct the user to wait a few more minutes and then log out and log in again. When you map a job role to one or more duty roles, the change may take effect a few minutes later, after Oracle Authorization Policy Manager refreshes the server cache.

- a. If you want to enable all of the available troubleshooting options, direct the user to select the **Enable all** checkbox.
- b. If you want to enable Database Trace, direct the user to select the **Database trace** checkbox.

For more information about Database Trace, including the options of capturing bind variables or wait events, see the chapter about application tracing tools in the *Oracle Database Performance Tuning Guide*.

- c. If you want to enable the PL/SQL profiler, direct the user to select **PL/SQL profiler**.

For more information, see the section about using the PL/SQL hierarchical profiler in the *Oracle Database Advanced Application Developer's Guide*.

- d. If you want to log information for the user's actions at a level of detail that is different from the logging done for other users at the same site, direct the user to select the **Applications logging** checkbox and to select the appropriate level of detail from the **Severity Level** drop-down list.

This user action is equivalent to an administrator setting the value of the `AFLOG_LEVEL` profile option for that user. For more information about the available logging severity levels, see [Section 13.7.3](#).

- e. If you want to limit the logging for the user's actions to operations that use one or more particular code modules, and if you know the name of those code modules, direct the user to enter the module names in the **Modules** field, separating the names using commas without spaces.

The user can use the % wildcard character to represent part of a module name or to represent all module names.

- f. Direct the user to click **OK**.

Note: Users who change the **Applications logging Severity Level** may need to log out from their Oracle Fusion application and log in again in order to have the changes take effect. It is recommended that you direct the user to do so at this point in the process.

6. Direct the user to try to reproduce the problem.
7. Verify that the additional information you wanted was collected.
8. Unless the user needs ongoing access to the Troubleshooting Options dialog box, use Oracle Identity Manager to revoke the user's assignment to the job role that you found or created in Step 1.

14.2.2 Configuring the Oracle Fusion Middleware Diagnostic Framework

As part of the infrastructure for creating and handling incidents, the Oracle Fusion Middleware Diagnostics Framework includes settings that you can configure to affect functionality areas such as the following:

- Whether multiple incidents that originate from Java or PL/SQL code can be created for the same problem within a particular time period, and, if so, how many (flood control)
- The total amount of space allocated for storage of information about incidents that originate from Java or PL/SQL code (older incidents are purged automatically when the allocated space limit is reached)

For more information about these Oracle Fusion Middleware Diagnostic Framework settings, see the "Configuring the Diagnostic Framework" section in the *Oracle Fusion Middleware Administrator's Guide*. For information about flood control and disk space management for incidents that originate from PL/SQL code, see [Table 14-1](#).

14.2.3 Precedence and Log Settings for Troubleshooting

During troubleshooting in a production environment, you will normally use profile options to increase or decrease the amount of information gathered in standard Oracle Fusion Applications log files. Where feasible, you will change profile option values at the `USER` level, rather than at the `SITE` level. This allows you to log more detailed information for a specific user who attempts to reproduce a problem, without cluttering your log file with unnecessarily detailed information for other users.

For Oracle Fusion applications that are written in Java or SOA, it is also possible for the `odlLevel` setting for the `oracle.apps` logger in the `logging.xml` configuration file to affect the level of detail that is logged. If both the `odlLevel` setting and the `AFLOG_LEVEL` profile option have defined values, then the minimum of those two values takes precedence.

Note: Unless Oracle Support advises you to change the value of the setting, it is strongly recommended that you keep your `odlLevel` setting for the `oracle.apps` logger set to a value of `All`.

14.2.4 Default System Settings for Incident Creation and QuickTrace

The default profile option configuration settings for incident creation and QuickTrace profile options are shown in [Table 14-1](#). For information about additional QuickTrace properties, see [Section 14.2.7](#).

Table 14–1 Profile Options for Oracle Fusion Applications Incidents and QuickTrace

Profile Option Name (and Display Names)	Environment	Description	Possible Values or Example	Applicable Profile Hierarchy Levels	Default Value
AFLOG_INCIDENT_ENABLED (Incident Enabled)	Java, PL/SQL, C, and SOA	Enables or disables incident creation for Oracle Fusion Applications. (This profile option does not affect Oracle Fusion Middleware incident creation functionality.) A value of Y allows incident creation at runtime. A value of N prevents incident creation. This profile option is available for <code>SITE</code> but is not available for <code>USER</code> . <code>AFLOG_INCIDENT_ENABLED</code> is a new profile that has been introduced in Fusion to support application incident creation.	Y, N	Site	Y
AFLOG_PLSQL_FLOOD_CONTROL_ENABLED (FND: Incident Flood Control Enabled)	PL/SQL	Determines whether to restrict the collection of diagnostics if multiple PL/SQL incidents occur within a set time period for the same problem. Specify Y to restrict diagnostic collection and N to leave it unrestricted.	Y, N	Site	Y
AFLOG_PLSQL_FLOOD_CONTROL_INC_COUNT (FND: Total number of incidents for problem key)	PL/SQL	When <code>AFLOG_PLSQL_FLOOD_CONTROL_ENABLED</code> is set to Y, the value of this setting determines the number of PL/SQL incidents with the same problem key for which diagnostics are collected during a set time period. If the number of incidents with the same problem key exceeds the value of this setting during the time period specified by <code>AFLOG_PLSQL_FLOOD_CONTROL_INC_TIMEPERIOD</code> , an incident is created, but no diagnostics are captured	Positive integers	Site	5
AFLOG_PLSQL_FLOOD_CONTROL_INC_TIMEPERIOD (FND: Flood control time period)	PL/SQL	When <code>AFLOG_PLSQL_FLOOD_CONTROL_ENABLED</code> is set to Y, the value of this setting specifies the time period in minutes during which diagnostics are not collected if the number of PL/SQL incidents that occur for the same problem exceeds the value of <code>AFLOG_PLSQL_FLOOD_CONTROL_INC_COUNT</code> .	Positive integers	Site	60

Table 14–1 (Cont.) Profile Options for Oracle Fusion Applications Incidents and QuickTrace

Profile Option Name (and Display Names)	Environment	Description	Possible Values or Example	Applicable Profile Hierarchy Levels	Default Value
AFLOG_PLSQL_MAX_TOTAL_INC_SIZE (FND: Total size of incident files)	PL/SQL	<p>Sets the amount of disk space in megabytes that is allocated for PL/SQL incident data storage. When incidents use more than this amount of disk space, the oldest incident that is no longer needed is automatically purged.</p> <p>Whether or not an older incident is still needed is determined by checking the incident's status in the Automatic Diagnostic Repository (ADR). If the status of the incident is 3, or KEEP, or tracking (in the Automatic Diagnostic Repository Command Interpreter, adrci), then that incident is not purged.</p> <p>If the allocated disk space is full and if all stored incidents are still needed, a new, empty incident directory called <code>inc_incident_ID_number</code> is created, and log file messages may refer to that location, but diagnostic log excerpts are not saved to that incident directory.</p>	Positive integers	Site	500 (MB)
AFLOG_QUICKTRACE_ENABLED (QuickTrace Enabled)	Java and SOA	Enables quick trace to occur at runtime.	Y, N	Site	Y
AFLOG_QUICKTRACE_LEVEL (QuickTrace Level)	Java and SOA	<p>Specifies the minimum level for tracing.</p> <p>This setting is similar to the <code>AFLOG_LEVEL</code> profile option for standard Oracle Fusion Applications logging to log files, but the two settings operate independently.</p>	1000 (SEVERE), 900 (WARNING), 800 (INFO), 700 (CONFIG), 500 (FINE), 400 (FINER), 300 (FINEST) and 0 (OFF)	Site	500 (FINE)

14.2.5 Adjusting Incident Settings for Troubleshooting

By default, incident creation functionality for Oracle Fusion Applications is enabled, and it should ordinarily remain enabled. However, if Oracle Support directs you to disable incident creation, you can use the `AFLOG_INCIDENT_ENABLED` profile option to do so. You can use the Manage Administrator Profile Values task in the Setup and Maintenance work area to set the value of this profile option. For more information about this setting, see [Section 14.2.4](#).

Incident creation functionality is also enabled by default for the Oracle Fusion Middleware Diagnostic framework. In addition, various Oracle Fusion Middleware profile options govern the details of how incidents are collected and stored. In general, troubleshooting will not require you to change these settings. For more information, see the "Configuring the Diagnostic Framework" section in the *Oracle Fusion Middleware Administrator's Guide*.

14.2.6 When Changes to Incident Configuration Settings Take Effect

Changes to incident profile options ordinarily take effect as follows:

- **User Sessions:** Changes to site level incident profile options take effect for any user session that is started after the setting has been changed. Changes to site level log file profile options do not affect any user sessions that are already in progress when the change is made.
- **PL/SQL and C processes such as scheduled jobs:** Changes to site level incident profile options take effect for any C or PL/SQL processes that are launched after the setting has been changed. Incident profile option changes do not affect C or PL/SQL processes that are already running.

14.2.7 Adjusting QuickTrace Configuration Settings for Troubleshooting

By default, QuickTrace functionality for Oracle Fusion Applications is enabled at the FINE logging level. QuickTrace should ordinarily remain enabled. However, if Oracle Support directs you to disable QuickTrace, or if you want to change the amount or organization of the information that QuickTrace collects, you can adjust the values of the following profile options either by using Fusion Applications Control or by using the Manage Administrator Profile Values task in the Setup and Maintenance work area:

- **QuickTrace Enabled** (AFLOG_QUICKTRACE_ENABLED profile option): Governs whether or not QuickTrace gathers any information.
- **QuickTrace Level** (AFLOG_QUICKTRACE_LEVEL profile option): Specifies the level of detail at which QuickTrace gathers information.

Note: Because the QuickTrace Enabled and QuickTrace Level profile options apply at the site level, either all of the site's Managed Servers capture QuickTrace information in memory buffers or none of the Managed Servers do so. Similarly, the QuickTrace Level value sets the level of detail that QuickTrace gathers for all of the site's Managed Servers.

For more information about these settings, see [Section 14.2.4](#).

In addition, [Table 14–2](#) shows QuickTrace property settings that you can change using Fusion Applications Control.

Table 14–2 QuickTrace Properties

Property Name	Description	Server Restart Required	Default
Buffer Size (bufferSize)	<p>Approximate size in bytes of each circular QuickTrace buffer. Each buffer's actual memory consumption may be less but not more than the specified value.</p> <p>By default, each WebLogic Server has a single <code>COMMON</code> QuickTrace buffer for storing multiple users' log records in memory.</p> <p>If you want store particular users' log records in individual buffers (one buffer for each user you specify), then you can allocate additional buffers of the same size by setting <code>Enable User Buffer</code> to true, restarting the server, and setting the value of <code>User Names for Reserve Buffer</code> to a comma-separated list of those users.</p>	Yes	5242880
Enable User Buffer (enableUserBuffer)	<p>When this property is set to false, each QuickTrace handler uses a single <code>COMMON</code> buffer to cache log messages for all users.</p> <p>When this property is set to true, each QuickTrace handler maintains an individual buffer for each user who is listed in the <code>User Names for Reserve Buffer</code> (<code>reserveBufferUserIDs</code>) property. Log messages for users who are not listed in that property are cached in the <code>COMMON</code> buffer.</p> <p>To determine in which buffer to place a particular log message, QuickTrace examines the log message for an application User Name value. If the log message does not include a User Name value, QuickTrace examines it for a User ID value provided by the server. If the log message does not include either a User Name or a User ID, then QuickTrace caches the message in the <code>COMMON</code> buffer.</p>	Yes	false
User Names for Reserve Buffer (reserveBufferUserIDs)	<p>Holds a list of application User Name or server User ID values separated by comma. When the <code>Enable User Buffer</code> (<code>enableUserBuffer</code>) property is set to true, for each listed application User Name or server User ID, any log messages associated with that user are cached in an individual QuickTrace user buffer. Log messages that are associated with other users are cached in the <code>COMMON</code> buffer.</p> <p>When the <code>Enable User Buffer</code> (<code>enableUserBuffer</code>) property is set to false, this property has no effect.</p>	No	null

You can also change the properties in [Table 14–2](#) by using the Oracle WebLogic Scripting Tool (WLST) `configureLogHandler` command. For more information, see the "configureLogHandler" section in the *Oracle Fusion Middleware WebLogic Scripting Tool Command Reference*.

To change QuickTrace profile option and property settings using Fusion Applications Control:

1. From the navigation pane, select a product family target such as **Financials**.

2. In the context pane, from the target type dropdown menu, choose **Logs > Log Configuration**.
3. In the Logging Profile Configuration dialog box, click the **Site-Level** tab.
4. Expand the **Advanced Configuration** area of the dialog box.
5. Expand the **Java Settings** portion of the **Advanced Configuration** area.
6. Expand the **Other QuickTrace Settings** portion of the **Java Settings** area.
7. In the **Java Settings** portion of the **Advanced Configuration** area, change the QuickTrace profile settings, as needed:
 - a. To disable QuickTrace, clear the **QuickTrace Enabled** checkbox. To enable QuickTrace, make sure the **QuickTrace Enabled** checkbox is selected.
 - b. To change the amount of information QuickTrace gathers, select a logging level from the **QuickTrace Level** drop-down list.
8. In the **Other QuickTrace Settings** portion of the **Java Settings** area, change the QuickTrace property settings, as needed:
 - a. If you intend to use any QuickTrace property value that is not the default value listed in [Table 14-2](#), select the **Override Settings** checkbox.
 - b. In the **Buffer Size** field, enter the desired QuickTrace buffer size in bytes.

Note: If you increase the value of the **Buffer Size** setting, QuickTrace uses more memory, which may affect system performance.

- c. If you want QuickTrace to store information in separate buffers for selected individual users, select the **Enable User Buffer** checkbox.
- d. In the **User Names for Reserve Buffer** field, enter comma-separated Oracle Fusion Applications user names for all of the users for whom QuickTrace should store information in individual buffers.

Note: The values in the **User Names for Reserve Buffer** field are used only if the **Enable User Buffer** checkbox is selected.

9. Click **Apply**.

Note: Changes to **QuickTrace Enabled** and **QuickTrace Level** settings take effect immediately for the user session in which the changes were made, but the changes do not affect other user sessions that are in progress when the changes are made. To use the changed settings for a different user session, log out from your Oracle Fusion application and log back in.

For each particular Managed Server that runs an Oracle Fusion application, changes to the QuickTrace properties **Buffer Size** and **Enable User Buffer** do not take effect until the server is restarted.

If a value of `true` is already in effect for the **Enable User Buffer** property, then changes to the **User Names for Reserve Buffer** property take effect as soon as you click **Apply**. Otherwise, you must set the **Enable User Buffer** property to `true` and restart the server in order to have changes to the **User Names for Reserve Buffer** property take effect.

14.2.8 Adjusting Standard Log Levels for Troubleshooting

If you encounter a situation in which an incident does not provide enough information to allow you to resolve the problem, then you may need to increase the standard logging level for the relevant Oracle Fusion application and the application user who will attempt to reproduce the problem. For information about the available log levels, see [Section 13.7.3](#). For information about setting the logging profile options, see [Section 13.7.5](#).

14.2.9 Adjusting Log Settings for Oracle Fusion Incentive Compensation Batch Jobs for Troubleshooting

The logging functionality for certain Oracle Fusion Incentive Compensation batch jobs is separate from the standard logging functionality for Oracle Fusion Applications. By default, the following kinds of Oracle Fusion Incentive Compensation batch jobs do not write log entries:

- Calculation
- Classification
- Collection
- Crediting
- Rollup

For optimum performance and log file sizes, it is recommended that you use the logging functionality for these areas only when troubleshooting an existing problem.

To configure Oracle Fusion Incentive Compensation batch jobs to write nonstandard log entries:

1. Navigate to the **Manage Messages** task flow in the **Setup and Maintenance** work area.
2. In the **Implementations** section of the **Tasks** pane, click **Manage Implementation Projects**.
3. In the **Manage Implementation Projects** pane, use standard Oracle search techniques to locate the implementation project that includes Oracle Fusion Incentive Compensation, then click the project name in the search results.

4. In the **Tasks Lists and Tasks** table, expand the **Define Incentive Compensation** entry.
5. Expand the **Define Incentive Compensation Shared Configuration** entry.
6. Locate the `Manage Incentive Compensation Profile Options` row of the table and click the **Go to Task** icon.
7. In the **Profiles** section of the **Tasks** pane, click **Manage Administrator Profile Values**.
8. In the **Search** area of the **Manage Profiles** pane, select `Incentive Compensation` from the **Application** dropdown list and click **Search**.
9. In the **Results** table, locate and select the row for the `CN_DEBUG` profile option.
10. If the **Enable Debug Mode Profile Values** table contains a record for the **Profile Level** and **Level Value** for which you want to collect log entries, select the record and click **Edit**; otherwise, click **New** and fill in the appropriate values, as follows:
 - If you want to collect log entries for a specific user's Oracle Fusion Incentive Compensation batch jobs, you would edit or create a record in which **Profile Level** is set to `User` and **Level Value** is set to the user name of the specific user. Set **Profile Value** to `Y`.
 - If you want to collect log entries for all of the Oracle Fusion Incentive Compensation batch jobs at the site, you would edit or create a record in which **Profile Level** is set to `Site` and **Level Value** is blank. Set **Profile Value** to `Y`.
11. Click **Save**.
12. Start a new user session of Oracle Fusion Incentive Compensation and try to replicate the problem.

For information about how to view the log entries, see [Section 14.5.2](#).

14.2.10 Adjusting Oracle Fusion General Ledger Log Settings for Troubleshooting

In the Oracle Fusion Financials product family, some logging functionality for the Oracle Fusion General Ledger application is separate from the standard logging functionality for Oracle Fusion Applications. By default, the following kinds of Oracle Fusion General Ledger batch jobs do not write log entries:

- `OpenPeriod`
- `Posting`
- `Translation`
- `Close Process - Create Income Statement Closing Journals`
- `Close Process - Create Balance Sheet Closing Journals`

For optimum performance and log file sizes, it is recommended that you use the logging functionality for these areas only when troubleshooting an existing problem.

To configure the Oracle Fusion General Ledger application to write nonstandard log entries:

1. From the **Setup and Maintenance** work area, navigate to the **Manage Administrator Profile Values** screen.
2. From the **Application** drop-down list in the **Search: Profile Option** pane, select **General Ledger**.

3. In the **Profile Display Name** field, enter %Debug% and click **Search**.
4. In the **Search Results** table, select `GL_DEBUG_MODE`.
5. If the **GL_DEBUG_MODE: Profile Values** table contains a record for the **Profile Level** and **Level Value** for which you want to collect log entries, select the record and click **Edit**; otherwise, click **New**.
6. Fill in the appropriate values, as follows:
 - If you want to collect log entries for a specific user's Oracle Fusion General Ledger batch jobs, edit or create a record in which **Profile Level** is set to `User` and **Level Value** is set to the user name of the specific user. Set **Profile Value** to `Y`.
 - If you want to collect log entries for all of the Oracle Fusion General Ledger batch jobs at the site, edit or create a record in which **Profile Level** is set to `Site` and **Level Value** is blank. Set **Profile Value** to `Y`.
7. Click **Save**.
8. Ask an appropriate user to start a new user session of Oracle Fusion General Ledger and try to replicate the problem.
9. Inspect the relevant log file for new entries.

If the value of the `AFLOG_FILENAME` profile option is set, then the value of the `AFLOG_FILENAME` profile option indicates the location of the log file.

If the value of the `AFLOG_FILENAME` profile option is not set, then the information is logged to the location designated for Oracle Enterprise Scheduler log files. For information about viewing Oracle Fusion Middleware log files, including log files for Oracle Enterprise Scheduler, see the "Viewing and Searching Log Files" section in the *Oracle Fusion Middleware Administrator's Guide*.

10. To help manage the amount of disk space used by the logged information, set the value of `GL_DEBUG_MODE` to `N` or delete any row of the **GL_DEBUG_MODE: Profile Values** table that you no longer need.

For more information about managing the amount of disk space used by Oracle Fusion General Ledger log entries, see [Section 13.7.4.3](#).

14.2.11 Adjusting Oracle Fusion Receivables AutoInvoice Log Settings for Troubleshooting

In the Oracle Fusion Financials product family, logging functionality for the AutoInvoice portion of the Oracle Fusion Receivables application is separate from the standard logging functionality for Oracle Fusion Applications. The amount of information that is logged for AutoInvoice depends on the value of the Log File Message Level system option setting for each business unit. The information is placed in the standard log file for Oracle Enterprise Scheduler.

Note: For optimum performance and log file sizes, it is recommended that you keep the Log File Message Level system option set to the lowest value that meets your everyday needs, and increase the value of the setting only when troubleshooting an existing problem.

The available values for the Log File Message Level setting are the integers 0 through 5. The value of 0 is recommended for day-to-day business needs. The value of 5 provides the maximum amount of information for troubleshooting.

Specifically, if you set Log File Message Level to a value of 0, the following kinds of information are logged:

- Product Version
- Program Name
- AutoInvoice Start Time
- AutoInvoice Concurrent Request Arguments
- Error and Warning Messages
- AutoInvoice End Time
- AutoInvoice Logical Steps

If you set Log File Message Level to a value of 1, the following additional information is gathered:

- Time-Stamped function labels

If you set Log File Message Level to a value of 2, the following additional information is gathered:

- Sizes of Allocated Arrays
- Dynamic SQL Statements
- Number of Rows Updated, Inserted and Deleted

If you set Log File Message Level to a value of 3, the following additional information is gathered:

- Method IV SQL Array Values

If you set Log File Message Level to a value of 4, the following additional information is gathered:

- Values of all variables that are used to call FND or Tax routines

If you set Log File Message Level to a value of 5, the following additional information is gathered:

- Bad lines
- Rejected lines

To change the amount of information that the Oracle Fusion Receivables application logs for the AutoInvoice functionality area in the current business unit:

1. In the Oracle Fusion Receivables application, select **Setup and Maintenance** from the Navigator menu.
2. Complete the following substeps to navigate to the **Edit System Options** screen for the **Manage Receivables System Options** task.
 - a. Expand the **Search: Tasks** pane.
 - b. Enter `Manage Receivables System Options` in the **Name** field and click **Search**.
 - c. In the `Manage Receivables System Options` row of the **Search Results** table, click **Go to Task**.

- d. In the **Search** area of the **Manage Receivables System Options** screen, select `Business Unit` from the drop-down list, enter the name of the business unit for which you are troubleshooting, and then click **Search**.
Alternately, you can click **Search** without specifying a business unit name to display a list of the business units you can choose from.
- e. In the **Search Results** table, click the name of the business unit for which you are troubleshooting.
3. In the **Edit System Options** screen, scroll down to display the **AutoInvoice** area of the screen, and then set **Log File Message Level** to the value that corresponds to the amount of information you want to gather.
For example, to gather the maximum amount of information for troubleshooting, set the value to 5. To gather the normal amount of information for day-to-day operations, set the value to 0.
4. Click **Save** to put the change into effect.

Log entries for the AutoInvoice portion of Oracle Fusion Receivables are placed in the standard log file for Oracle Enterprise Scheduler. For information about viewing Oracle Fusion Middleware log files, including log files for Oracle Enterprise Scheduler, see the "Viewing and Searching Log Files" section in the *Oracle Fusion Middleware Administrator's Guide*.

Note: After you have gathered the detailed information that you need, remember to repeat this procedure and return the **Log File Message Level** setting to the value that meets your day-to-day information needs.

14.2.12 Disabling Logging of a Particular Message

You may occasionally encounter a situation in which one of your Oracle Fusion applications repeatedly logs the same error message. In this situation, it is best to discover and correct the cause of the error. However, on rare occasions, Oracle Support may recommend that you use the `FND_MESSAGES.LOGGABLE_ALERTABLE` attribute to prevent the repeating message from being logged while efforts are underway to resolve the issue.

Caution: Do not change the `FND_MESSAGES.LOGGABLE_ALERTABLE` attribute unless Oracle specifically advises you to do so, as this involves permanent loss of data.

Do not change the `FND_MESSAGES.LOGGABLE_ALERTABLE` attribute directly in the database. Instead, use the Manage Messages task flow.

To disable logging for a particular message:

1. Navigate to the **Manage Messages** task flow in the **Setup and Maintenance** work area.
2. In the **Message Name** field in the **Search** pane, enter the name of the error message that you want to temporarily stop logging.
3. Select the relevant Oracle Fusion application from the **Application** drop-down list.
4. If you know it, select the relevant code module name from the **Module** drop-down list.

5. Click **Search**.
6. Complete the following substeps to set the value of the `FND_MESSAGES.LOGGABLE_ALERTTABLE` attribute to `N` for the error message that you want to temporarily stop logging:
 - a. In the search results table, select the error message that you want to temporarily stop logging.
 - b. Click the edit icon in the table header.
 - c. In the **Edit Message** screen, clear the **Logging Enabled** checkbox.
 - d. Click **Save**, **Save and Create Another**, or **Save and Close** to put the change into effect immediately:

14.2.13 Logging More Detailed Information for a Particular Code Module

If you contact Oracle Support for assistance in resolving a problem, the support representative may ask you to gather more detailed log information on the operations of one or more specific code modules in your Oracle Fusion applications. You can do this by setting the value of the `AFLOG_MODULE` profile option to one or more module names that Oracle Support specifies.

Caution: Unless Oracle Support specifically advises you to change the value of the `AFLOG_MODULE` profile option at the `Site` level, be sure to adjust the value only for a particular user who will attempt to reproduce the problem you are investigating. Specifying one or more particular modules at the `Site` level prevents logging on any modules that are not specified.

For more information about using administrative screens to set the value of this profile option either for a specific user or at the `Site` level, see [Section 13.7.5](#). For more information about having a specific user set the user's own profile option using **Troubleshooting Options**, see [Section 14.2.1](#).

14.2.14 Configuring and Using Profile Options for Troubleshooting

You may find it useful to adjust one or more of the following profile options during troubleshooting:

- To log more detailed information about a particular user's operations, you can create or change the value of the `AFLOG_LEVEL` profile option for that user. For more information about setting this profile option, see [Section 13.7](#).
- To log more detailed information for an entire site, you can create or change the value of the `AFLOG_LEVEL` profile option for the site. For more information about setting this profile option, see [Section 13.7](#).
- To specify the Oracle Fusion Applications code modules for which standard log entries will be written at the level of detail specified by the `AFLOG_LEVEL` profile option for a particular user, you can set the `AFLOG_MODULE` profile option for the same user. (To minimize the number of unrelated log entries that will be placed in the log file during troubleshooting, it is recommended that you specify a particular module whenever feasible.) For more information about setting the `AFLOG_MODULE` profile option, see [Section 13.7](#).

- To turn on nonstandard logging for troubleshooting Oracle Fusion Incentive Compensation, you can set the value of the application-specific `CN_DEBUG` profile option to `Y`, either for a single user or for an entire site. For more information about setting this profile option, see [Section 14.2.9](#).

14.3 Configuring the Diagnostic Testing Framework for Troubleshooting

When you investigate an Oracle Fusion Applications problem, you may want to assign job roles to some users, temporarily, that will allow those users to run diagnostic tests. For information about how to do this, see [Section 13.9.1](#).

There is generally no need to change other configuration settings in order to use the Oracle Fusion Applications Diagnostic Testing Framework for troubleshooting. When Java or SOA error-handling code in an Oracle Fusion application creates an incident, the Diagnostic Testing Framework automatically runs any Oracle Fusion Applications diagnostic tests that are associated with the particular error message that is logged. There is no configuration setting in the Diagnostic Testing Framework for turning off this automatic execution of diagnostic tests.

Note: (Oracle Fusion Applications developers use the `APPS_MSG_ID` diagnostic test tag to associate particular error messages with appropriate diagnostic tests. You cannot delete or modify the values of diagnostic test tags that are supplied by Oracle.)

14.4 Investigating, Reporting, and Solving a Problem

From time to time, you may receive a user complaint or an automatically generated incident notification about a problem with a particular Oracle Fusion application. For any problem that does not have an immediate and obvious solution, you can diagnose and resolve the problem more effectively by becoming familiar with incident functionality and by using a systematic process for your troubleshooting process.

This section contains the following topics:

- [Process for Investigating, Reporting, and Solving a Problem](#)
- [Managing Oracle Fusion Applications Problems and Incidents](#)
- [Working with Automatically Created QuickTrace Dumps](#)
- [Working with Manually Created QuickTrace Dumps](#)
- [Working with Other Diagnostic Dumps](#)

14.4.1 Process for Investigating, Reporting, and Solving a Problem

When you become aware of a problem with a particular Oracle Fusion application, you can use a systematic process to help diagnose and address the problem. While any problem may have unique features, the information that you gather systematically is likely to be useful during the troubleshooting process. The recommended process varies somewhat depending on whether or not your Oracle Fusion Applications environment includes Oracle Enterprise Manager Cloud Control.

This section contains the following topics:

- [Process for Investigating, Reporting, and Solving a Problem Using Cloud Control](#)
- [General Process for Investigating, Reporting, and Solving a Problem](#)

14.4.1.1 Process for Investigating, Reporting, and Solving a Problem Using Cloud Control

If your Oracle Fusion Applications environment includes Cloud Control, you can use the **Problem Analysis** and **Analyze Logs** options in Cloud Control to help you focus your investigation on the log messages and metrics that are most likely to help you determine the cause of the problem. You can also use **Problem Analysis** to make quick assessments of the health of your system, when no problems have been reported. For more information about using **Problem Analysis** for general monitoring, see [Chapter 10](#).

To help diagnose and address a problem in an Oracle Fusion Applications environment that includes Cloud Control, use the following steps:

1. If an Oracle Fusion application user reports an error message to you, make a note of any specific information that the error message may have provided, such as:
 - Incident ID
 - Server Domain
 - Server Instance
 - Application Name
2. If there is an alert for a particular metric, then use the **Problem Analysis** option in Cloud Control to scan performance metric charts, target status tables, and relevant log portions for indications of problem causes.

For more information about using the **Problem Analysis** option, see [Section 14.4.6](#).

3. If no incident was generated, create an incident manually.

Using incidents for all problems helps you organize the information that you gather during the troubleshooting process. For information about creating an incident manually, see [Section 14.4.2.4](#).

4. If the cause of the problem is still uncertain, manually perform any relevant diagnostic dumps, assigning the dump results to the incident you created.
5. Examine any incident information that has been collected, such as any QuickTrace dumps or diagnostic test results that are associated with the incident, whether those test results and dumps were collected automatically or manually.

For more information about working with QuickTrace dumps, see [Section 14.4.3](#).

You can view the results of diagnostic tests that are associated with an incident either from the Incident Detail page in Support Workbench or from the Diagnostic Dashboard. For more information about viewing incident details from Support Workbench, see [Section 14.4.2.1](#). For more information about using the Diagnostic Dashboard to view diagnostic test results, see [Section 13.10.9.1](#).

6. If the cause of the problem is still uncertain, view recent information in the standard log files for any of the following that may be relevant:
 - Oracle Fusion Applications
 - Oracle Fusion Middleware
 - Oracle Business Intelligence
 - Oracle Enterprise Scheduler Service
 - Oracle Service-Oriented Architecture (SOA) Suite

For more information about locating log files for Oracle Fusion Applications, see [Section 13.6.4](#). For more information about locating log files for Oracle Fusion Middleware components, see the "Log File Location" table in the *Oracle Fusion Middleware Administrator's Guide*. For more information about locating log files for Oracle Enterprise Scheduler, see [Section 5.17](#), [Section 13.7.4.3](#), or [Section 14.5.3.2](#).

7. If you want to run more diagnostic tests for the relevant Oracle Fusion application or attach the results of such manually-run diagnostic tests to the incident, navigate from Cloud Control to the Diagnostic Dashboard to perform those tasks.

For information about navigating from Cloud Control to the Diagnostic Dashboard, see [Section 13.9.2.2](#).

For information about using the Diagnostic Dashboard to run diagnostic tests, see [Section 13.10](#).

For information about the individual diagnostic tests that are provided with this release, see the *Oracle Fusion Applications Common User Guide* in Oracle Fusion Applications Help.

For information about adding diagnostic test results to an incident, see [Section 14.6.3](#).

8. If necessary, change the setting that governs the amount of information to be logged, try to replicate the problem, and inspect the newly logged information.

The steps for increasing the amount of information to be logged may vary depending on the application involved:

- For most Oracle Fusion applications, standardized logging mechanisms place information in log files that you can view from Fusion Applications Control. You can increase the amount of information to be gathered for a particular user or for a whole site. For more information, see [Section 14.2.1](#) and [Section 13.7.5](#).
 - If the problem occurred in an application that uses nonstandard mechanisms for logging, such as certain Oracle Fusion Incentive Compensation batch jobs, certain Oracle Fusion General Ledger batch jobs, or the AutoInvoice portion of the Oracle Fusion Receivables application, then you may need to turn on a logging facility that is normally disabled and specify the kind of information you want to gather. For more information, see [Section 13.7.6](#).
9. If the cause of the problem is still uncertain, click **My Oracle Support Knowledge** in the **Problem Details** or **Incident Details** page and log in to the My Oracle Support knowledge base to search for related information and solutions.
 10. If you have not succeeded in resolving the problem, use Support Workbench to package the problem or incident along with all related dump files and diagnostic reports and then send the package to Oracle Support as a Service Request.

For more information, see [Section 14.4.2.5](#) and the topic "Problem Details Page" in the Cloud Control online help.

14.4.1.2 General Process for Investigating, Reporting, and Solving a Problem

To help diagnose and address a problem in an Oracle Fusion Applications environment that does not include Cloud Control, use the following steps:

1. If an Oracle Fusion application user reports an error message to you, make a note of any specific information that the error message may have provided, such as:
 - Incident ID

- Server Domain
 - Server Instance
 - Application Name
2. If no incident was generated, create an incident manually and then manually perform any relevant diagnostic dumps, assigning the dump results to the incident you created.

Using incidents for all problems helps you organize the information that you gather during the troubleshooting process. For information about creating an incident manually, see [Section 14.4.2.6](#).

3. Examine any information that was collected and stored in the incident directory, such as QuickTrace dumps and the results of any diagnostic tests that may have been run, whether the tests and dumps were run automatically or were created manually.

For more information about incident directory locations, see [Section 14.4.2](#). For more information about how to work with QuickTrace dumps, see [Section 14.4.3](#).

4. If the cause of the problem is still uncertain, view recent information in the standard log files for any of the following that may be relevant:

- Oracle Fusion Applications
- Oracle Fusion Middleware components
- Oracle Business Intelligence
- Oracle Enterprise Scheduler Service
- Oracle Service-Oriented Architecture (SOA) Suite

For more information about locating log files for Oracle Fusion Applications, see [Section 13.6.4](#). For more information about locating log files for Oracle Fusion Middleware components, see the "Log File Location" table in the *Oracle Fusion Middleware Administrator's Guide*. For more information about locating log files for Oracle Enterprise Scheduler, see [Section 5.17](#), [Section 13.7.4.3](#), or [Section 14.5.3.2](#).

5. Locate diagnostic tests that are related to the feature in which the problem occurred, run the diagnostic tests that seem most likely to be relevant, and inspect the results.

For information about starting and using the Diagnostic Dashboard, see [Section 13.9.2](#) and [Section 13.10](#). For information about the individual diagnostic tests that are provided with this release, see the *Oracle Fusion Applications Common User Guide* in Oracle Fusion Applications Help.

6. If necessary, change the setting that governs the amount of information to be logged, try to replicate the problem, and inspect the newly logged information.

The steps for increasing the amount of information to be logged may vary depending on the application involved:

- For most Oracle Fusion applications, standardized logging mechanisms place information in log files that you can view from Fusion Applications Control. You can increase the amount of information to be gathered for a particular user or for a whole site. For more information, see [Section 14.2.1](#) and [Section 13.7.5](#).
- If the problem occurred in an application that uses nonstandard mechanisms for logging, such as certain Oracle Fusion Incentive Compensation batch jobs, certain Oracle Fusion General Ledger batch jobs, or the AutoInvoice portion of

the Oracle Fusion Receivables application, then you may need to turn on a logging facility that is normally disabled and specify the kind of information you want to gather. For more information, see [Section 13.7.6](#).

7. If the cause of the problem is still uncertain, search the My Oracle Support knowledge base for related information and solutions.
8. If you have not succeeded in resolving the problem, use `adrci` to package the problem or incident along with all related dump files and diagnostic reports and then send the package to Oracle Support as a Service Request.

For more information about adding files to an incident and packaging the incident into a zip file, see the "Packaging an Incident" section in the *Oracle Fusion Middleware Administrator's Guide*.

14.4.2 Managing Oracle Fusion Applications Problems and Incidents

All Oracle Fusion Applications customers can use the Automatic Diagnostic Repository Command Interpreter (`adrci`) to work with incidents. However, if your Oracle Fusion Applications environment includes Cloud Control, it is recommended that you use Support Workbench to work with incidents, as it provides more convenient ways to perform operations that otherwise require knowledge of specific (`adrci`) commands.

This section contains the following topics:

- [Viewing Incidents Using Support Workbench](#)
- [Viewing Incidents Using the Automatic Diagnostic Repository Command Interpreter \(`adrci`\)](#)
- [Recovering from Incidents Generated During SOA Operations](#)
- [Creating Incidents Manually Using Support Workbench](#)
- [Packaging Incidents Using Support Workbench](#)
- [Creating and Packaging Incidents Using the Automatic Diagnostic Repository Command Interpreter \(`adrci`\)](#)

Regardless of which tool you use to work with incidents, for Oracle Fusion Applications, Oracle Fusion Middleware, and Oracle Database, incidents are stored in the `incident` subdirectories of Automatic Diagnostic Repository (ADR) home directories. However, the location of the appropriate ADR home directory for a given incident depends on the environment from which the incident was created. In all cases, incidents are stored separately from log files.

Paths to directories that contain incident information begin with the path to the ADR base directory, where *domain* is the full path to an Oracle Fusion Applications domain, including the name of the physical Oracle WebLogic Server domain:

```
(UNIX) domain/servers/managed_server_name/adr
(Windows) domain\servers\managed_server_name\adr
```

The ADR base directory path is followed by the path to one of several possible `incident` directories:

For an incident that involves Oracle Fusion Middleware code, the incident information is stored under the following location, where, *domain_name* is the name of the Oracle Fusion Applications domain and *server_name* is the name of the server where the incident occurred:

```
(UNIX) ADR_base_directory/diag/ofm/domain_name/server_name/incident
```

(Windows) `ADR_base_directory\diag\ofm\domain_name\server_name\incident`

For an incident that involves Java code or a SOA composite in an Oracle Fusion application, the incident information is stored under the following location, where `application_name` is the name of the Oracle Fusion application that was executing when the incident occurred:

(UNIX) `ADR_base_directory/diag/ofm/fusionapps/application_name/incident`
 (Windows) `ADR_base_directory\diag\ofm\fusionapps\application_name\incident`

For an incident that involves PL/SQL code in an Oracle Fusion application, the incident information is stored under the following location:

(UNIX) `ADR_base_directory/diag/ofm/db_server_name/database_instance/incident`
 (Windows) `ADR_base_directory\diag\ofm\database_server\database_instance\incident`

For an incident that involves Oracle Database code, the incident information is stored under the following location:

(UNIX) `ADR_base_directory/diag/rdbms/db_server_name/database_instance/incident`
 (Windows) `ADR_base_directory\diag\rdbms\database_server\database_instance\incident`

For an incident that involves C code in an Oracle Fusion application, the incident information is stored on the Java application server that runs C batch jobs using Oracle Enterprise Scheduler. On that server, the incidents are stored under the following location, where `ess_hosted_app_name` is the name of the Oracle Fusion application that launches batch jobs:

(UNIX) `ADR_base_directory/diag/ofm/fusionapps/ess_hosted_app_name/incident`
 (Windows) `ADR_base_directory\diag\ofm\fusionapps\ess_hosted_app_name\incident`

The exact name of the application that launches batch jobs depends upon which product family is in use. Typical values include **FinancialsEssApp**, **HcmEssApp**, **CrmEssApp**, **ProjectFinancialsEssApp**, **ProcurementEssApp**, and **ScmEssApp**. You can determine the correct `ess_hosted_app_name` name for your Oracle Fusion application by using Fusion Applications Control for the application. In the navigation pane, if you expand the folder that is labeled with the product family name and then expand the **Fusion Applications** folder, the `ess_hosted_app_name` is the name under **Fusion Applications** that contains **EssApp**.

14.4.2.1 Viewing Incidents Using Support Workbench

If your Oracle Fusion Applications environment includes Cloud Control, it is recommended that you use Support Workbench to work with incidents, as it is simpler to use than `adrci`.

Note: To use Support Workbench with Oracle Fusion Applications, you must obtain and install Cloud Control.

As part of viewing problem or incident details using Support Workbench, you must supply credentials both to connect to the Oracle WebLogic Server on which the relevant Oracle Fusion Application is running, and to connect to the host machine on which that Oracle WebLogic Server is running.

To view Oracle Fusion Applications incidents using Support Workbench:

1. In Cloud Control, select one of the following Oracle Fusion Applications target types:

- Oracle Fusion Applications instance
- Product family
- Cluster application
- J2EE application

Note: You can find cluster application and J2EE targets in navigation hierarchical trees under the product family name and the Fusion Applications folder. If the tree you are viewing includes Oracle Fusion Applications instance names, you must click a product family name to navigate to a page that displays cluster application and J2EE targets in the navigation tree.

2. From the dynamic drop-down menu, choose **Diagnostics > Support Workbench**.

If you selected a J2EE application as your target, the **Login to Support Workbench** page appears. Skip to Step 6.

If you selected a Fusion Instance, Fusion Product Family, or Fusion Cluster Application as your target, the rollup page for Support Workbench appears, summarizing the Support Workbench problems that are related to your target.

Note: If you want to return to this Support Workbench rollup page at a later time, from another Support Workbench page (such as **Incident Details** or the page that summarizes problems for a selected Fusion J2EE Application target), you can do so by clicking the **Diagnostics** link in the **Related Links** section of the page that corresponds to your selected target type:

- **Fusion Instance Diagnostics**
 - **Fusion Product Family Diagnostics**
 - **Fusion Cluster Application Diagnostics**
-
-

3. From the **View By** drop-down list, select **Deployed Application**.

Note: In order for this page to display links for the number of problems and incidents for Oracle Fusion Applications targets, you must view the data by deployed application, rather than by product. The links for number of problems and incidents that appear in the **Support Workbench Infrastructure Problems and Incidents** table at the bottom of this page are for Oracle WebLogic Server targets, rather than Oracle Fusion Applications targets.

4. From the **View Data** drop-down list, select the timeframe that contains the problem or incident that interests you.
5. In the **Problems and Incidents** table, click a link for the number of problems or incidents that interests you.
The **Login to Support Workbench** page appears.
6. In the **Host Credentials** area of the page, use one of the following methods to select credentials for connecting to the machine that hosts the Oracle WebLogic

Server on which the Oracle Fusion application was running when the problem or incident occurred.

- If appropriate credentials have previously been saved as preferred credentials, select **Preferred** in the **Credential** field.
- If appropriate credentials have previously been saved as named credentials, select **Named** in the **Credential** field and then select the specific credentials from the **Credential Name** drop-down list.
- If appropriate credentials have not previously been set, select **New** in the **Credential** field and complete the following sub-steps:
 - a. In the **UserName** and **Password** fields, enter credential values for connecting to the host machine.
 - b. Inspect the default credential name supplied by Support Workbench in the **Save As** field, and modify the credential name if you prefer.
 - c. If you want the credentials you are entering to be the preferred credentials, select the **Set as Preferred Credentials** checkbox.

The Preferred Host Credential is checked against the host machine's user-names and passwords, but the preference is only set for the currently selected WebLogic server, even if multiple WebLogic servers run on that host machine.
 - d. If you want to verify that the credentials you entered work correctly, click **Test**.

7. In the **WebLogic Server Credentials** area of the page, use one of the following methods to select credentials for connecting to the Oracle WebLogic Server where the Oracle Fusion application was running when the problem or incident occurred.

Note: If you are not prompted for **WebLogic Server Credentials**, that indicates that the relevant Oracle WebLogic Server is currently unavailable. Some operations cannot be done when the server is unavailable, such as creating user-reported problems and performing additional diagnostic dumps. For information about restarting an Oracle WebLogic Server, see the "Starting and Stopping Servers" chapter in *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server*.

- If appropriate credentials have previously been saved as preferred credentials, select **Preferred** in the **Credential** field.
- If appropriate credentials have previously been saved as named credentials, select **Named** in the **Credential** field and then select the specific credentials from the **Credential Name** drop-down list.
- If appropriate credentials have not previously been set, select **New** in the **Credential** field and complete the following sub-steps:
 - a. In the **UserName** and **Password** fields, enter credential values for connecting to the Oracle WebLogic Server.
 - b. Inspect the default credential name supplied by Support Workbench in the **Save As** field, and modify the credential name if you prefer.

- c. If you want the credentials you are entering to be the preferred credentials for the currently selected Oracle WebLogic server, select the **Set as Preferred Credentials** checkbox.

8. Click **Continue**.

A Support Workbench page that summarizes problems related to your target now appears.

9. Click the **Problems** tab and select the relevant time period from the **View** drop-down list.
10. In the **Details** column of the table, click **Show** for the problem listing that interests you.
11. In the list of incidents for the problem you selected, click the **Incidents** link for the incident that occurred at the time that interests you.

The **Incident Details** page appears.

12. In the **Dump Files** table, inspect the list of dump files that are associated with this incident, click the **View Contents** link for any dump file that interests you, and use standard operating system steps to indicate whether to open the file or download it.
13. If you want to create an additional dump file, click the **Additional Diagnostics** tab, select the kind of dump you want to create, and click **Run**.
14. In the **Run User Action** page, enter any required input parameters and choose when to create the dump, and click **Submit**.
15. Click **OK** to dismiss the confirmation message that appears when the dump operation completes.
16. To view the new dump file, repeat Step 12.
17. To create additional dump files, repeat Steps 13 through 15.

For information about using Support Workbench to package an incident and send it to Oracle Support, see [Section 14.4.2.5](#).

14.4.2.2 Viewing Incidents Using the Automatic Diagnostic Repository Command Interpreter (adrci)

If your Oracle Fusion Applications environment does not include Cloud Control, then you must use the Automatic Diagnostic Repository Command Interpreter (adrci) to work with incidents.

To view incidents using the Automatic Diagnostic Repository Command Interpreter (adrci):

1. Log in to the server that contains the Automatic Diagnostic Repository (ADR) that was mentioned in any incident notifications you received.

This may be a repository for either Oracle Database or Oracle Fusion Middleware. The notification may be either a message received by an Oracle Fusion Applications user or an alert email that Oracle Enterprise Manager sent to designated administrators.

2. At an operating system prompt, navigate to the directory that contains a copy of the Automatic Diagnostic Repository Command Interpreter (adrci) that is associated with the ADR home for the incident.
 - If the incident notification referred to an Automatic Diagnostic Repository for Oracle Database, navigate to `DB_ORACLE_HOME/bin` (for UNIX) or `DB_ORACLE_`

`HOME\bin` (for Windows), where `DB_ORACLE_HOME` is the home directory for Oracle Database. For information about home directories for Oracle Database, see the "Overview of Oracle Database Installation" chapter in the *Oracle Database Installation Guide for Microsoft Windows*.

- If the incident notification referred to an Automatic Diagnostic Repository for Oracle Fusion Middleware, navigate to `WL_HOME/server/adr` (for UNIX) or `WL_HOME\server\adr` (for Windows), where `WL_HOME` is the home directory for the Oracle WebLogic Server. For more information about the structure of home directories for Oracle Fusion Applications, see [Section 1.2.3](#).
3. At the prompt, enter `adrci`.
 4. Determine where the incident information is stored.
 - a. At an `adrci` prompt, enter `show base` to display the path to the Automatic Diagnostic Repository (ADR) base directory.
 - b. At an `adrci` prompt, enter `show homes` to display the path to the ADR home directory, under the ADR base directory.
 - c. Make a note of the location of the incident directory, which is a subdirectory of the ADR home directory.
 5. At an `adrci` prompt, enter `show incident`.

This command lists all the incidents that have been created in the incident subdirectory of the ADR home directory.

6. If you already know the ID number of the incident you are investigating, skip to Step 8.

Otherwise, inspect the appropriate `diagnostic.log` file and make a note of either the `incident_ID` value or the `incidentCustomId` value in the relevant log message.

Note: Due to a technical limitation, incidents that are created from PL/SQL code log `incidentCustomId` values rather than `incident_ID` values.

For more information about viewing `diagnostic.log` files, see [Section 13.6](#).

7. If the `diagnostic.log` file supplies the relevant `incident_ID` value, skip to Step 8.

Otherwise, enter the following command at an `adrci` prompt, substituting the appropriate `incidentCustomId` value, and inspect the output for the `incident_ID` value:

```
show incident -mode detail -p "ERROR_ARG1='incidentCustomID'";
```

8. To view detailed information about the incident for which you now know the incident ID number, enter the following command at an `adrci` prompt, substituting the appropriate incident ID number:

```
show incident -mode DETAIL -p "incident_id=incident_id"
```

9. To display a list of the dump files that are associated with the `incident_ID` incident, enter the following command at an `adrci` prompt:

```
show tracefile -i incident_ID
```

10. To view information in a particular dump file for a particular incident, enter one of the following commands at an `adrci` prompt:

- To view a copy of the information in your default editor, enter:

```
show trace filename
```

- To view the information without using an editor:

```
show trace filename -term
```

14.4.2.3 Recovering from Incidents Generated During SOA Operations

In some cases in which an Oracle Fusion Applications incident is generated during an operation that involves a SOA composite, the application code may let you manually recover the operation and resume processing at the point where the incident occurred.

To determine whether you can recover or resume processing of an interrupted SOA composite operation:

1. Log on to the server where the incident was logged as a user that has the Oracle Fusion Middleware role of Administrator.

2. Navigate to the directory where the information about the incident is stored.

For information about typical locations for incidents, see [Section 14.4.2](#).

3. Inspect the `readme.txt` file in the incident directory and make a note of the `composite_instance_id` and `composite_name` values.

For example, the `readme.txt` file might include the following lines:

```
composite_instance_id: 10009
composite_name: FinApInvTransactionsInvoiceApprovalComposite
```

4. In the navigation pane of Fusion Applications Control, expand the hierarchical tree as needed to locate and select the SOA composite that was mentioned in the `readme.txt` file for the incident.

For example, you might expand the `Farm_soa_domain_new` listing, and then expand `SOA`, and then expand `soa_infra (soa_server1)` and then expand `default` and finally select `FinApInvTransactionsInvoiceApprovalComposite`.

5. If it is not already selected, click the **Dashboard** tab.
6. In the **Recent Instances** table in the content pane, click the **Instance ID** value that matches the `composite_instance_id` value that you saw in the log.

The **Flow Trace** page appears.

7. If you have not already done so, use the information on the **Flow Trace** page to determine how to correct the problem that generated the incident, then repeat steps as needed to return to the **Flow Trace** page.

For more information about troubleshooting SOA operations, see the "Troubleshooting Oracle SOA Suite" chapter in the *Oracle Fusion Applications Administrator's Troubleshooting Guide*.

8. Inspect the **Recovery** column of the **Faults** table for a row that contains a **Recoverable** link.
 - If the **Recovery** column contains a **Recoverable** link, then you may be able to resume processing of the SOA operation at the point where it previously

failed. Click **Recoverable** to display more details about the recoverable instance and continue to Step 9.

- If the **Recovery** column does not contain a **Recoverable** link, then the SOA operation is not recoverable. Do not proceed with these instructions. Instead, have an appropriate Oracle Fusion Applications user restart the sequence of operations that led to the incident, beginning as close to the point of failure as is practical.
9. In the list of faults for the instance in question, select the row that contains the word **Recoverable** in the **Recovery** column, and then click **Recover** in the page area below the list.
 10. In the Confirmation dialog box, click **Yes**.

If the recovery process succeeds, the following message is displayed:

Action completed successfully.

Click **OK** to dismiss this message, which marks the completion of the recovery task.

If the recovery process does not succeed, the list of faults with the word **Recoverable** in the **Recovery** column remains visible. This indicates that you have not succeeded in addressing the underlying cause of the original problem. Double-check your fix for the problem or try another approach. When you believe that you have addressed the problem successfully, return to the **Flow Trace** page and repeat Step 8 through Step 10.

14.4.2.4 Creating Incidents Manually Using Support Workbench

If you discover an issue with your Oracle Fusion applications that you want to gather data about, and if the system has not created an incident for it automatically, then it is recommended that you create an incident for it, yourself. Using incidents for all problems helps you organize the information that you gather during the troubleshooting process.

If your Oracle Fusion Applications environment includes Cloud Control, it is recommended that you use Support Workbench to work with incidents.

Note: To use Support Workbench with Oracle Fusion Applications, you must obtain and install Cloud Control.

If your Oracle Fusion Applications environment does not include Cloud Control, then you must use the Automatic Diagnostic Repository Command Interpreter (`adrci`) to work with incidents, instead of Support Workbench. For more information, see [Section 14.4.2.6](#).

When you use Support Workbench to create an incident manually, the process of creating the incident automatically creates a new problem listing that is associated with that incident.

To create a new incident and a new problem listing, manually, using Support Workbench:

1. If you have not already done so, in Cloud Control, select the Fusion J2EE Application target that is associated with the issue for which you are creating an incident.

2. From the dynamic drop-down menu, choose **Diagnostics > Support Workbench**.
3. In the **Host Credentials** area of the page, use one of the following methods to supply credentials for connecting to the machine that hosts the Oracle WebLogic Server that is associated with the incident you are creating.
 - If appropriate credentials have previously been saved as preferred credentials, select **Preferred** in the **Credential** field.
 - If appropriate credentials have previously been saved as named credentials, select **Named** in the **Credential** field and then select the specific credentials from the **Credential Name** drop-down list.
 - If appropriate credentials have not previously been set, select **New** in the **Credential** field and complete the following sub-steps:
 - a. In the **UserName** and **Password** fields, enter credential values for connecting to the host machine.
 - b. Inspect the default credential name supplied by Support Workbench in the **Save As** field, and modify the credential name if you prefer.
 - c. If you want the credentials you are entering to be the preferred credentials for the relevant host, select the **Set as Preferred Credentials** checkbox.
 - d. If you want to verify that the credentials you entered work correctly, click **Test**.
4. In the **WebLogic Server Credentials** area of the page, use one of the following methods to supply credentials for connecting to the Oracle WebLogic Server that is associated with the incident you are creating.

Note: If a **WebLogic Server Credentials** area is not visible on the page, that indicates that the relevant Oracle WebLogic Server is currently unavailable. Some operations cannot be done when the server is unavailable, such as creating user-reported problems and performing additional diagnostic dumps. For information about restarting an Oracle WebLogic Server, see the "Starting and Stopping Servers" chapter in *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server*.

- If appropriate credentials have previously been saved as preferred credentials, select **Preferred** in the **Credential** field.
 - If appropriate credentials have previously been saved as named credentials, select **Named** in the **Credential** field and then select the specific credentials from the **Credential Name** drop-down list.
 - If appropriate credentials have not previously been set, select **New** in the **Credential** field and complete the following sub-steps:
 - a. In the **UserName** and **Password** fields, enter credential values for connecting to the Oracle WebLogic Server.
 - b. Inspect the default credential name supplied by Support Workbench in the **Save As** field, and modify the credential name if you prefer.
 - c. If you want the credentials you are entering to be the preferred credentials, select the **Set as Preferred Credentials** checkbox.
5. Click **Continue**.

A Support Workbench page that summarizes problems related to your target now appears.

6. On the page that summarizes problems for the selected target, scroll down to the **Related Links** section and click **Create User-Reported Problem**.
7. On the **Create User-Reported Problem** page, select the **Other** issue type and click **Continue with Creation of Problem**.
8. On the **Problem Details** page, perform the following sub-steps, as needed:
 - a. If you want to associate the new incident with an existing Oracle Service Request, click **SR# Edit**, enter the Service Request number, and click **OK**.
 - b. If you want to associate the new incident with an existing bug, click **Bug#**, enter the bug number, and click **OK**.
 - c. On the **Incidents** tab, if you want to see more information about the incident you are creating, such as user impact, checker findings, and purge date, click **Show** in the appropriate row of the table, or click **Show All Details** in the table header.
 - d. If you want to change which incidents are displayed in the table, select a new value from the **Data Dumped** drop-down list and click **Go**.
 - e. If you want to see the list of Oracle Fusion Middleware diagnostic dumps that are created automatically when you manually create an incident, or if you want to run any additional Oracle Fusion Middleware diagnostic dumps, select the listing for your incident in the table and click **View** to display the **Incident Details** page.

For more information about using this page, see the topic "Incident Details Page" in the Cloud Control online help. For more information about Oracle Fusion Middleware dumps, see the "Diagnosing Problems" chapter in the *Oracle Fusion Middleware Administrator's Guide*.

14.4.2.5 Packaging Incidents Using Support Workbench

If you discover an issue with your Oracle Fusion applications that you want to gather data about, and if the system has not created an incident and problem for it automatically, then it is recommended that you create an incident for it, yourself.

If your Oracle Fusion Applications environment includes Cloud Control, it is recommended that you use Support Workbench to work with incidents. For more information about using Support Workbench to create incidents manually, see [Section 14.4.2.4](#). The process of manually creating an incident using Support Workbench automatically creates an associated problem listing.

Note: To use Support Workbench with Oracle Fusion Applications, you must obtain and install Cloud Control.

If your Oracle Fusion Applications environment does not include Cloud Control, then you must use the Automatic Diagnostic Repository Command Interpreter (`adrci`) to work with incidents, instead of Support Workbench. For more information, see [Section 14.4.2.6](#).

As part of creating and packaging incidents using Support Workbench, you must supply credentials both to connect to the Oracle WebLogic Server on which the

relevant Oracle Fusion Application is running, and to connect to the host machine on which that Oracle WebLogic Server is running.

Various kinds of incidents can be packaged in Support Workbench. For Oracle Fusion Applications performance incidents, in particular, it is recommended that you use the wizard that is specifically designed for packaging Oracle Fusion Applications incidents, because that wizard greatly simplifies the process of gathering the information that Oracle Support is likely to need when working on your problem. However, you can also use general Support Workbench functionality to package Oracle Fusion Applications incidents.

This section contains the following topics:

- [Packaging Oracle Fusion Applications Incidents Using a Wizard](#)
- [Packaging Oracle Fusion Applications Incidents Using General Support Workbench Functionality](#)

14.4.2.5.1 Packaging Oracle Fusion Applications Incidents Using a Wizard If your Oracle Fusion Applications environment includes Cloud Control, it is recommended that you package Oracle Fusion Applications performance incidents using the special Support Workbench wizard for this incident type. The wizard provides a convenient way to include various file types that Oracle Support may need to resolve your problem.

Note: To use the wizard for packaging Oracle Fusion Applications performance incidents, the relevant incident must already exist. If you need to create an incident manually, see [Section 14.4.2.4](#).

Note: When you plan to package an Oracle Fusion Applications incident, it is recommended that you perform the following actions before you begin the process of packaging the incident.

- Use the Diagnostic Dashboard to perform any diagnostic tests that are specific to that Oracle Fusion application and that are relevant for the selected incident. For more information, see [Section 13.10](#).
- Associate the diagnostic test results with the selected incident. For more information, see [Section 14.6.3](#).

The packaging procedure does not automatically include the results of manually run Oracle Fusion Applications diagnostic tests unless those results are already associated with the incident being packaged.

To package incidents using the Support Workbench wizard for Oracle Fusion Applications incidents:

1. If you have not already done so, in Cloud Control, select a target that is associated with the problem or incident that you want to package.

Make sure that you choose one of the following target types:

- An Oracle Fusion Applications instance target
- An Oracle Fusion Applications product family target
- An Oracle Fusion Applications cluster application target

Note: The Support Workbench wizard for Oracle Fusion Applications incidents is unavailable when an Oracle Fusion application instance (**J2EE Fusion Application**) target is selected.

2. From the dynamic drop-down menu, choose **Diagnostics > Support Workbench**.
3. In the Support Workbench rollup page for your target, scroll down until the **Problem Summary** table is visible and select the problem that is associated with the incident that you want to package.
4. Click **Create Package**.
5. In the **Host Credentials** area of the **Create Package: Credentials** page, use one of the following methods to select credentials for connecting to the machine that hosts the Oracle WebLogic Server on which the Oracle Fusion application was running when the problem or incident occurred.
 - If appropriate credentials have previously been saved as preferred credentials, select **Preferred** in the **Credential** field.
 - If appropriate credentials have previously been saved as named credentials, select **Named** in the **Credential** field and then select the specific credentials from the **Credential Name** drop-down list.
 - If appropriate credentials have not previously been set, select **New** in the **Credential** field and complete the following sub-steps:
 - a. In the **UserName** and **Password** fields, enter credential values for connecting to the host machine.
 - b. Inspect the default credential name supplied by Support Workbench in the **Save As** field, and modify the credential name if you prefer.
 - c. If you want the credentials you are entering to be the preferred credentials for the Oracle WebLogic Server, select the **Set as Preferred Credentials** checkbox.
 - d. If you want to verify that the credentials you entered work correctly, click **Test**.
6. In the **WebLogic Server Credentials** area of the **Create Package: Credentials** page, use one of the following methods to select credentials for connecting to the Oracle WebLogic Server where the Oracle Fusion application was running when the problem or incident occurred.

Note: If a **WebLogic Server Credentials** area is not visible on the page, that indicates that the relevant Oracle WebLogic Server is currently unavailable. Some operations cannot be done when the server is unavailable, such as creating user-reported problems and performing additional diagnostic dumps. For information about restarting an Oracle WebLogic Server, see the "Starting and Stopping Servers" chapter in *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server*.

- If appropriate credentials have previously been saved as preferred credentials, select **Preferred** in the **Credential** field.

- If appropriate credentials have previously been saved as named credentials, select **Named** in the **Credential** field and then select the specific credentials from the **Credential Name** drop-down list.
 - If appropriate credentials have not previously been set, select **New** in the **Credential** field and complete the following sub-steps:
 - a. In the **UserName** and **Password** fields, enter credential values for connecting to the Oracle WebLogic Server.
 - b. Inspect the default credential name supplied by Support Workbench in the **Save As** field, and modify the credential name if you prefer.
 - c. If you want the credentials you are entering to be the preferred credentials for the Oracle WebLogic Server, select the **Set as Preferred Credentials** checkbox.
7. Click **Next**.
8. In the **Create Package: Package Details** page, specify the a name and description for the package you are creating, select the incidents to include in the package, and then click **Next**.
9. In the **Create Package: Dumps** page, select various kinds of dump files to add to the package and then click **Next**.

Note: Unless Oracle Support has given you other instructions, it is recommended that you select all of the types of dump fields that are listed on the page. If you select **Heap Dumps**, you must supply relevant file names and locations. If you select **AWR Reports**, you must specify the name of the relevant database instance and credentials for working with that database.

10. In the **Create Package: Review** page, decide whether or not to customize the package:
- If you need to customize the package, select the **Customize Package** checkbox to display the Customize Package page, and proceed to Step 11.
 - If you do not need to customize the package, click **Proceed to Send** and skip to Step 17.
11. On the **Customize Package: PackageName** page, review the current contents of the package, decide how you want to customize the package, and use the tabs, buttons, and links on the page to make the desired changes.

For more information about the settings and controls on this page, see the "Customize Package Page" topic in the Cloud Control online help.

For more information about customizing actions that you start on this page but complete on other pages, see any of the following Cloud Control online help topics that interest you:

- Add Incidents Page
- Add Problems to Package Page
- Exclude Problems Page
- Package Manifest Page
- Additional Dumps and Test Cases Page

- Add External Files Page
 - Copy Out Files Page
 - Copy In Files Page
12. When you have made the desired changes, click the **Finish Contents Preparation** link in the **Packaging Tasks** area of the page.
After a short wait, the page displays information about any additional relevant files that were added automatically.
 13. In the **Packaging Tasks** area of the page, click the **Generate Upload File** link.
 14. On the **Generate Upload File: *PackageName*** page, select whether the upload file should contain the entire contents of the package, or only an incremental addition made since the last upload file was generated.
For more information about the settings on this page, see the topic "Generate Upload File Page" in the Cloud Control online help.
 15. On the same page, specify whether to generate the package immediately or at a specified later time, and then click **Submit**.
It may take several minutes to generate the upload file. When the process is complete, a confirmation message is displayed. Click **OK** to dismiss it.
 16. On the **Customize Package: *PackageName*** page, click either the **Send Upload Files** link or the **Send to Oracle** button in the **Packaging Tasks** area of the page.
 17. On the **Send to Oracle** page, supply the necessary credentials for logging in to My Oracle Support, make sure that the package you want to upload is selected, and specify whether to send the package immediately or at a specified later time.
For more information about the settings on this page, see the topic "Send to Oracle Page" in the Cloud Control online help.
 18. When you are satisfied with the composition of the package and the details for sending it to Oracle Support, click **Submit**.
After a brief wait, Support Workbench displays a message about whether or not the selected package was sent to Oracle successfully.
If the send attempt was not successful, address any problem that the message describes, then repeat the send attempt.

14.4.2.5.2 Packaging Oracle Fusion Applications Incidents Using General Support Workbench Functionality If your Oracle Fusion Applications environment includes Cloud Control, it is recommended that you package Oracle Fusion Applications performance incidents using the special Support Workbench wizard for that incident type, as described in [Section 14.4.2.5.1](#). For Oracle Fusion Applications incidents that do not concern performance, you can use general Support Workbench functionality to package Oracle Fusion Applications incidents.

Note: When you plan to package an Oracle Fusion Applications incident, it is recommended that you perform the following actions before you begin the process of packaging the incident.

- If no incident exists for the issue that you are investigating, create one manually. For more information, see [Section 14.4.2.4](#).
- Use the Diagnostic Dashboard to perform any diagnostic tests that are specific to that Oracle Fusion application and that are relevant for the selected incident. For more information, see [Section 13.10](#).
- Associate the diagnostic test results with the selected incident. For more information, see [Section 14.6.3](#).

The packaging procedure does not automatically include the results of manually run Oracle Fusion Applications diagnostic tests unless those results are already associated with the incident being packaged.

To package incidents using general Support Workbench functionality:

1. In Cloud Control, if you have not yet selected the J2EE Fusion Application target that is associated with the problem or incident that you want to package, select that target and skip to Step 2.

If you have already selected the J2EE Fusion Application target that is associated with the problem and have already started working with Support Workbench, then complete the following substeps:

- a. Navigate to either the applicable **Problem Details** page or the Support Workbench rollup page that summarizes the problems that are related to your target.
- b. If you navigated to the Support Workbench rollup page, skip to Step 6; otherwise, continue to the following sub-step.
- c. If you navigated to the applicable **Problem Details** page, click either **Quick Package** or the **Package the Problem** link in the **Investigate and Resolve** area of the page.

For more information about the Quick Package button and the Package the Problem link, see the topic "Problem Details Page" in the Cloud Control online help.

- d. If you clicked the **Package the Problem** link, then skip to Step 9; otherwise, continue to the following sub-step.
 - e. If you clicked **Quick Package**, then skip to Step 10.
2. From the dynamic drop-down menu, choose **Diagnostics > Support Workbench**.
 3. In the **Host Credentials** area of the page, use one of the following methods to select credentials for connecting to the machine that hosts the Oracle WebLogic Server on which the Oracle Fusion application was running when the incident or issue occurred.
 - If appropriate credentials have previously been saved as preferred credentials, select **Preferred** in the **Credential** field.
 - If appropriate credentials have previously been saved as named credentials, select **Named** in the **Credential** field and then select the specific credentials from the **Credential Name** drop-down list.

- If appropriate credentials have not previously been set, select **New** in the **Credential** field and complete the following sub-steps:
 - a. In the **UserName** and **Password** fields, enter credential values for connecting to the host machine.
 - b. Inspect the default credential name supplied by Support Workbench in the **Save As** field, and modify the credential name if you prefer.
 - c. If you want the credentials you are entering to be the preferred credentials for the relevant host, select the **Set as Preferred Credentials** checkbox.
 - d. If you want to verify that the credentials you entered work correctly, click **Test**.
- 4. In the **WebLogic Server Credentials** area of the page, use one of the following methods to select credentials for connecting to the Oracle WebLogic Server where the Oracle Fusion application was running when the problem or incident occurred.

Note: If a **WebLogic Server Credentials** area is not visible on the page, that indicates that the relevant Oracle WebLogic Server is currently unavailable. Some operations cannot be done when the server is unavailable, such as creating user-reported problems and performing additional diagnostic dumps. For information about restarting an Oracle WebLogic Server, see the "Starting and Stopping Servers" chapter in *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server*.

- If appropriate credentials have previously been saved as preferred credentials, select **Preferred** in the **Credential** field.
 - If appropriate credentials have previously been saved as named credentials, select **Named** in the **Credential** field and then select the specific credentials from the **Credential Name** drop-down list.
 - If appropriate credentials have not previously been set, select **New** in the **Credential** field and complete the following sub-steps:
 - a. In the **UserName** and **Password** fields, enter credential values for connecting to the Oracle WebLogic Server.
 - b. Inspect the default credential name supplied by Support Workbench in the **Save As** field, and modify the credential name if you prefer.
 - c. If you want the credentials you are entering to be the preferred credentials, select the **Set as Preferred Credentials** checkbox.
- 5. Click **Continue**.
A Support Workbench page that summarizes problems related to your target now appears.
- 6. Click the **Problems** tab and select the relevant time period from the **View** drop-down list.
- 7. In the table that lists problems, select the checkbox for one or more relevant problems.
- 8. Click **Package**.

9. On the **Package: Select Packaging Mode** page, select the mode of packaging that you want to use, and then click **Continue**.

Note: If you selected multiple problems in Step 7, you must select the **Custom Packaging** mode at this point.

For more information about packaging modes, see the topic "Package Page" in the Cloud Control online help.

If you select **Quick Packaging**, continue to Step 10.

If you select **Custom Packaging**, skip to Step 15.

10. On the **Quick Packaging: Create New Package** page, supply information about the name and description for the package, indicate whether you want the package to be sent to Oracle Support automatically, and then click **Next**.

For more information about the fields on this page, see the topic "Quick Packaging Wizard" in the Cloud Control online help.

11. On the **Quick Packaging: View Contents** page, review the information about the incidents that will be included in the package, and click **Next** when you are satisfied.

12. On the **Quick Packaging: View Manifest** page, you can review information about the specific files to be included in the package.

Click **Next** when you are ready to proceed to the next step in the packaging process.

13. On the **Quick Packaging: Schedule** page, specify whether to generate the package immediately or at a specified later time.

If, earlier, you elected to send the package to Oracle Support automatically, the timing you specify for generating the package also applies to sending the package.

For more information about the fields on this page, see the topic "Quick Packaging Wizard" in the Cloud Control online help.

14. When you are satisfied with the composition of the package, click **Submit**.

If you selected **Quick Packaging** in Step 9, skip the remaining steps of this procedure.

15. If you selected **Custom Packaging** in Step 9, on the **Custom Packaging: Select Package** page, either select an existing package to modify, or supply information about the name and description for a new package, and then click **OK**.

For more information about the settings on this page, see the topic "Custom Packaging Page" in the Cloud Control online help.

After a short wait, the page displays a message confirming either that the selected problem was added to an existing package, or that a new package was created.

16. On the **Customize Package: PackageName** page, review the current contents of the package, decide how you want to customize the package, and use the tabs, buttons, and links on the page to make the desired changes.

For more information about the settings and controls on this page, see the "Customize Package Page" topic in the Cloud Control online help.

For more information about customizing actions that you start on this page but complete on other pages, see any of the following Cloud Control online help topics that interest you:

- Add Incidents Page
 - Add Problems to Package Page
 - Exclude Problems Page
 - Package Manifest Page
 - Additional Dumps and Test Cases Page
 - Add External Files Page
 - Copy Out Files Page
 - Copy In Files Page
17. When you have made the desired changes, click the **Finish Contents Preparation** link in the **Packaging Tasks** area of the page.
- After a short wait, the page displays information about any additional relevant files that were added automatically.
18. In the **Packaging Tasks** area of the page, click the **Generate Upload File** link.
19. On the **Generate Upload File: *PackageName*** page, select whether the upload file should contain the entire contents of the package, or only an incremental addition made since the last upload file was generated.
- For more information about the settings on this page, see the topic "Generate Upload File Page" in the Cloud Control online help.
20. On the same page, specify whether to generate the package immediately or at a specified later time, and then click **Submit**.
- It may take several minutes to generate the upload file. When the process is complete, a confirmation message is displayed. Click **OK** to dismiss it.
21. On the **Customize Package: *PackageName*** page, click the **Send Upload Files** link in the **Packaging Tasks** area of the page.
22. On the **Send to Oracle** page, supply the necessary credentials for logging in to My Oracle Support, make sure that the package you want to upload is selected, and specify whether to send the package immediately or at a specified later time.
- For more information about the settings on this page, see the topic "Send to Oracle Page" in the Cloud Control online help.
23. When you are satisfied with the composition of the package and the details for sending it to Oracle Support, click **Submit**.
- After a brief wait, Support Workbench displays a message about whether or not the selected package was sent to Oracle successfully.
- If the send attempt was not successful, address any problem that the message describes, then repeat the send attempt.

14.4.2.6 Creating and Packaging Incidents Using the Automatic Diagnostic Repository Command Interpreter (adrci)

If you discover an issue with your Oracle Fusion applications that you want to gather data about, and if the system has not created an incident for it automatically, then it is recommended that you create an incident for it, yourself.

If your Oracle Fusion Applications environment includes Cloud Control, it is recommended that you use Support Workbench to create and package incidents. For more information, see [Section 14.4.2.4](#) and [Section 14.4.2.5](#).

If your Oracle Fusion Applications environment does not include Cloud Control, you must use WLST to create incidents and `adrci` to package them. For more information, see the "Creating an Incident Manually" section in the *Oracle Fusion Middleware Administrator's Guide*.

After creating the incident manually, you can add files to it. For example, you might want to add Oracle Fusion Applications diagnostic test results, Oracle Fusion Middleware diagnostic dumps, excerpts from downloaded log files, and Remote Diagnostic Agent (RDA) reports to the incident. You might also want to add a `Readme.txt` file containing information for Oracle Support. After you add files to an incident, you can package the incident into a zip file for transmission to Oracle Support. For more information about adding files to an incident and packaging the incident into a zip file, see the "Packaging an Incident" section in the *Oracle Fusion Middleware Administrator's Guide*.

14.4.3 Working with Automatically Created QuickTrace Dumps

As part of assisting you with resolving incidents, Oracle representatives often examine any dump files that were automatically created when the incident occurred, including QuickTrace dump files.

When a QuickTrace dump file is automatically created as part of an incident, the file is stored in the incident directory along with any other dump files that are related to that incident. QuickTrace dump file names begin with `odl.quicktrace`.

This section contains the following topics:

- [Viewing Automatically Created QuickTrace Dump Files Using Oracle WebLogic Scripting Tool](#)
- [Viewing Automatically Created QuickTrace Dump Files Using the Automatic Diagnostic Repository Command Interpreter \(adrci\)](#)
- [Viewing Automatically Created QuickTrace Dump Files Using Support Workbench](#)

14.4.3.1 Viewing Automatically Created QuickTrace Dump Files Using Oracle WebLogic Scripting Tool

You can use the following standard Oracle WebLogic Scripting Tool commands to view information about automatically created QuickTrace dump files:

- Use the `listIncidents` command to view a list of incidents.
- Use the `showIncident` command to view details about a particular incident, including a list of dump files associated with that incident.
- Use the `getIncidentFile` command to view a specified dump file such as a QuickTrace dump file.

For more information about using these commands, see the "Viewing Problems and Incidents" section and the "Working with Diagnostic Dumps" section in the *Oracle Fusion Middleware Administrator's Guide*.

14.4.3.2 Viewing Automatically Created QuickTrace Dump Files Using the Automatic Diagnostic Repository Command Interpreter (adrci)

You can use the following standard Automatic Diagnostic Repository Command Interpreter (adrci) commands to view information about automatically created QuickTrace dump files:

- Use the `show homes` command to view a list of Automatic Diagnostic Repository (ADR) home directories. Depending on the configuration of your system, QuickTrace dumps may be located under one or more of these ADR home directories.
- Use the `set home directory_path` command to indicate which repository you want to work with, replacing `directory_path` with one of the locations from the output of the `show home` command.
- Use the `show incdir -i incdir_incident_number` command to view a list of files associated with the incident that `incident_number` identifies, including any QuickTrace dump files for that incident. QuickTrace dump file names start with `odl_quicktrace` and end with `.dmp`.
- Use the `show tracefile -i incident_number` command to view a list of just the dump files associated with the incident that `incident_number` identifies.
- Use the `show trace filename` command to view the contents of the QuickTrace dump file that you specify.

For more information about using `adrci`, see [Section 14.4.2](#) and the chapter about the ADR Command Interpreter in *Oracle Database Utilities*.

14.4.3.3 Viewing Automatically Created QuickTrace Dump Files Using Support Workbench

If your Oracle Fusion Applications environment includes Cloud Control, then you can use Support Workbench to view Oracle Fusion Applications incident information, including automatically created QuickTrace dumps.

1. In Oracle Enterprise Manager, select an Oracle Fusion Applications instance, product family, or cluster application.
2. From the dynamic drop-down menu, choose **Diagnostics > Support Workbench**.
3. From the **View Data** drop-down list, select the timeframe that contains the problem for which you want to inspect a QuickTrace dump.
4. Scroll down to the **Support Workbench Problems Summary** table and click the **Incident** count for the problem that interests you.
5. Click the **Incidents** tab in the **Problem Details** page.
6. Click the **Summary** of the incident that interests you.
7. In the **Guided Resolution** area of the **Incident Details** page, click **View Diagnostic Data**.
8. In the **Dump Files** table, locate a row that contains a **File Name** that starts with `odl_quicktrace`.
9. Scroll to the right and click the **View Contents** icon in that row.
10. Inspect the information in the **View Log File** page for the dump file, and click **Download** if you want to download the file.

14.4.4 Working with Manually Created QuickTrace Dumps

Oracle Support may occasionally ask you to perform a manual dump of QuickTrace buffers as part of work to resolve a problem that did not generate an incident automatically.

To create a QuickTrace dump manually:

1. Decide whether you want to associate the QuickTrace dump file with an existing incident at the time when you perform the dump:
 - If you want to associate a new QuickTrace dump file with an existing incident at the time when you perform the dump, make a note of the relevant incident ID. This can either be an automatically created incident or a manually created incident. For more information about creating incidents manually, see [Section 14.4.2.4](#) and [Section 14.4.2.6](#).
 - If you do not want to associate a new QuickTrace dump with an incident at the present time, decide where you will store the dump file.
2. Use the Oracle WebLogic Scripting Tool `executeDump` command to perform an `odl.quicktrace dump`, choosing appropriate syntax for your needs:
 - If you want to associate the dump file with an existing incident, use the `id` argument. This allows you to retrieve the dump file, later, using the `getIncidentFile` argument
 - If you would rather associate the dump file with an incident at a later time, use the `outputFile` argument to specify an output file location at the time when you execute the dump.

For more information about using the `executeDump` command and its arguments, see the "Executing Dumps" section in the *Oracle Fusion Middleware Administrator's Guide*.

14.4.5 Working with Other Diagnostic Dumps

In some troubleshooting situations, you may find it useful to locate or obtain Oracle Fusion Middleware diagnostic dumps such as Oracle WebLogic Server diagnostic images, Java Virtual Machine thread dumps, and Dynamic Monitoring Service metric dumps.

Oracle Fusion Applications incidents that are generated from Java code (including SOA) automatically execute these kinds of dumps and include the results in the incident data. Oracle Fusion Middleware diagnostic dumps are not available for Oracle Fusion Applications incidents that are generated from PL/SQL or C code.

If you troubleshoot an Oracle Fusion Applications problem that did not generate an incident automatically, it is recommended that you create an incident manually, then execute diagnostic dumps manually, and examine the dump files for information that may help you to resolve the problem. If you need to work with Oracle Support to resolve the problem, it is recommended that you include manually executed dump results in the incident data package that you send to Oracle for review. For information about how to list available diagnostic dumps, view dump descriptions, and manually execute selected dumps for Oracle Fusion Middleware, see the "Working with Diagnostic Dumps" section in the *Oracle Fusion Middleware Administrator's Guide*.

14.4.6 Working with Cloud Control Problem Analysis and the Analyze Logs Options

If your Oracle Fusion Applications environment includes Cloud Control, you can use the **Problem Analysis** and **Analyze Log** options in Cloud Control to help you inspect metrics, target status information, and logs during troubleshooting.

To inspect metrics, status information, and logs using Cloud Control:

1. Use one of the following methods to select a target for which you want to use problem analysis or analyze logs:
 - Navigate to **Incident Manager** and select the incident that interests you.
 - Navigate to the home page of an Oracle Fusion Applications product family and locate a metric chart in the **Performance Indicators** section.
 - From the home page of any Fusion Instance, Product Family, Product, Fusion Cluster Application, or Fusion J2EE Application, choose **Monitoring > Key Performance Indicators** or **Monitoring > System Performance** and then locate a metric chart.
2. Use one of the following methods to display **Problem Analysis** pages for Oracle Fusion Applications targets:
 - In Incident Manager, click **Problem Analysis** in the **Guided Resolution** pane.
 - On a home page or a metrics page such as a **Key Performance Indicators** or **System Performance** page, under or adjacent to any metric chart, click a link to display the **Additional Information** pop-up and then click **Problem Analysis**.

The name and location of the link that displays the **Additional Information** popup varies depending on the page and metric chart you are viewing. For example, some typical link names are:

- **Request Processing Time** (for **Response and Load** metric charts)
- **Payables Invoice Validation Run Time** (for an Oracle Fusion Financials **Key Performance Indicator** metric chart)

In the **Problem Analysis** page, the **Source Metric** chart displays the type of metric for which the incident was automatically created or the type of metric for which you clicked **Request Processing Time**.

3. Near the top of the **Metric Dependency Charts** pane, click the time period for which you want the charts to display data, or click and drag within a metric chart to indicate the time period you want to inspect.
4. Inspect the charts for unusual increases in recorded metrics such as request processing time, CPU usage in storage units or percent utilized, number of requests per minute, Java Virtual Machine heap memory usage, or server memory usage.

If you find that request processing time increased due to a high number of requests per minute, you may need to increase the capacity of your system.

5. If the metric charts do not indicate the cause of the problem, scroll down to the **Related Information** pane and inspect the **List of Related Targets To Analyze** table for information about target health (status) and recent configuration changes.
6. If you want to see a reminder of the topology of the components for which data is being displayed, click the **Related Targets Topology** tab.
7. If the **List of Related Targets to Analyze** table does not indicate the cause of the problem, scroll up and click a link adjacent to one of the metrics and then click

Analyze Logs in the **Additional Information** pop-up to display log messages for the selected target and its members during the selected time period.

8. To display log messages in the selected time period for all targets related to the relevant incident or starting metric, scroll up to the **Metric Dependency Charts** pane header and click **View Related Log Messages**.
9. Inspect any log messages that are displayed for possible causes of problems.

14.5 Troubleshooting Using Log Settings

You can use log settings to increase the amount of information that a particular Oracle Fusion application gathers about its own operations. This is particularly useful with problems that are easily reproducible—you can change log settings, reproduce the problem, then return log settings to their normal levels so that you do not use disk space for unnecessarily detailed information that is not relevant to the problem.

This section contains the following topics:

- [Troubleshooting Using Standard Log Files](#)
- [Viewing Special Log Output for Oracle Fusion Incentive Compensation](#)
- [Correlating Log Messages Across Log Files and Components](#)
- [Downloading Log Files](#)

When you are troubleshooting a problem in an Oracle Fusion Applications environment, you may also find it helpful to use the logging capabilities of Oracle Fusion Middleware. At times, Oracle Support may ask you to use certain of these capabilities when attempting to reproduce a problem, particularly if you have a dedicated testing environment available.

14.5.1 Troubleshooting Using Standard Log Files

Most code modules in Oracle Fusion Applications use a standard diagnostic logging framework in which profile options determine how much information is logged for a particular site or a particular user. For information about adjusting site-level profile options or adjusting user-level profile options as an administrator, see [Section 13.7.5](#). For information about allowing selected users to adjust the level of information that is logged for their own sessions, see [Section 14.2.1](#).

14.5.2 Viewing Special Log Output for Oracle Fusion Incentive Compensation

For information about configuring the special logging functionality for certain Oracle Fusion Incentive Compensation batch jobs, see [Section 14.2.9](#). After you have enabled this functionality and attempted to replicate the Oracle Fusion Incentive Compensation problem that you are troubleshooting, you must perform some database queries to view the resulting log entries. For these queries, you must specify the name of the Oracle Fusion Incentive Compensation process that you ran when attempting to replicate the reported problem. The process names that you can use in the query are as follows:

- CALCULATION MAIN PROCESS
- CLASSIFICATION
- COLLECT_STATISTICS
- Crediting

- GEN_RULE_ENGINE
- GET_WINNERS
- PopulateSrpBatches
- PopulateSrpTables
- REVERT TRANSACTIONS
- ROLLUP
- Rollup within Crediting

Note: These process names are case-sensitive. For more information about how Oracle Fusion Incentive Compensation processes operate, see product-specific documentation in Oracle Fusion Applications Help.

To view special log entries for Oracle Fusion Incentive Compensation batch jobs:

1. Use a SQL client to connect to Oracle Database and log in as a user that has permission to view the Fusion database schema.
2. Execute the following SQL command, replacing *process_name* with the name of the Oracle Fusion Incentive Compensation process that you ran when attempting to replicate the reported problem:

```
select process_audit_id from cn_process_audits_all
where process_type = process_name
order by creation_date desc;
```

3. Make a note of the process audit ID that is returned in the first row of the output.
4. Execute the following SQL command, replacing *selected_process_audit_ID* with the value that you noted in Step 3.

```
select message_type, message_text from cn_process_audit_lines_all
where process_audit_id = selected_process_audit_ID
order by process_audit_line_id asc;
```

5. If you have not already done so, place the output of Step 4 in a file that you can inspect with a text editor or send to Oracle Support if requested.
6. When you are finished gathering log entries, reset the value of the CN_DEBUG profile option setting to N.

For information about setting the value of the CN_DEBUG profile option, see [Section 14.2.9](#).

14.5.3 Correlating Log Messages Across Log Files and Components

Sometimes, as part of the process of troubleshooting Oracle Fusion Applications, you may need to find all of the log entries that result from a particular action. If the action involves multiple Oracle Fusion applications or multiple Oracle Fusion Middleware components, then the log entries are placed in multiple log files. You can use certain values related to the particular action to locate the relevant log messages in those log files.

This section contains the following topics:

- [Values for Correlating Log Messages](#)

- [Correlating Log Messages for Oracle Enterprise Scheduler Jobs](#)
- [Correlating Log Messages for SOA Services and BPEL Processes](#)
- [Correlating Log Messages for Other Process Types](#)

14.5.3.1 Values for Correlating Log Messages

You can use a particular action's Execution Context Identifier (ECID) and Relationship Identifier (RID) values to locate the log entries that result from that action.

A unique ECID value is assigned to each user request. However, depending on what is involved in the request, a log entry that is related to that request may have the original ECID of the request, or a different ECID of a child process spawned by the original request. For example, if you submit a job to Oracle Enterprise Scheduler, the request to schedule the job and the job, itself, always have different ECIDs.

For a particular ECID, if the request requires processing by multiple components, by multiple modules within a component, or using multiple threads, then a unique RID value may be assigned when the work passes to a new thread or process.

In log entries, the ECID and RID values are typically separated by a comma. ECID values can include alphanumeric characters, colons, periods, and hyphens. RID values can include numerals and colons. The RID value indicates the relationship of the current thread or process to its creator. The first RID created for a request is RID = 0. Each shift in generation is represented by a colon. For example, the RID value for the seventh child process of the third child process of the root context for a request is represented as 0:3:7.

14.5.3.2 Correlating Log Messages for Oracle Enterprise Scheduler Jobs

When a job is submitted to Oracle Enterprise Scheduler, the request to schedule the job and the job, itself, always have different Execution Context Identifiers (ECIDs).

The ECID for the job-scheduling request is displayed in the requesting user session at the time when the job is submitted. However, the ECID of the Oracle Enterprise Scheduler job, itself, is not assigned until the job starts, and it is not displayed in the requesting user session. To locate log entries related to the Oracle Enterprise Scheduler job, itself, you must use the job Request ID value.

To correlate log messages for Oracle Enterprise Scheduler jobs using Fusion Applications Control:

1. In the navigation pane, expand the **Farm** listing and then the **Scheduling Services** listing and then determine whether **Scheduling Services** contains Oracle Enterprise Scheduler clusters that contain individual Oracle Enterprise Scheduler servers, or individual Oracle Enterprise Scheduler servers without clusters.
2. To open the ESSAPP deployment administration page, click the listing for an Oracle Enterprise Scheduler cluster such as **ESSAPP (ESSCluster)**, if one is available, or an individual Oracle Enterprise Scheduler server such as **ESSAPP (ess_server1)**.

Note: If you can select a cluster, that will allow you to search all servers in the cluster simultaneously. This is an advantage because each job request is logged only on the Oracle Enterprise Scheduler server where that request is actually processed.

3. If the job request for the action that you are troubleshooting appears in the list of **Top 10 Running Job Requests**, click its **Request ID** link to display the request details, then skip to Step 8.
4. In the context pane, from the dynamic **Scheduling Service Group** or **Scheduling Service** target menu, choose **Job Requests > Search Job Request**.
5. Use standard search techniques to locate the job request for the action that you are troubleshooting.

For information about searching Oracle Enterprise Scheduler job requests, see [Section 5.8.2](#).

6. In the search results, make a note of the **Request ID** value for the job that interests you.
7. In the search results, click the **Request ID** link to display the request details.
8. In the **Request Details** pane, select **Request Log** from the Action drop-down list.
The **Log Messages** page displays all the log messages for the ECID associated with the selected job request.
9. Click any **Log File** link to display a list of the related log messages in the **View Log Messages** page.

At this point, you can inspect the log messages on your screen or download the log file and search it for all messages containing the relevant ECID.

Note: If the log level for a job request is set to `FINER` or `FINEST` when the job runs, then the diagnostic log entries for the job request are written to the log file specified by the `AFLOG_FILENAME` profile option, as well as to the request log file that is available to the user who submitted the job request.

14.5.3.3 Correlating Log Messages for SOA Services and BPEL Processes

If you need to correlate log messages that are related to SOA services or BPEL processes, you obtain the Execution Context Identifier (ECID) by first locating the relevant SOA composite instance.

To correlate log messages for SOA Services or BPEL processes using Fusion Applications Control:

1. In the navigation pane, expand the **Farm** listing, and then the **SOA** listing, and then the **soa_infra** listing, and then and select **default**.
2. In either pane, click the SOA composite instance name for which you want to locate log entries.
3. Make sure that your browser is set to allow the display of popup windows.

For example, in Firefox, choose **Options > Content** from the **Tools** menu and either make sure that the **Block pop-up windows** checkbox is cleared or add any relevant hosts to **Exceptions**.

Similarly, in Internet Explorer, choose **Pop-up Blocker > Turn Off Pop-up Blocker** from the **Tools** menu.

4. In the **Recent Instances** table, click the relevant **Instance ID** link to display the Flow Trace popup window.

The **ECID** is displayed in the upper right corner of the window.

5. Make a note of the **ECID** value and close the Flow Trace popup window.
6. To view the relevant log entries using Fusion Applications Control, complete the following substeps. To view the log entries using the Oracle WebLogic Scripting Tool (WLST), skip to Step 7.
 - a. In the navigation pane, select an appropriate target.
 - b. From the dynamic target menu, choose **Logs > View Log Messages**.
 - c. On the Log Messages page, in the **Search** area, click **Add Field**.
 - d. In the popup list, select the **ECID** checkbox and click **Add**.
 - e. Enter the **ECID** value in the **ECID** field.
 - f. Adjust other search field values appropriately and click **Search**.

Note: You can leave any search fields blank; blank fields are not used as search criteria.

Skip the remainder of this procedure.

7. In the Oracle WebLogic Scripting Tool (WLST), enter a command like the following example, but substitute the ECID value that you previously noted.

```
displayLogs (ecid='0000H19TwKUCs1T6uBi8UH181kWX000002')
```

14.5.3.4 Correlating Log Messages for Other Process Types

Many operations within Oracle Fusion applications do not use Oracle Enterprise Scheduler or SOA services or BPEL processes. For these operations, you can correlate log entries using the Execution Context Identifier (ECID).

To locate log entries that result from one particular Oracle Fusion application action that does not use Oracle Enterprise Scheduler or SOA services or BPEL processes:

1. Use Fusion Applications Control to search the log file that contains messages for the Oracle Fusion application in which the action was taken.

Use any search criteria that will allow you to locate a log entry that is associated with the particular action. The associated ECID and RID values are included in the detailed log entry information. For more information about searching for particular information in an Oracle Fusion Applications log file, see [Section 13.6.2](#) or [Section 13.6.3](#).

2. When you have the ECID and RID values for the action, use one of the following methods to view all of the log entries associated with the action.
 - In Fusion Applications Control, on the Log Messages page, in the **Search** area, click **Add Field** to make sure that the **ECID** search field is displayed, then enter the **ECID** value in that field and click **Search**.
 - In Fusion Applications Control, on the page of Log Messages search results, with a relevant log message selected, select **by ECID** from the **View Related Messages** list.
 - In the Oracle WebLogic Scripting Tool (WLST), enter a command like the following example, but substitute the ECID value that you previously noted.

```
displayLogs (ecid='0000H19TwKUCs1T6uBi8UH181kWX000002')
```

For more information about using ECIDs and RIDs to correlate log entries, see the "Correlating Messages Across Log Files and Components" section in the *Oracle Fusion Middleware Administrator's Guide*.

14.5.4 Downloading Log Files

As part of the process of troubleshooting Oracle Fusion Applications problems, an Oracle Support representative may occasionally request a copy of a particular log file. You can add downloaded log files to incidents in order to package those log files with incident information that you send to Oracle Support.

For more information about using Support Workbench to package incidents along with external files such as log files, see [Section 14.4.2.5.2](#). For more information about adding files to an incident and packaging the incident into a zip file using `adrci`, see the "Packaging an Incident" section in the *Oracle Fusion Middleware Administrator's Guide*.

To download an Oracle Fusion Applications log file using Fusion Applications Control:

1. From the navigation pane, select the target for which you want to download a log file.
2. In the context pane, from the dynamic target menu, choose **Logs > View Log Messages** to display the log entries for the target you selected.
3. Use one of the following methods to locate the log file that you want to download:
 - Inspect the **Log File** column in the table of log entries that appears by default. If the column contains a link for the log file that you want to download, click the link and proceed to Step 4.
 - Perform a standard log file search using criteria that you expect to list the desired log file in the search results. When the **Log File** column contains a link for the log file that you want to download, click the link and proceed to Step 4.
 - Expand the **Selected Targets** section of the **Search** area and click the **Target Log Files** icon for the target most likely to be associated with the log file you want to download. In the list of log files that appears, select the file that you want to download and proceed to Step 4.
4. Click **Download** and specify where you want the log file to be saved.

14.6 Troubleshooting Using the Diagnostic Testing Framework

Some Oracle Fusion Applications diagnostic tests are designed to help you to monitor the health of your system. Other diagnostic tests are designed to help you troubleshoot problems or to gather data that Oracle Support may need when helping you to resolve a problem. For basic information about using Diagnostic Testing Framework features that are useful for both troubleshooting and monitoring system health, see [Chapter 13](#). For information about the individual diagnostic tests that are provided with this release, see the *Oracle Fusion Applications Common User Guide* in Oracle Fusion Applications Help.

This section contains the following topics:

- [Working with Automatically Launched Diagnostic Tests](#)
- [Registering Diagnostic Tests](#)
- [Sending Diagnostic Test Results to Oracle Support](#)

- [Purging the Results of Selected Diagnostic Test Runs](#)

Note: The user name that you use to sign in to an Oracle Fusion application affects which diagnostic operations are available to you. Be sure that you sign in using an account that is assigned to the job roles for the diagnostic operations that you need. For more information, see [Section 13.9.1](#).

14.6.1 Working with Automatically Launched Diagnostic Tests

To make sure that certain diagnostic tests run automatically when certain error conditions occur, Oracle developers set the value of each test's `APPS_MSG_ID` tag to match the identifier of any error message that should trigger the automatic execution of that test. There is no configuration setting for disabling this automatic execution of diagnostic tests.

If an error has resulted in the automatic creation of an incident, then the results of any automatically run diagnostic tests are automatically associated with the incident. The identity of the user who received the error message is also recorded.

14.6.2 Registering Diagnostic Tests

Oracle supplies you with a number of diagnostic tests for Oracle Fusion Applications in seed data. Diagnostic tests that are supplied in seed data do not require you to register them.

However, if you report an Oracle Fusion application issue to Oracle Support, you may receive additional diagnostic tests to help resolve the issue. Depending on the exact circumstances, those diagnostic tests may or may not require registration. If registration is required for the additional diagnostic tests, Oracle will provide you with detailed registration instructions to follow.

14.6.3 Sending Diagnostic Test Results to Oracle Support

From time to time, you may run a diagnostic test and then want to send the results of the test to Oracle Support. To accomplish this, you attach the results of the test to the relevant incident.

Note: You can also attach other types of files to incidents, such as log file excerpts or `Readme.txt` files containing information for Oracle Support. For more information about using Support Workbench to package incidents along with files such as log file excerpts or `readme.txt` files, see [Section 14.4.2.5.2](#). For more information about using `adrci` to add such files to incidents, see the "Packaging an Incident" section in the *Oracle Fusion Middleware Administrator's Guide*.

To attach diagnostic test results to an incident:

1. Sign in to the Oracle Fusion application that generated the incident, using an account that has the `Diagnostic Advanced User` duty role or the `Diagnostic Administrator` duty role.
2. In the Oracle Fusion application, choose **Troubleshooting > Run Diagnostic Tests** from the **Help** menu to display the Diagnostic Dashboard instance for that application. For more information, see [Section 13.9.2](#).

3. If you are not already displaying the **Diagnostic Test Run Status** pane, expand the **Tasks** panel in the Regional area of your screen and click **Run Status**.
4. In the **Diagnostic Test Run Status** table, locate and select the test run that has the report that you want to add to an incident.
5. Click **Add to Incident** to display the Select Incident dialog box.
6. Click the Select ADR Home icon to the right of the **ADR Home** field.
7. In the Search and Select: ADR Home dialog box, locate and select the table row for the **ADR_HOME** and server that your Oracle Fusion application is currently using.

If you wish, you can use any of the following techniques to locate the row that describes your **ADR_HOME** and server:

- In the **ADR Home** field, enter any part of your **ADR_HOME** name and then click **Search**.
- In the **Server** field, enter any part of the name of the server that is running your Oracle Fusion application and then click **Search**.
- In both the **ADR Home** field and the **Server** field, enter the values you want to search for, then select either **All** (to display rows that match both search strings) or **Any** (to display rows that match either search string) and click **Search**.

Note: In the Search and Select: ADR Home dialog box, search operations are case-sensitive.

8. With the correct table row highlighted, click **OK** to return to the Select Incident dialog box.
9. To display a list of known problems for your selected ADR home directory, click the Select Problem icon to the right of the **Problem** field.
10. In the Search and Select: Problem dialog box, locate and select the problem that describes the incident to which you want to add test results.
11. With the correct problem highlighted, click **OK** to return to the Select Incident dialog box.
12. From the **Incident** field, select the incident to which you want to add test results, and then click **OK**.

The test report results are added to the incident and a confirmation message is displayed.

14.6.4 Purging the Results of Selected Diagnostic Test Runs

From time to time, you may want to remove diagnostic test run results from your database, to keep the Run Status table from becoming too large. For information about deleting diagnostic test results, see [Section 13.10.11](#).

Part IV

Advanced Administration

This part describes advanced administration tasks that involve reconfiguring the Oracle Fusion Applications environment.

Part IV contains the following chapters:

- [Chapter 15, "Backing Up and Recovering Oracle Fusion Applications"](#)
- [Chapter 16, "Configuring High Availability and Scaling Out Oracle Fusion Applications"](#)
- [Chapter 17, "Moving Components for Oracle Fusion Applications Across Environments"](#)

Backing Up and Recovering Oracle Fusion Applications

This chapter describes recommended back up and recovery strategies and procedures for recovering your Oracle Fusion Applications environment from different types of failures and outages, such as data loss or corruption, host failure, or media failure.

This chapter contains the following topics:

- [Introduction to Backup and Recovery](#)
- [Overview of Backing Up Your Environment](#)
- [Overview of Recovering Your Environment](#)
- [Prerequisites for Using Cloud Control to Back Up or Restore Your Environment](#)
- [Backup and Recovery Recommendations for Oracle Fusion Applications](#)
- [Performing a Backup](#)
- [Creating a Record of Your Oracle Fusion Applications Configuration](#)
- [Recovering Using Cloud Control](#)
- [Recovering After Data Loss, Corruption, or Media Failure](#)
- [Recovering After Loss of Host](#)
- [A Case Study: Recovering Oracle Fusion Customer Relationship Management](#)

15.1 Introduction to Backup and Recovery

An Oracle Fusion Applications environment can consist of different Oracle Fusion Applications product families. It is built on Oracle Fusion Middleware, which contains Oracle WebLogic Server domains with Java components, such as Oracle SOA Suite, and system components, such as Oracle HTTP Server. It also contains a separate Oracle WebLogic Server domain with Identity Management components, such as Oracle Internet Directory and Oracle Virtual Directory.

Oracle Fusion Applications and some Oracle Fusion Middleware components use Oracle Database instances to store data and metadata.

Note that Oracle Fusion Applications uses at least three databases:

- A database that holds the LDAP repository
- A database that holds the Oracle Fusion Applications data
- A database that holds the Oracle Identity Manager repository

The installations of an Oracle Fusion Applications environment are interdependent in that they contain configuration information, applications, and data that are kept in synchronization.

It is, therefore, important to consider your entire Oracle Fusion Applications environment when performing backup and recovery. You should back up your entire Oracle Fusion Applications environment as soon as you have installed and configured it, then periodically. If a loss occurs, you can restore your environment to a consistent state.

15.2 Overview of Backing Up Your Environment

You should back up your environment when you install and configure Oracle Fusion Applications and on a regular basis. You can back up your full environment or you can back up only parts of it. You can perform the backups in online or offline mode.

This section includes the following topics:

- [Tools to Use to Back Up Your Environment](#)
- [Modes of Backup](#)
- [Types of Backups](#)
- [Recommended Backup Strategy](#)

15.2.1 Tools to Use to Back Up Your Environment

To back up your Oracle Fusion Applications environment, you can use:

- Oracle Enterprise Manager Cloud Control
 - With Cloud Control, you can back up the directories and databases that comprise a Oracle Fusion Applications installation. Because Cloud Control uses Oracle Secure Backup to perform Oracle Fusion Applications backups, all backups are done to tape, which is the media supported by Oracle Secure Backup.
- File copy utilities such as `copy`, `xcopy`, `tar`, or `jar`. Make sure that the utilities:
 - Preserve symbolic links
 - Support long file names
 - Preserve the permissions and ownership of the files

For example:

- On Windows, for online backups, you can use `copy`; for offline backups, you can use `copy`, `xcopy`, `tar`, or `jar`. Ensure that the backup utility you use can support long and unicode file names and extensions. Many of the early archiving utilities did not have this support.

Note that for some versions of Windows, any file name with more than 256 characters will fail. You can use the `xcopy` command with the following switches to work around this issue:

```
xcopy /s/e "C:\Temp\*.*)" "C:\copy"
```

See the `xcopy` help for more information about syntax and restrictions.

- On UNIX, for online and offline backups, you can use `tar`.

See [Section 15.2.2](#) for descriptions of online and offline backups.

- Oracle Recovery Manager (RMAN) to back up database-based metadata repositories and any databases used by Oracle Fusion Applications. With RMAN, you can perform full backups or incremental backups. See *Oracle Database Backup and Recovery User's Guide* for information about using RMAN to back up a database.

If you want to retain your backups for a longer duration, consider backing up to tape, for example using Oracle Secure Backup.

For optimized backup time and fast restore times for the file system, you can use a storage snapshot feature provided by a storage vendor. Snapshots are point-in-time, read-only copies of the file system. Snapshots normally use copy-on-write mechanisms and hence do not occupy any space at the beginning. When new data is written, the old copy is written to the snapshot. Most of the storage vendors support an unlimited number of snapshots and allow creating snapshots manually or automatically without any user interventions. Snapshots can be taken quickly, thus reducing the time taken when performing a backup. Taking regular snapshots of the Oracle Fusion Applications configuration helps you restore the Oracle Fusion Applications environment quickly in the event of configuration corruption or data loss. Snapshots provide rollback capabilities. When a rollback occurs, any newer snapshots (and clones of newer snapshots) are destroyed, and the active data reverts to the state when the snapshot was taken.

15.2.2 Modes of Backup

You can back up your Oracle Fusion Applications environment offline or online:

- With an **offline backup**, you must shut down the environment before performing the backup. When you perform an offline backup, the Administration Server, all Managed Servers in the domain, and all system components in the Oracle instances should be shut down.

Back up the environment offline immediately after installation and after applying upgrades.

- With an **online backup**, you do not shut down the environment before backing up the files. To avoid an inconsistent backup, do not make any configuration changes until the backup is completed. To ensure that no changes are made in the WebLogic Server domain, lock the WebLogic Server configuration, as described in "Locking the WebLogic Server Configuration" in the *Oracle Fusion Middleware Administrator's Guide*.

During an online backup, applications can continue to run during the backup, so your business is not affected.

15.2.3 Types of Backups

You should back up your Oracle Fusion Applications file system and your Oracle Fusion Applications databases.

For the Oracle Fusion Applications file system, you can perform a full backup or you can perform a partial backup. See [Figure 1–6](#) for a graphic of the directory structure in the Oracle Fusion Applications file system.

To perform a full backup of the file system, you should back up the **binary files**, as well as the **configuration files**.

Binary files are static files and directories that do not change frequently. These include:

- The Applications base directory, which is the top-level directory containing Oracle Fusion Applications, Middleware homes, and Oracle homes.
- The Middleware home (*MW_HOME*). A Middleware home consists of a WebLogic Server home, an Oracle Common home, and optionally one or more Oracle homes and one or more Oracle instances.
- OraInventory
- On UNIX, the `OraInst.loc` file, which is located in the following directory:

(Linux and IBM AIX) /etc
(Other UNIX systems) /var/opt/oracle

- On UNIX, the `oratab` file, which is located in the following directory:

/etc

- The `beahomelist` file, which is located at:

(UNIX) `user_home/boa/beahomelist`
(Windows) `C:\boa\beahomelist`

- On Windows, the following registry key:

`HKEY_LOCAL_MACHINE\Software\oracle`

In addition, for system components, such as Oracle HTTP Server, you must back up the following Windows Registry key:

`HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services`

Configuration files are those files that change frequently. Back up these files when you perform a full backup and on a regular basis. Configuration files include:

- The Applications configuration directory, which is the top-level directory containing domains and Oracle instances.
- Domain directories of the Administration Server and the Managed Servers. The Oracle Fusion Applications environment can consist of multiple domains (for example, `CRMDomain` and `FinancialDomain`). Each domain consists of an Administration Server and one or more Managed Servers.

Unless stated in the backup recommendations for a particular type of Oracle Fusion application, you do not need to back up Managed Server directories separately because the Administration Server contains information about all of the Managed Servers in its domain.

- All Oracle instance homes, which reside by default in the `MW_HOME`, but can be configured to be in a different location.

For the Oracle Fusion Applications database and related databases, you can perform full or incremental backups. You use Oracle Recovery Manager to back up an Oracle Database instance.

Note that you must keep the databases synchronized when you restore them.

15.2.4 Recommended Backup Strategy

This section outlines the recommended strategy for performing backups. Using this strategy ensures that you can perform the recovery procedures in this book.

Note: Store your backups in a secure location, that is, not on the same hardware that contains your Oracle Fusion Applications environment.

- **Perform a full offline backup:** Back up the binary files and directories and the configuration files described in [Section 15.2.3](#). If the Applications base directory is shared, you only need to back it up once. If the Applications base directory is not shared, back it up on each host in your Oracle Fusion Applications environment. Perform a full offline backup at the following times:
 - Immediately after you install Oracle Fusion Applications
 - Immediately after an operating system software upgrade.
 - Immediately before upgrading your Oracle Fusion Applications environment
 - Immediately after upgrading your Oracle Fusion Applications environment
 - Immediately before patching your Oracle Fusion Applications environment.
- **Perform an online backup of configuration files:** Back up the configuration files described in [Section 15.2.3](#). Backing up the configuration files enables you to restore your environment to a consistent state as of the time of your most recent configuration and metadata backup. You can back up the configuration files at the following granularity:
 - The Applications configuration directory.
 - The domain.
 - The instance. You should back up the files in the following directories:
 - For the CRM Domain: `instance/domains/host_name/CRMDomain`
 - For the HCM Domain: `instance/domains/host_name/HCMDomain`
 - For the FIN Domain: `instance/domains/host_name/FinancialDomain`
 - For the PRJ Domain: `instance/domains/host_name/PRJDomain`
 - For the Webtier home: `instance/CommonDomain_webtier`
 - For the BI Instance: `instance/BIInstance`

To avoid an inconsistent backup, do not make any configuration changes until the backup completes.

Perform an online backup of configuration files at the following times:

- On a regular basis. Oracle recommends that you back up configuration files nightly.
 - Prior to making configuration changes to an Administration Server, a Managed Server, or application.
 - Immediately after patching your Oracle Fusion Applications environment.
 - After making configuration changes to an Administration Server, a Managed Server, Oracle instance, or application.
 - After a major change to the deployment architecture, such as creating servers or clusters.
- **Perform a full or incremental backup of your databases:** Use RMAN to backup your databases. See the *Oracle Database Backup and Recovery User's Guide* for information about using RMAN and for suggested methods of backing up the databases.

Note the following recommendations:

- To ensure that no changes are made in the WebLogic Server domains, lock the WebLogic Server configuration for all the domains in Oracle Fusion Applications environment, as described in "Locking the WebLogic Server Configuration" in the *Oracle Fusion Middleware Administrator's Guide*.
- When you create the backup, name the archive file with a unique name. Consider appending the date and time to the name. For example, if you create a backup of the Applications base directory on March 2 2012, name the backup:

```
ApplBase_backup_030212.tar
```

15.3 Overview of Recovering Your Environment

If your environment suffers from critical failures that involve actual data corruption, data loss, or loss of host, you must recover all or part of your environment.

This section includes the following topics:

- [Tools to Use to Recover Your Environment](#)
- [Types of Recovery](#)
- [Recommended Recovery Strategies](#)

15.3.1 Tools to Use to Recover Your Environment

To recover your Oracle Fusion Applications environment, you can use:

- Oracle Enterprise Manager Cloud Control
With Cloud Control, you can restore the directories and databases that comprise a Oracle Fusion Applications installation.
- File copy utilities, such as `copy`, `xcopy` or `tar`.

When you restore the files, use your preferred tool to extract the compressed files.

For example, for online recovery on Windows, you can use `copy`; for offline recovery on Windows, you can use `copy`, `xcopy`, or `jar`. Ensure that the utility you use can support long and unicode file names and extensions. Many of the early archiving utilities did not have this support."

For example, for UNIX, you can use `tar`.

- Oracle Recovery Manager (RMAN) to recover database-based metadata repositories and Oracle Fusion Applications databases. See the *Oracle Database Backup and Recovery User's Guide* for information about using RMAN to recover a database.

15.3.2 Types of Recovery

Recovery strategies enable you to recover from critical failures that involve actual data corruption, data loss, or loss of host. Depending on the type of loss, you can recover your Oracle Fusion Applications environment in part or in full. You can recover the following:

- The Applications base directory
- The Oracle WebLogic Server domains containing Oracle Fusion Applications
- The Administration Server for the Oracle Fusion Applications domains

- Managed Servers
- The Middleware home
- An Oracle home
- An Oracle instance home
- Any database used by Oracle Fusion Applications. You use Oracle Recovery Manager (RMAN) to recover an Oracle Database instance. See *Oracle Database Backup and Recovery User's Guide* for information about using RMAN to recover a database.

Note that with Oracle Fusion Applications, you can install all or some of the files on shared storage. Some of the procedures will differ depending on whether all files are on shared storage. The options are:

- All files are on a shared file system.
- All binary files are on a shared file system as well as all Administration Servers. Managed Servers are on local file systems.

15.3.3 Recommended Recovery Strategies

Note the following key points about recovery:

- All or part of your Oracle Fusion Applications environment must be offline while you are performing recovery.

Stop the relevant processes. The processes you stop depends on the granularity of the recovery. For example, if you are recovering only one domain, shut down the corresponding Administration Server and Managed Servers.

- Rename existing files and directories before you begin restoring the files from backup so that you do not unintentionally override necessary files.
- Although, in some cases, it may appear that only one or two files are lost or corrupted, you should restore the directory structure for the entire element, such as an Oracle instance home or a component, rather than just restoring one or two files. In this way, you are more likely to guarantee a successful recovery.
- If you need to recover a database, perform a complete recovery to recover the database to the most current state. However, there may be some situations where you do not have all the required logs to accomplish complete recovery. In that situation, perform an incomplete recovery (point-in-time recovery) to recover the database as close to the current time as possible. (You can use point-in-time recovery if the database is configured in Archive Log Mode. This is typically a time right before the database failure occurred.) Oracle recommends using Archive Log Mode for production databases.

You must keep the databases synchronized. See [Section 15.9.9](#) for procedures for recovering them to the same point in time and reconciling differences.

15.4 Prerequisites for Using Cloud Control to Back Up or Restore Your Environment

Before you back up or restore an Oracle Fusion Applications environment using Oracle Enterprise Manager Cloud Control, ensure the following prerequisites are met:

- Cloud Control 12c Agents must be installed on all hosts in the Oracle Fusion Applications environment where components (such as Oracle WebLogic Server

Administration and Managed Servers) are running, and the hosts must be discovered in Cloud Control. (Although it is possible to monitor an Oracle Fusion Applications environment without having Agents on every host, backup and restore requires a local Agent on each host). For more information on installing and configuring Cloud Control, refer to the *Oracle Enterprise Manager Cloud Control Basic Installation Guide*.

- An Oracle Secure Backup domain target must be discovered in Cloud Control. This requires an Oracle Secure Backup server installation, and all hosts that are part of the Oracle Fusion Applications environment must be configured as Oracle Secure Backup client hosts. For more information on installing Oracle Secure Backup, refer to the *Oracle Secure Backup Installation and Configuration Guide*.
- A Fusion Instance target representing the Oracle Fusion Applications environment must be discovered in Cloud Control. The Fusion Instance target must be discovered through a local Agent that is running on one of the hosts in the Oracle Fusion Applications environment.
- You must create a backup configuration using Cloud Control. A backup configuration contains the settings that will be used for the database and file backups, and is required to perform an Oracle Fusion Applications backup. See [Section 15.6.1.1](#) for more details on creating a backup configuration.
- Before you perform a backup or restore through Cloud Control, you must properly shut down all Oracle WebLogic Server instances and other components in the Oracle Fusion Applications environment. Refer to [Section 15.9](#) and [Section 15.10](#) for details on what components must be shut down and how to shut them down. Optionally, scripts that perform the required shutdown of components on each host can be invoked as user-specified pre-backup or pre-restore scripts within the Cloud Control backup procedure. (See [Section 15.6.1.2](#) for more details on user-specified scripts.)
- If the Oracle Fusion Applications repository or other databases are included in an Oracle Fusion Applications restore, the databases must be started in mounted mode prior to submitting the restore procedure. If a complete restore of the databases is required, it may be necessary to first perform a SPFILE or controlfile restore or both for the individual databases through the database Perform Recovery wizard, which can be accessed from the Availability menu on the Cloud Control database home page. (Go to the home page of the database that must be restored, select **Availability**, then select **Backup and Recovery**, then select **Perform Recovery**.)

For information about using Cloud Control to back up and recover your Oracle Fusion Applications environment, see:

- [Performing a Backup Using Cloud Control](#)
- [Recovering Using Cloud Control](#)

15.5 Backup and Recovery Recommendations for Oracle Fusion Applications

The topics in this section describe backup and recovery recommendations for specific Oracle Fusion Applications components and for components that are related to Oracle Fusion Applications.

The topics include information about configuration files for each application or component. Note that the list of files is not an exhaustive list. You do not back up or recover the individual files. Generally, you back up or recover the Applications base

directory, the Applications configuration directory, Middleware home, domain, Oracle home, or an Oracle instance, as described in [Section 15.6](#), [Section 15.9](#), and [Section 15.10](#).

The configuration files and database schemas are for information purposes only. You must back up the entire database.

This section includes the following topics:

- [Backup and Recovery Recommendations for Oracle Fusion Customer Relationship Management](#)
- [Backup and Recovery Recommendations for Oracle Fusion Financials](#)
- [Backup and Recovery Recommendations for Oracle Fusion Human Capital Management](#)
- [Backup and Recovery Recommendations for Oracle Fusion Supply Chain Management](#)
- [Backup and Recovery Recommendations for Oracle Fusion Project](#)
- [Backup and Recovery Recommendations for Oracle Fusion Procurement](#)
- [Backup and Recovery Recommendations for Oracle Fusion Incentive Compensation](#)
- [Backup and Recovery Recommendations for Oracle Fusion Applications Technology](#)
- [Backup and Recovery Recommendations for Oracle Fusion Setup](#)
- [Backup and Recovery Recommendations for Oracle Enterprise Scheduler](#)
- [Backup and Recovery Recommendations for Oracle Enterprise Crawl and Search Framework](#)
- [Backup and Recovery Recommendations for Oracle Authorization Policy Manager](#)
- [Backup and Recovery Recommendations for Java Servers for Oracle E-Mail and Web Marketing](#)

15.5.1 Backup and Recovery Recommendations for Oracle Fusion Customer Relationship Management

This section describes the Oracle Fusion Customer Relationship Management data that must be backed up and restored.

Configuration Files

Most configuration data is stored in the Oracle WebLogic Server domain.

In addition, Oracle Fusion Customer Relationship Management uses configuration data for Java servers for Oracle E-Mail and Web Marketing. See [Section 15.5.13](#) for information about backing up and recovering Oracle E-Mail and Web Marketing.

Dependencies on Oracle Fusion Middleware Components

Oracle Metadata Services, Oracle Business Intelligence metadata, Oracle Real-Time Decisions, an LDAP provider such as Oracle Internet Directory, Oracle WebCenter Portal (including tagging, group spaces, and the Oracle WebCenter Portal schemas), portlet metadata, Oracle SOA Suite, Oracle ADF user interface, Oracle Essbase, Oracle WebCenter Content, Oracle Enterprise Crawl and Search Framework metadata and data, including Oracle Secure Enterprise Search, Oracle Product Data Quality

Dependencies on Third-Party Products

None

Database Repository Dependencies

The databases containing Oracle Fusion Applications, the LDAP store, and Oracle Fusion Middleware schemas, including Oracle WebCenter Content

Backup Recommendations

Back up the Applications base directory and the Administration Server domain in which Oracle Fusion Customer Relationship Management is deployed. Back up the standalone Java servers used for Oracle E-Mail and Web Marketing, as described in [Section 15.5.13](#).

Back up the databases containing the Oracle Fusion Applications, the LDAP store, and Oracle Fusion Middleware schemas, including those for Oracle WebCenter Content.

Back up Oracle WebCenter Content, as described in "Backup and Recovery Recommendations for Oracle WebCenter Content" in the *Oracle Fusion Middleware Administrator's Guide*.

Recovery Recommendations

Recover the Managed Server to which the application is deployed. If needed, recover the standalone Java servers used for Oracle E-Mail and Web Marketing, as described in [Section 15.5.13](#).

Depending upon the extent of failure, recovery should be performed at the desired granularity. For the steps to recover Oracle Fusion Customer Relationship Management, including for loss of host, see [Section 15.9.10.1](#).

Recover Oracle WebCenter Content, as described in "Backup and Recovery Recommendations for Oracle WebCenter Content" in the *Oracle Fusion Middleware Administrator's Guide*.

Recover the databases containing the Oracle Fusion Applications, Oracle Fusion Middleware schemas, including those for Oracle WebCenter Content.

If needed, perform a complete recovery to recover the databases to the most current state. However, there may be some situations where you do not have all the required logs to accomplish complete recovery. In that situation, perform an incomplete recovery (point-in-time recovery) to recover the database to as close to the current time as possible. Then, reconcile the databases, as described in [Section 15.9.9](#).

15.5.2 Backup and Recovery Recommendations for Oracle Fusion Financials

This section describes the Oracle Fusion Financials data that must be backed up and restored.

Configuration Files

Configuration data is stored in the Oracle WebLogic Server domain.

Dependencies on Oracle Fusion Middleware Components

Oracle Application Development Framework, Oracle BPEL Process Manager, Oracle WebCenter Portal (portlets), Oracle Metadata Services, Oracle WebCenter Content: Imaging, Oracle Enterprise Scheduler, Oracle BI Enterprise Edition, Oracle Essbase, and Applications Technology Group

Dependencies on Third-Party Products

None

Database Repository Dependencies

The databases containing the Oracle Fusion Applications, the LDAP store, Oracle Essbase, and Oracle Fusion Middleware schemas

Backup Recommendations

Back up the Applications base directory and the Administration Server domain in which Oracle Fusion Financials is deployed.

Back up Oracle WebCenter Content, as described in "Backup and Recovery Recommendations for Oracle WebCenter Content" in the *Oracle Fusion Middleware Administrator's Guide*.

Back up the databases containing the Oracle Fusion Applications, Oracle Essbase, and Oracle Fusion Middleware schemas, including Oracle WebCenter Content.

Recovery Recommendations

Recover the Managed Server to which the application is deployed. For information about recovering Oracle Essbase, see "Recovering Oracle Essbase After Loss of Host" in the *Oracle Fusion Middleware Administrator's Guide*.

Depending upon the extent of failure, recovery should be performed at the desired granularity. See [Section 15.9](#).

Recover Oracle WebCenter Content, as described in "Backup and Recovery Recommendations for Oracle WebCenter Content" in the *Oracle Fusion Middleware Administrator's Guide*.

Recover the databases to the most recent point in time, if needed.

If needed, perform a complete recovery to recover the databases to the most current state. However, there may be some situations where you do not have all the required logs to accomplish complete recovery. In that situation, perform an incomplete recovery (point-in-time recovery) to recover the database to as close to the current time as possible. Then, reconcile the databases, as described in [Section 15.9.9](#).

15.5.3 Backup and Recovery Recommendations for Oracle Fusion Human Capital Management

This section describes the Oracle Fusion Human Capital Management data that must be backed up and restored.

Configuration Files

Most configuration data is stored in the Oracle WebLogic Server domain.

Dependencies on Oracle Fusion Middleware Components

Oracle Application Development Framework, Oracle BPEL Process Manager, Oracle WebCenter Portal (portlets), Oracle Metadata Services, Oracle Enterprise Scheduler, Oracle Business Intelligence Publisher, an LDAP provider such as Oracle Internet Directory, Oracle Fusion Middleware Extensions for Applications, Oracle WebCenter Content

Dependencies on Third-Party Products

None

Database Repository Dependencies

The databases containing Oracle Fusion Applications, the LDAP store, and Oracle Fusion Middleware schemas, including Oracle WebCenter Content

Backup Recommendations

Back up the Applications base directory and the Administration Server domain in which Oracle Fusion Human Capital Management is deployed.

Back up Oracle WebCenter Content, as described in "Backup and Recovery Recommendations for Oracle WebCenter Content" in the *Oracle Fusion Middleware Administrator's Guide*.

Back up the databases containing the LDAP store, Oracle Fusion Applications, Oracle Fusion Middleware schemas, including those for Oracle WebCenter Content.

Recovery Recommendations

Recover the Managed Server to which the application is deployed.

Depending upon the extent of failure, recovery should be performed at the desired granularity. See [Section 15.9](#).

Recover Oracle WebCenter Content, as described in "Backup and Recovery Recommendations for Oracle WebCenter Content" in the *Oracle Fusion Middleware Administrator's Guide*.

Recover the databases to the most recent point in time, if needed. Note that the databases, including the LDAP store, the database containing Oracle Fusion Human Capital Management schemas and the Oracle Fusion Middleware schemas, including those for Oracle WebCenter Content must be kept synchronized. If you restore one, restore the others to the same point in time.

There may be some situations where you do not have all the required logs to accomplish complete recovery. In that situation, perform an incomplete recovery (point-in-time recovery) to recover the database to as close to the current time as possible. Then, reconcile the databases, as described in [Section 15.9.9](#).

15.5.4 Backup and Recovery Recommendations for Oracle Fusion Supply Chain Management

This section describes the Oracle Fusion Supply Chain Management data that must be backed up and restored.

Configuration Files

Configuration data is stored in the Oracle WebLogic Server domain.

Dependencies on Oracle Fusion Middleware Components

Oracle Application Development Framework, Oracle SOA Suite, Oracle Enterprise Scheduler

Dependencies on Third-Party Products

None

Database Repository Dependencies

The databases containing Oracle Fusion Applications, the LDAP store, and Oracle Fusion Middleware schemas, including Oracle WebCenter Content

Backup Recommendations

Back up the Applications base directory and the Administration Server domain in which Oracle Fusion Supply Chain Management is deployed.

Back up Oracle WebCenter Content, as described in "Backup and Recovery Recommendations for Oracle WebCenter Content" in the *Oracle Fusion Middleware Administrator's Guide*.

Back up the databases containing the Oracle Fusion Applications, the LDAP store, and Oracle Fusion Middleware schemas, including those for Oracle WebCenter Content.

Recovery Recommendations

Recover the Managed Server to which the application is deployed.

Recover Oracle WebCenter Content, as described in "Backup and Recovery Recommendations for Oracle WebCenter Content" in the *Oracle Fusion Middleware Administrator's Guide*.

Recover the databases to the most recent point in time, if needed.

Depending upon the extent of failure, recovery should be performed at the desired granularity. See [Section 15.9](#). For the steps specific to recovering from loss of host, see [Section 15.10.6.1](#).

15.5.5 Backup and Recovery Recommendations for Oracle Fusion Project

This section describes the Oracle Fusion Project data that must be backed up and restored.

Configuration Files

Most configuration data is stored in the Oracle WebLogic Server domain. In addition, configuration data is stored in the Oracle Essbase database.

Dependencies on Oracle Fusion Middleware Components

Oracle Metadata Services, an LDAP provider such as Oracle Internet Directory, Oracle WebCenter Portal (including tagging, group spaces, and the Oracle WebCenter Portal schemas), Oracle SOA Suite, Oracle Application Development Framework user interface, Oracle Enterprise Crawl and Search Framework metadata and data, Oracle Enterprise Scheduler, Oracle Essbase

Dependencies on Third-Party Products

None

Database Repository Dependencies

The databases containing the Oracle Fusion Applications, Oracle Essbase, the LDAP store, and Oracle Fusion Middleware schemas, including Oracle WebCenter Content

Backup Recommendations

Back up the Applications base directory and the Administration Server domain in which Oracle Fusion Project is deployed.

Back up Oracle WebCenter Content, as described in "Backup and Recovery Recommendations for Oracle WebCenter Content" in the *Oracle Fusion Middleware Administrator's Guide*.

Back up the Oracle Essbase database used by Oracle Fusion Project. Back up the databases containing the Oracle Fusion Applications, the LDAP store, and Oracle Fusion Middleware schemas.

Recovery Recommendations

Recover the Managed Server to which the application is deployed.

Recover Oracle WebCenter Content, as described in "Backup and Recovery Recommendations for Oracle WebCenter Content" in the *Oracle Fusion Middleware Administrator's Guide*.

Recover the databases to the most recent point in time, if needed.

Depending upon the extent of failure, recovery should be performed at the desired granularity. See [Section 15.9](#).

15.5.6 Backup and Recovery Recommendations for Oracle Fusion Procurement

This section describes the Oracle Fusion Procurement data that must be backed up and restored.

Configuration Files

Most configuration data is stored in the Oracle WebLogic Server domain. In addition, the certificate file, which is determined by the Profile Option, contains a certificate for SSL connections made to supplier web sites in order to let requisitioners search or browse catalogs.

Dependencies on Oracle Fusion Middleware Components

Oracle Application Development Framework, Oracle BPEL Process Manager, Oracle B2B, Oracle Fusion Middleware Extensions for Applications, Oracle Enterprise Scheduler, Oracle Business Intelligence Publisher, Oracle Streams Advanced Queuing (AQ) Java Message Service (JMS), JMS queues with Oracle B2B

Dependencies on Third-Party Products

CUPS software, a portable printing layer for UNIX systems

Database Repository Dependencies

The databases containing Oracle Fusion Applications, the LDAP store, and Oracle Fusion Middleware schemas, including Oracle WebCenter Content

Backup Recommendations

Back up the Applications base directory and the Administration Server domain in which Oracle Fusion Procurement is deployed. Back up the certificate file.

Back up Oracle WebCenter Content, as described in "Backup and Recovery Recommendations for Oracle WebCenter Content" in the *Oracle Fusion Middleware Administrator's Guide*.

Back up the databases containing the Oracle Fusion Applications, the LDAP store, and Oracle Fusion Middleware schemas, including those for Oracle WebCenter Content.

Recovery Recommendations

Recover the Managed Server to which the application is deployed.

Depending upon the extent of failure, recovery should be performed at the desired granularity. For the steps to recover Oracle Fusion Procurement, see [Section 15.9.10.2](#). For the steps specific to recovering from loss of host, see [Section 15.10.6.2](#).

For information about recovering JMS, see "Backup and Recovery Recommendations for Oracle WebLogic Server JMS" in the *Oracle Fusion Middleware Administrator's Guide*.

Recover Oracle WebCenter Content, as described in "Backup and Recovery Recommendations for Oracle WebCenter Content" in the *Oracle Fusion Middleware Administrator's Guide*.

Recover the databases to the most recent point in time, if needed.

If needed, perform a complete recovery to recover the databases to the most current state. However, there may be some situations where you do not have all the required logs to accomplish complete recovery. In that situation, perform an incomplete recovery (point-in-time recovery) to recover the database to as close to the current time as possible. Then, reconcile the databases, as described in [Section 15.9.9](#).

15.5.7 Backup and Recovery Recommendations for Oracle Fusion Incentive Compensation

This section describes the Oracle Fusion Incentive Compensation data that must be backed up and restored.

Configuration Files

Configuration data is stored in the Oracle WebLogic Server domain.

Dependencies on Oracle Fusion Middleware Components

Oracle Application Development Framework, Oracle SOA Suite, Oracle Data Integrator, Oracle BI Enterprise Edition, Oracle Enterprise Scheduler, Oracle Enterprise Crawl and Search Framework, Oracle Secure Enterprise Search

Dependencies on Third-Party Products

None

Database Repository Dependencies

The databases containing Oracle Fusion Applications, the LDAP store, and Oracle Fusion Middleware schemas, including Oracle WebCenter Content

Backup Recommendations

Back up the Applications base directory and the Administration Server domain in which Oracle Fusion Incentive Compensation is deployed.

Back up Oracle WebCenter Content, as described in "Backup and Recovery Recommendations for Oracle WebCenter Content" in the *Oracle Fusion Middleware Administrator's Guide*.

Back up the databases containing the Oracle Fusion Applications, the LDAP store, and Oracle Fusion Middleware schemas, including those for Oracle WebCenter Content.

Recovery Recommendations

Recover the Managed Server to which the application is deployed.

Depending upon the extent of failure, recovery should be performed at the desired granularity. See [Section 15.9](#).

Recover Oracle WebCenter Content, as described in "Backup and Recovery Recommendations for Oracle WebCenter Content" in the *Oracle Fusion Middleware Administrator's Guide*.

Recover the databases to the most recent point in time, if needed.

If needed, perform a complete recovery to recover the databases to the most current state. However, there may be some situations where you do not have all the required logs to accomplish complete recovery. In that situation, perform an incomplete recovery (point-in-time recovery) to recover the database to as close to the current time as possible. Then, reconcile the databases, as described in [Section 15.9.9](#).

15.5.8 Backup and Recovery Recommendations for Oracle Fusion Applications Technology

This section describes the Oracle Fusion Applications Technology data that must be backed up and restored.

Configuration Files

Configuration data is stored in the Oracle WebLogic Server domain.

Dependencies on Oracle Fusion Middleware Components

Oracle Application Development Framework, Oracle SOA Suite, Federated Worklist, Oracle WebCenter Portal (including portlets, tagging, group spaces, and forums), Oracle WebCenter Portal, MDS, portlet metadata, Oracle Enterprise Crawl and Search Framework metadata and data, and Oracle Secure Enterprise Search, Oracle BI Enterprise Edition, Oracle Fusion Topology Manager, Applications Technology Group

Dependencies on Third-Party Products

None

Database Repository Dependencies

The databases containing Oracle Fusion Applications, the LDAP store, and Oracle Fusion Middleware schemas, including Oracle WebCenter Content

Backup Recommendations

Back up the Applications base directory and the Administration Server domain in which Oracle Fusion Applications Technology is deployed.

Back up Oracle WebCenter Content, as described in "Backup and Recovery Recommendations for Oracle WebCenter Content" in the *Oracle Fusion Middleware Administrator's Guide*.

Back up the databases containing the Oracle Fusion Applications, the LDAP store, and Oracle Fusion Middleware schemas, including those for Oracle WebCenter Content.

Recovery Recommendations

Recover the Managed Server to which the application is deployed.

Depending upon the extent of failure, recovery should be performed at the desired granularity. See [Section 15.9](#).

Recover Oracle WebCenter Content, as described in "Backup and Recovery Recommendations for Oracle WebCenter Content" in the *Oracle Fusion Middleware Administrator's Guide*.

Recover the databases to the most recent point in time, if needed.

If needed, perform a complete recovery to recover the databases to the most current state. However, there may be some situations where you do not have all the required logs to accomplish complete recovery. In that situation, perform an incomplete recovery (point-in-time recovery) to recover the database to as close to the current time as possible. Then, reconcile the databases, as described in [Section 15.9.9](#).

15.5.9 Backup and Recovery Recommendations for Oracle Fusion Setup

This section describes the Oracle Fusion Setup data that must be backed up and restored.

Configuration Files

Configuration data is stored in the Oracle WebLogic Server domain.

Dependencies on Oracle Fusion Middleware Components

Oracle Application Development Framework, Oracle BPEL Process Manager, Oracle WebCenter Content, Oracle WebCenter Portal (portlets), Oracle Enterprise Scheduler, Identity Governance Framework

Dependencies on Third-Party Products

None

Database Repository Dependencies

The databases containing Oracle Fusion Applications, the LDAP store, and Oracle Fusion Middleware schemas, including Oracle WebCenter Content

Backup Recommendations

Back up the Applications base directory and the Administration Server domain in which Oracle Fusion Setup is deployed.

Back up Oracle WebCenter Content, as described in "Backup and Recovery Recommendations for Oracle WebCenter Content" in the *Oracle Fusion Middleware Administrator's Guide*.

Back up the databases containing the Oracle Fusion Applications, the LDAP store, and Oracle Fusion Middleware schemas, including those for Oracle WebCenter Content.

Recovery Recommendations

Recover the Managed Server to which the application is deployed.

Depending upon the extent of failure, recovery should be performed at the desired granularity. See [Section 15.9](#).

Recover Oracle WebCenter Content, as described in "Backup and Recovery Recommendations for Oracle WebCenter Content" in the *Oracle Fusion Middleware Administrator's Guide*.

Recover the databases containing the Oracle Fusion Applications, the LDAP store, and Oracle Fusion Middleware schemas, including those for Oracle WebCenter Content.

If needed, perform a complete recovery to recover the databases to the most current state. However, there may be some situations where you do not have all the required logs to accomplish complete recovery. In that situation, perform an incomplete recovery (point-in-time recovery) to recover the database to as close to the current time as possible. Then, reconcile the databases, as described in [Section 15.9.9](#).

15.5.10 Backup and Recovery Recommendations for Oracle Enterprise Scheduler

This section describes the Oracle Enterprise Scheduler data that must be backed up and restored.

Configuration Files

Configuration data is stored in the Oracle WebLogic Server domain.

Dependencies on Oracle Fusion Middleware Components

None

Dependencies on Third-Party Products

None

Database Repository Dependencies

The databases containing Oracle Fusion Applications, Oracle Enterprise Scheduler, LDAP store, and Oracle Fusion Middleware schemas, including Oracle WebCenter Content

Backup Recommendations

Back up the Oracle home and the domain home.

Back up the databases containing the Oracle Fusion Applications, the LDAP store, and Oracle Fusion Middleware schemas.

Recovery Recommendations

Recover the domain home and the Oracle home, as needed.

Depending upon the extent of failure, recovery should be performed at the desired granularity. For the steps to recover Oracle Enterprise Scheduler, including for loss of host, see [Section 15.9.11.2](#).

Recover the databases to the most recent point in time, if needed.

If needed, perform a complete recovery to recover the databases to the most current state. However, there may be some situations where you do not have all the required logs to accomplish complete recovery. In that situation, perform an incomplete recovery (point-in-time recovery) to recover the database to as close to the current time as possible. Then, reconcile the databases, as described in [Section 15.9.9](#).

15.5.11 Backup and Recovery Recommendations for Oracle Enterprise Crawl and Search Framework

This section describes the Oracle Enterprise Crawl and Search Framework data that must be backed up and restored.

Configuration Files

Configuration data is stored in the Oracle WebLogic Server domain.

Dependencies on Oracle Fusion Middleware Components

Oracle Application Development Framework

Dependencies on Third-Party Products

None

Database Repository Dependencies

The database containing Oracle Fusion Applications schemas

Backup Recommendations

Back up the domain in which Oracle Enterprise Crawl and Search Framework is deployed.

Back up the database containing the Oracle Fusion Applications schemas.

Recovery Recommendations

Recover the domain in which Oracle Enterprise Crawl and Search Framework is deployed.

Depending upon the extent of failure, recovery should be performed at the desired granularity. For the steps to recover Oracle Enterprise Crawl and Search Framework, including for loss of host, see [Section 15.9.11.3](#).

Recover the database containing the Oracle Fusion Applications schemas.

If needed, perform a complete recovery to recover the databases to the most current state. However, there may be some situations where you do not have all the required logs to accomplish complete recovery. In that situation, perform an incomplete recovery (point-in-time recovery) to recover the database to as close to the current time as possible. Then, reconcile the databases, as described in [Section 15.9.9](#).

15.5.12 Backup and Recovery Recommendations for Oracle Authorization Policy Manager

This section describes the Oracle Authorization Policy Manager data that must be backed up and restored.

Configuration Files

The configuration files are located in the Oracle instance home.

Dependencies on Oracle Fusion Middleware Components

An LDAP provider, such as Oracle Internet Directory

Dependencies on Third-Party Products

None

Database Repository Dependencies

The databases used by Oracle Authorization Policy Manager and the LDAP store

Backup Recommendations

Back up the Oracle Authorization Policy Manager Domain home and its Oracle home.

Back up the database used by Oracle Authorization Policy Manager and the LDAP store.

Recovery Recommendations

Recover the domain in which Oracle Authorization Policy Manager is deployed.
Recover the Oracle home, if necessary.

Depending upon the extent of failure, recovery should be performed at the desired granularity. For the steps to recover Oracle Authorization Policy Manager, including for loss of host, see [Section 15.9.11.4](#).

Recover the databases used by Oracle Authorization Policy Manager and the LDAP store, if needed, to the same point in time.

If needed, perform a complete recovery to recover the databases to the most current state. However, there may be some situations where you do not have all the required logs to accomplish complete recovery. In that situation, perform an incomplete recovery (point-in-time recovery) to recover the database to as close to the current time as possible. Then, reconcile the databases, as described in [Section 15.9.9](#).

15.5.13 Backup and Recovery Recommendations for Java Servers for Oracle E-Mail and Web Marketing

Oracle E-Mail and Web Marketing is provided with Oracle Fusion Customer Relationship Management. Oracle E-Mail and Web Marketing provides three components: Email Sending Daemon (ESD), Click Thru Daemon (CTD), and Bounce Handling Daemon (BHD).

This section describes the Oracle E-Mail and Web Marketing data that must be backed up and restored.

Configuration Files

Configuration data is stored in the Oracle WebLogic Server domain.

Dependencies on Oracle Fusion Middleware Components

None

Dependencies on Third-Party Products

None

Database Repository Dependencies

None

Backup Recommendations

Back up the Managed Servers to which ESD and CTD are deployed. Back up the directory for BHD.

Recovery Recommendations

Restore the Managed Servers to which ESD and CTD are deployed. Restore the directory for BHD.

Depending upon the extent of failure, recovery should be performed at the desired granularity. For the steps to recover Oracle E-Mail and Web Marketing, including for loss of host, see [Section 15.9.11.5](#).

15.6 Performing a Backup

You can perform a full offline backup or an online or offline backup of configuration files.

This section includes the following topics:

- [Performing a Backup Using Cloud Control](#)

- [Performing a Full Offline Backup Using the Command Line](#)
- [Performing an Online Backup of Configuration Files Using the Command Line](#)

15.6.1 Performing a Backup Using Cloud Control

You can use Cloud Control to back up Oracle Fusion Applications, as described in the following topics:

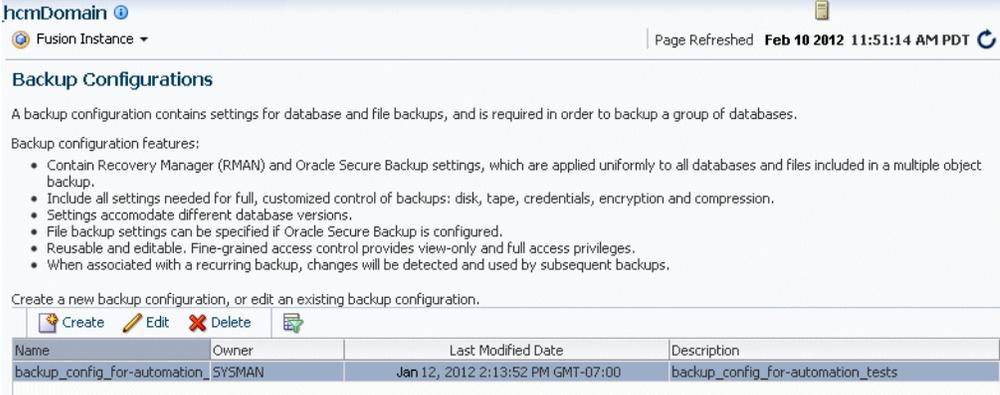
- [Configuring Cloud Control Backups](#)
- [Backing Up Oracle Fusion Applications Using Cloud Control](#)

15.6.1.1 Configuring Cloud Control Backups

You must create a backup configuration before you perform a backup or restore using Cloud Control. A backup configuration contains settings for database and file backups. You create the configuration using the Backup Configurations page. You can create multiple backup configurations, using different settings. Then, you can use the configurations in subsequent backup and restore operations.

To create a new backup configuration using Cloud Control:

1. From the **Targets** menu, select **Fusion Applications**.
The Fusion Applications target home page is displayed.
2. Click the **Fusion Instance** that you want to back up.
The Fusion Instance home page is displayed.
3. From the **Fusion Instance** menu, choose **Control**, then **Backup Configurations**.
The Backup Configurations page is displayed, as shown in the following figure:



hcmDomain Fusion Instance Page Refreshed Feb 10 2012 11:51:14 AM PDT

Backup Configurations

A backup configuration contains settings for database and file backups, and is required in order to backup a group of databases.

Backup configuration features:

- Contain Recovery Manager (RMAN) and Oracle Secure Backup settings, which are applied uniformly to all databases and files included in a multiple object backup.
- Include all settings needed for full, customized control of backups: disk, tape, credentials, encryption and compression.
- Settings accommodate different database versions.
- File backup settings can be specified if Oracle Secure Backup is configured.
- Reusable and editable. Fine-grained access control provides view-only and full access privileges.
- When associated with a recurring backup, changes will be detected and used by subsequent backups.

Create a new backup configuration, or edit an existing backup configuration.

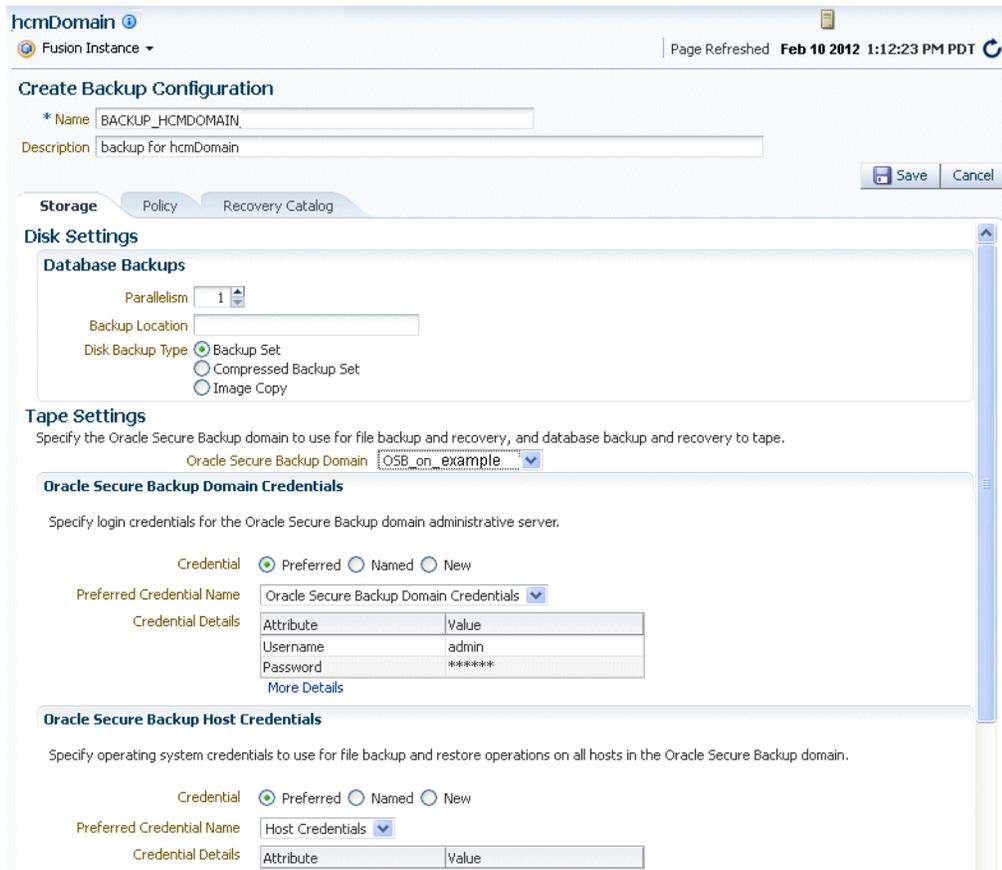
Create Edit Delete

Name	Owner	Last Modified Date	Description
backup_config_for-automation	SYSMAN	Jan 12, 2012 2:13:52 PM GMT-07:00	backup_config_for-automation_tests

4. Click **Create**.
The Create Backup Configurations page is displayed. It contains three tabs: **Storage**, **Policy**, and **Recovery Catalog**. You can use the default settings or customize individual settings. For more information on the database backup settings in these pages, refer to the *Oracle Database Backup and Recovery Reference*. (For information on the Oracle Secure Backup settings in these pages, refer to the *obtool* backup command options in the *Oracle Secure Backup Reference*.)
5. On the **Storage** tab, you specify settings related to disk and tape backup. Because Cloud Control always backs up Oracle Fusion Applications to tape, the disk settings are not applicable.

To specify tape settings, first select an **Oracle Secure Backup Domain**, after which separate tape settings for database and files are shown. Database settings include datafile and archived log copies, backup type, and media management parameters. File settings include media family and devices. Specify **Oracle Secure Backup** domain and host credentials.

The following figure shows the **Storage** tab:



6. On the **Policy** tab, you specify policy settings for database and tape backups. Database settings include backup retention, compression, and encryption. Separate database encryption settings are specified for disk and tape backups. File settings include Oracle Secure Backup job and encryption settings.

The following figure shows the **Policy** tab:

hcmDomain
Fusion Instance

Page Refreshed Feb 10 2012 1:12:23 PM PDT

Create Backup Configuration

* Name: BACKUP_HCMDOMAIN
Description: backup for hcmDomain

Save Cancel

Storage Policy Recovery Catalog

Database Backups

Maximum Backup Piece (File) Size: 20 MB

Backup Policy

Backup all files including read-only and offline files in the whole database backup even though they have not changed since the last backup
 Override tablespace exclusion

Retention

Retain backups forever
 Retain backups for specified number of days: 30

Compression Algorithm

Specify the compression algorithm that will be used for both disk and tape compressed backup sets.
The compression algorithm choices are applicable to 11.2 and higher database versions with a COMPATIBLE setting of 11.2 or higher. If 11.1 databases or databases with a COMPATIBLE setting lower than 11.2 are included in a backup that uses this backup configuration, the algorithm choice will be mapped to an available compression algorithm for those databases.

TIP
Algorithm Name: Low
 Optimize for Load

Encryption

Disk Settings

Do not encrypt backups
 Encrypt backups using Recovery Manager (RMAN) encryption

Oracle Secure Backup Tape Settings

Do not encrypt backups

7. If you are using a recovery catalog for database backups, specify the catalog on the **Recovery Catalog** tab, which is shown in the following figure:

hcmDomain
Fusion Instance

Page Refreshed Feb 10 2012 1:12:23 PM PDT

Create Backup Configuration

* Name: BACKUP_HCMDOMAIN
Description: backup for hcmDomain

Save Cancel

Storage Policy Recovery Catalog

Recovery Catalog

Specify the recovery catalog that will be used for all backups that will use this configuration. Each database that will be backed up using this configuration must be individually registered with the recovery catalog. If a catalog is not specified, the respective controlfiles will be used.

TIP Create and configure recovery catalogs from [Recovery Catalogs](#).

Use Control File
 Use Recovery Catalog

Recovery Catalog: [Dropdown]

8. Click **Save**.

15.6.1.2 Backing Up Oracle Fusion Applications Using Cloud Control

By default, Cloud Control backs up the core directories and databases that comprise a Oracle Fusion Applications installation. You can include additional files and databases in the backup. It also allows you to execute user-defined pre-backup and post-backup scripts.

Because Cloud Control uses Oracle Secure Backup to perform the backups, all backups are done to tape, which is the media supported by Oracle Secure Backup.

To back up a Oracle Fusion Applications environment using Cloud Control:

1. Ensure that you have met the prerequisites, as described in [Section 15.4](#).
2. From the **Targets** menu, choose **Fusion Applications**.

The Fusion Applications target home page is displayed.

3. Click the **Fusion Instance** that you want to back up.

The Fusion Instance home page is displayed.

4. From the **Fusion Instance** menu, choose **Control**, then **Schedule Backup**.

The Schedule Backup wizard is displayed, with the first step, Backup Scope, as shown in the following figure:

5. In the Backup Scope step:

- a. Select the **Oracle Secure Backup** domain to use for the backup.

- b. Select the scope of the backup. You can select:

- **Software and Configuration Directories:** Backs up all directories in the Oracle Fusion Applications installation, including the Applications Home directory (the root of the Oracle Fusion Applications installation) and any other directories (for example, Domain homes) that may be installed in a different location. Cloud Control will determine the distinct file systems that comprise the installation across all hosts in the environment, and will include the top-level directories needed to constitute a complete backup.
- **Configuration Directories:** Backs up only the configuration-related directories, including the Domain home of the Administrative Server for each Oracle WebLogic Server domain. You can select the following options:

All: Back up all configuration directories.

Custom: Narrow the scope of the backup by selecting specific Oracle Fusion Applications components. (On the next page, you can further select which of their associated directories to include in the backup.)

Click **Next**.

6. In the Directories step, specify the Oracle Fusion Applications directories to include in the backup. By default, the directories that correspond to the backup scope choice (made on the Backup Scope page) are included. If you chose the Custom option, you can remove directories from the list. (The list of directories that correspond to the backup scope are automatically determined by Cloud Control. It is not possible to individually select or remove subdirectories from the backup.) Click **Add** to specify additional files or directories to include in the backup from outside the core Oracle Fusion Applications installation. Any directories or files that reside on a discovered Oracle Secure Backup client host can be included.

Click **Next**.

7. In the Databases step, click **Add** to specify additional databases to add to the backup. By default, the Oracle Fusion Applications repository database is included in the backup.

Click **Next**.

8. In the Configuration step, select the backup configuration to use for the backup. (See the [Section 15.6.1.1](#) for details on creating a backup configuration.)

Click **Next**.

9. In the Options step:

- a. Select the backup type:

- **Full:** A full backup of the databases and files will be performed. If the database portion of the backup is to be used as the base of an incremental backup strategy, select the **Use as the base of an incremental backup strategy** option.
- **Incremental:** An incremental backup of the databases and files will be performed.

See the *Oracle Database Backup and Recovery User's Guide* for a description of the backup types.

b. Select the backup mode:

- **Online:** The databases will remain running during the backup. You must manually quiesce or shut down the Oracle Fusion Applications processes prior to submitting the backup, as this is not done automatically by Cloud Control.
- **Offline:** The databases will be shutdown automatically prior to the backup. You must manually shut down the Oracle Fusion Applications processes prior to submitting the backup, as this is not done automatically by Cloud Control.

See [Section 3.3](#) for information about stopping Oracle Fusion Applications processes. Alternatively, you can specify scripts to shutdown the processes in Step 10.

Click **Next**.

- 10.** In the **Scripts** step, specify the host name and path of scripts to be run before and after the backup. Post-backup scripts can be specified to run upon success or failure of the backup. The scripts can be run on any discovered host. The scripts must be executable files that exist on the specified hosts. Any type of executable file can be specified.

Click **Next**.

- 11.** In the **Credentials** step, if one or more databases is included in the backup, specify the database and host credentials that will be used to back up the databases.

Click **Next**.

- 12.** In the **Schedule** step, specify the procedure name, description, and scheduling information.

Click **Next**.

- 13.** In the **Review** step, verify all the input specified in the wizard. Make changes as needed and click **Submit**. A Cloud Control deployment procedure is submitted to perform the backup, and a confirmation page appears with a link to the deployment procedure status page.

15.6.2 Performing a Full Offline Backup Using the Command Line

To perform a full offline backup, you copy the file system artifacts and database repositories corresponding to Oracle Fusion Applications. You use your preferred tool for archiving and compressing, as described in [Section 15.2](#). Ensure that the tool you are using preserves the permissions of the files.

To perform a full offline backup:

- 1.** Stop all processes. See [Section 3.3.2.2](#).
- 2.** Back up the Applications base directory on all hosts. For example:

```
(UNIX) tar -cpf ApplBase_backup_030212.tar APPLICATIONS_BASE/*
(Windows) jar cf ApplBase_backup_030212.tar APPLICATIONS_BASE\*
```

3. Back up the Applications configuration directory on all hosts. For example:

```
(UNIX) tar -cpf ApplConfig_backup_030212.tar APPLICATIONS_CONFIG/*
(Windows) jar cf ApplConfig_backup_030212.tar APPLICATIONS_CONFIG\*
```

4. If a domain is not located within the Applications configuration home, back up the domains separately. This backs up the Managed Servers that are running Java components such as Oracle SOA Suite and Oracle WebCenter Portal.

For example:

```
(UNIX) tar -cpf domain_home_backup_030212.tar APPLICATIONS_CONFIG/instance/domains/domain_name/*
(Windows) jar cf domain_home_backup_030212.jar APPLICATIONS_CONFIG\instance\domains\domain_name\*
```

In most cases, you do not need to back up the Managed Server directories separately, because the Administration Server domain contains information about the Managed Servers in its domain. See [Section 15.2.4](#) for information about what you need to back up.

5. If the Oracle instance home is not located within the Applications base home, back up the Oracle instance home. The Oracle instance home contains configuration information about system components, such as the Global Order Promising component of Oracle Fusion Supply Chain Management.

For example:

```
(UNIX) tar -cpf instance_home_backup_030212.tar ORACLE_INSTANCE/*
(Windows) jar cf instance_home_backup_030212.jar ORACLE_INSTANCE\*
```

6. If a Managed Server is not located within the domain, back up the Managed Server directory. For example:

```
(UNIX) tar -cpf man_server1_backup_030212.tar APPLICATIONS_CONFIG/instance/domains/domain_name/servers/server_name/*
(Windows) jar cf man_server1_backup_030212.jar APPLICATIONS_CONFIG\instance\domains\domain_name\servers\server_name\*
```

7. Back up the OraInventory directory. For example:

```
(UNIX) tar -cpf Inven_home_backup_030212 /scratch/oracle/OraInventory
(Windows) jar cf Inven_home_backup_030212.jar C:\Program Files\Oracle\Inventory
```

8. On UNIX, back up the OraInst.loc file, which is located in the following directory:

```
(Linux and IBM AIX) /etc
(Other UNIX systems) /var/opt/oracle
```

9. On UNIX, back up the oratab file, which is located in the following directory:

```
/etc
```

Note that the oratab file is located on the database host.

10. Back up the databases using the Oracle Recovery Manager (RMAN). For detailed steps, see the *Oracle Database Backup and Recovery User's Guide*.

11. On Windows, export the following registry key:

```
HKEY_LOCAL_MACHINE\Software\oracle
```

In addition, for system components, export the following Windows Registry key:

```
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services
```

To export a key, use the following command:

```
regedit /E filename key
```

For example:

```
regedit /E C:\oracleregistry.reg HKEY_LOCAL_MACHINE/oracle
```

You can also use the Registry Editor to export the key. See the Registry Editor Help for more information.

12. Create a record of your Oracle Fusion Applications environment. See [Section 15.7](#).

15.6.3 Performing an Online Backup of Configuration Files Using the Command Line

You should perform a backup of configuration files on a regular basis and at the times described in [Section 15.2.4](#).

To back up configuration files:

1. To avoid an inconsistent backup, do not make any configuration changes until the backup is completed. To ensure that no changes are made in the WebLogic Server domain, lock the WebLogic Server configuration, as described in "Locking the WebLogic Server Configuration" in the *Oracle Fusion Middleware Administrator's Guide*.
2. Back up the domain directories. This backs up the Managed Servers that are running Java components such as Oracle SOA Suite and Oracle WebCenter Portal. For example:

```
(UNIX) tar -cpf domain_home_backup_030212.tar APPLICATIONS_
CONFIG/instance/domains/domain_name/*
(Windows) jar cf domain_home_backup_030212.jar APPLICATIONS_
CONFIG\instance\domains\domain_name\*
```

3. Back up the Oracle instance home. This backs up the system components, such as Oracle HTTP Server. For example:

```
(UNIX) tar -cpf instance_home_backup_030212.tar ORACLE_INSTANCE/*
(Windows) jar cf instance_home_backup_030212.jar ORACLE_INSTANCE\*
```

4. Back up the databases using the Oracle Recovery Manager (RMAN). For detailed steps, see the *Oracle Database Backup and Recovery User's Guide*.

15.7 Creating a Record of Your Oracle Fusion Applications Configuration

In the event that you need to restore and recover your Oracle Fusion Applications environment, it is important to have all the necessary information at your disposal. This is especially true in the event of a hardware loss that requires you to reconstruct all or part of your Oracle Fusion Applications environment on a new disk or host.

You should maintain an up-to-date record of your Oracle Fusion Applications environment that includes the information listed in this section. You should keep this information both in hardcopy and electronic form. The electronic form should be

stored on a host or email system that is completely separate from your Oracle Fusion Applications environment.

Your Oracle Fusion Applications hardware and software configuration record should include:

- The following information for each host in your environment:
 - Host name
 - Virtual host name (if any)
 - Domain name
 - IP address
 - Hardware platform
 - Operating system release level and patch information
- The following information for each Oracle Fusion Applications installation in your environment:
 - Installation type (for example, Oracle SOA Suite, Oracle Fusion Supply Chain Management)
 - Host on which the installation resides
 - User name, user ID number, group name, group ID number, environment profile, and type of shell for the operating system user that owns the Oracle home (*/etc/passwd* and */etc/group* entries)
 - Directory structure, mount points, and full path for the Applications base directory, Applications configuration directory, Middleware home, Oracle Common home, Oracle homes, Oracle WebLogic Server domain homes, and the Oracle instance homes
 - Amount of disk space used by the installation
 - Port numbers used by the installation
- The following information for the databases containing the metadata for components and any other databases used by Oracle Fusion Applications:
 - Host name
 - Database version and patch level
 - Base language
 - Character set
 - Global database name
 - SID
- The following information about backups:
 - The time of the backup.
 - The contents of the backup. For example, a full backup, a backup of a domain.
 - The tool used to create the backup.
 - Where the backup is stored.

15.8 Recovering Using Cloud Control

You can recover an Oracle Fusion Applications environment with Cloud Control. Cloud Control restores the directories and databases of a Oracle Fusion Applications backup. All directories and databases of the backup are restored by default. Optionally, you can select individual directories and databases to be restored.

The restore capability provided by Cloud Control is one part of a larger process that must be performed to recover a Oracle Fusion Applications environment or individual Oracle Fusion Applications components. After you restore Oracle Fusion Applications directories and databases using Cloud Control, you must perform additional manual steps depending on the component that you are restoring. Those steps are described in [Section 15.9](#) and [Section 15.10](#).

To recover Oracle Fusion Applications with Cloud Control:

1. Ensure that you have met the prerequisites, as described in [Section 15.4](#).
2. From the **Targets** menu, choose **Fusion Applications**.

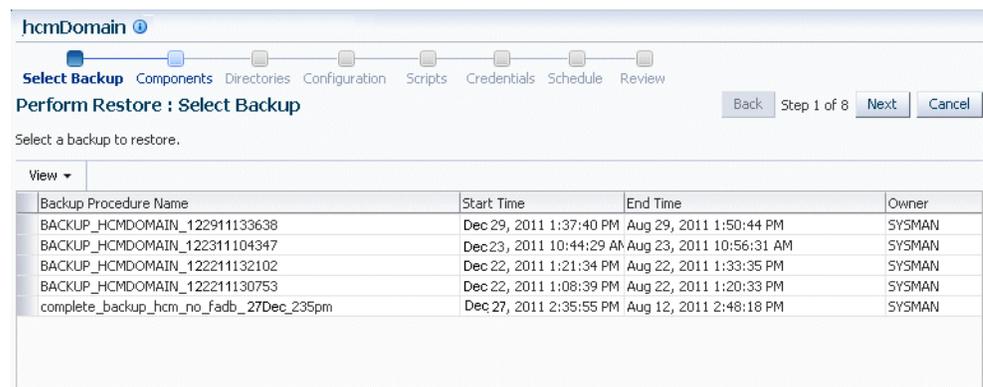
The Fusion Applications target home page is displayed.

3. Click the **Fusion Instance** that you want to recover.

The Fusion Instance home page is displayed.

4. From the **Fusion Instance** menu, choose **Control**, then **Perform Restore**.

The Perform Restore wizard is displayed, as shown in the following figure:



5. In the Select Backup step, select the backup to restore from the list of available backups for the Fusion Instance target.

Click **Next**.

6. In the Components step, the Oracle Fusion Applications components and databases included in the selected backup are listed. Select one or more components to include in the restore. The directories that correspond to the selected components will be shown on the next page. If multiple components are selected that have common directories, those common directories will be included along with any directories that are specific to individual components.

Click **Next**.

7. In the Directories step, the directories that correspond to the selected Oracle Fusion Applications components are listed in the first table. By default, they will all be restored to their original locations. You can remove directories that should not be included in the restore from the table. You can change the destination host name and path where the directories will be restored. (Note that Oracle Fusion

Applications directories must be restored to their original paths to successfully recover the environment. However, restoring to different paths is supported to provide the option of restoring to a temporary staging location in preparation for eventual movement to the final location.)

While you can remove complete directories from the list, you cannot select or remove individual subdirectories or files under the top-level directories shown in the list. If a more fine-grained restore granularity is desired, individual files and subdirectories can be restored by using the Perform Restore wizard available from the Oracle Secure Backup Domain home page. (From that home page, select **Manage**, then **File System Backup/Restore**, and then **Perform Restore**.)

In addition, any additional files or directories included in the backup are listed in a separate table, and will be restored to their original locations by default. Directories and files can be removed from the table and you can change the destination host name and path.

Click **Next**.

8. In the Configuration step, the original backup configuration used to create the backup is displayed. If the backup configuration is still available it will be used by default and no user action is required on this page. If the configuration is not available, select an alternate backup configuration that has settings similar to backup configuration used to create the backup.

Click **Next**.

9. In the Scripts step, specify pre-restore and post-restore scripts that will be executed as part of the restore operation. Scripts can be run on any discovered host. The scripts must be executable files that exist on the specified hosts.

The Oracle Fusion Applications components being restored must be shut down prior to being restored. Because this is not done automatically by Cloud Control, the appropriate shutdown scripts can be executed as part of the restore procedure by specifying them as pre-restore scripts.

Click **Next**.

10. In the Credentials step, if one or more databases is included in the restore operation, specify both database and host credentials that will be used to restore the databases.

Click **Next**.

11. In the Schedule step, specify the procedure name, description, and scheduling information.

Click **Next**.

12. In the Review step, verify all the input specified in the wizard. Make changes as needed and click **Submit**.

15.9 Recovering After Data Loss, Corruption, or Media Failure

This section describes recovery strategies for outages that involve actual data loss or corruption, or media failure where the disk cannot be restored. It also describes recovery strategies for applications that are no longer functioning properly. This type of failure requires some type of data restoration before the Oracle Fusion Applications environment can be restarted and continue with normal processing.

Depending on the extent of the failure, you can recover the Applications base directory, the Middleware homes, the Administration Server, a Managed Server, or the

database. Some Oracle Fusion Applications components require additional steps, which are described in subsequent sections.

This section includes the following topics:

- [Recovering the Applications Base Directory](#)
- [Recovering a Middleware Home](#)
- [Recovering an Oracle WebLogic Server Domain](#)
- [Recovering an Oracle Home](#)
- [Recovering an Oracle Instance Home](#)
- [Recovering the Administration Server Configuration](#)
- [Recovering a Managed Server](#)
- [Recovering the Databases](#)
- [Reconciling the Data](#)

This section also includes the following topics which describe additional considerations you may need to take for particular components:

- [Recovering Oracle Fusion Applications](#)
- [Recovering Components Related to Oracle Fusion Applications](#)

See the "Recovering Your Environment" section in the *Oracle Fusion Middleware Administrator's Guide* for information about recovering Oracle Fusion Middleware components, such as Oracle SOA Suite.

15.9.1 Recovering the Applications Base Directory

To recover an Applications base directory that was corrupted or from which files were deleted:

1. Stop all relevant processes. That is, stop all processes that are related to Oracle Fusion Applications, such as the Administration Server, Node Manager, Managed Servers, and Oracle instances, as described in [Section 3.3.2.2](#).
2. Recover the Applications base directory from the backup file. For example:


```
(UNIX) tar -xf ApplBase_backup_030212.tar
(Windows) jar xtf ApplBase_backup_030212.jar
```
3. Start all relevant processes. That is, start all processes that run in the Applications base home, as described in [Section 3.3.2.1](#).

15.9.2 Recovering a Middleware Home

To recover a Middleware home that was corrupted or from which files were deleted:

1. Stop all relevant processes. That is, stop all processes that run in the Middleware home, such as the Administration Server, Node Manager, Managed Servers, and Oracle instances, as described in [Section 3.3.2.2](#).
2. Recover the Middleware home directory from the backup file. For example:


```
cd MW_HOME
(UNIX) tar -xf mw_home_backup_030212.tar
(Windows) jar xtf mw_home_backup_030212.jar
```

3. Start all relevant processes. That is, start all processes that run in the Middleware home, as described in [Section 3.3.2.1](#).

15.9.3 Recovering an Oracle WebLogic Server Domain

To recover an Oracle WebLogic Server domain that was corrupted or deleted from the file system:

1. Stop all relevant processes. That is, stop all processes that are related to the domain, such as the Administration Server and Managed Servers, as described in [Section 3.3.2.2](#).
2. Recover the domain directory from the backup file:

```
cd DOMAIN_HOME
(UNIX) tar -xf domain_backup_030212.tar
(Windows) jar xtf domain_backup_030212.jar
```
3. Start all relevant processes. That is, start all processes that are related to the domain, as described in [Section 3.3.2.1](#).
4. If you cannot start the Administration Server, recover it, as described in [Section 15.9.6](#).
5. If you cannot start a Managed Server, recover it, as described in [Section 15.9.7](#).

15.9.4 Recovering an Oracle Home

To recover an Oracle home from the backup file:

1. Recover the Oracle home to the original directory from a backup file. For example:

```
cd ORACLE_HOME
(UNIX) tar -xf Oracle_home_backup_030212.tar
(Windows) jar xtf Oracle_home_backup_030212.jar
```
2. Restart the Managed Server to which applications are deployed, using the WLST start command. For example:

```
wls:/mydomain/serverConfig> start('myserver','Server')
```

15.9.5 Recovering an Oracle Instance Home

An Oracle instance home contains configuration information for system components, such as Oracle HTTP Server or Oracle Internet Directory.

This section includes the following topics:

- [Recovering After Oracle Instance Home Deleted from File System](#)
- [Recovering After Oracle Instance Home Deregistered](#)

15.9.5.1 Recovering After Oracle Instance Home Deleted from File System

To recover an Oracle instance home that was corrupted or deleted from the file system:

1. Stop all relevant processes. That is, kill all processes that are related to that Oracle instance.
2. Recover the Oracle instance home directory from a backup file. For example:

```
cd ORACLE_INSTANCE
(UNIX) tar -xf instance_home_backup_030212.tar
(Windows) jar xtf instance_home_backup_030212.jar
```

3. Start all relevant processes. That is, start all processes that are related to that Oracle instance:

```
opmnctl startall
```

15.9.5.2 Recovering After Oracle Instance Home Deregistered

An Oracle instance must be registered with the domain. To recover an Oracle instance home that was deregistered from the domain:

1. Recover the Oracle instance home directory from a backup file. For example:

```
cd ORACLE_INSTANCE
(UNIX) tar -xf instance_home_backup_030212.tar
(Windows) jar xtf instance_home_backup_030212.jar
```

2. Register the Oracle instance, along with all of its components, with the Administration Server, using the `opmnctl registerinstance` command. For example:

```
opmnctl registerinstance -adminHost admin_server_host
                        -adminPort admin_server_port -adminUsername username
                        -adminPassword password
                        -oracleInstance ORACLE_INSTANCE_dir -oracleHome ORACLE_HOME_dir
                        -instanceName Instance_name -wlsServerHome Middleware_Home
```

15.9.6 Recovering the Administration Server Configuration

If the Administration Server configuration has been lost because of file deletion or file system corruption, the Administration Server console continues to function if it was already started when the problem occurred. The Administration Server directory is regenerated automatically, except for security information. As a result, whenever you start the Administration Server, it prompts for a user name and password. To prevent this, you can recover the configuration.

Caution: Performing a domain-level recovery can impact other aspects of a running system and all of the configuration changes performed after the backup was taken will be lost.

To recover the Administration Server configuration:

1. Stop all processes, including the Administration Server, Managed Servers, and Node Manager, if they are started, as described in [Section 3.3.2.2](#).
2. Recover the Administration Server configuration by recovering the domain home backup to a temporary location. Then, restore the `config` directory to the following location:

```
(UNIX) DOMAIN_HOME/config
(Windows) DOMAIN_HOME\config
```

3. Start the Administration Server as described in [Task 3, "Start the Administration Servers"](#) in [Section 3.3.2.1](#).
4. Verify that the Administration Server starts properly and is accessible.
5. Start other processes, as described in [Section 3.3.2.1](#).

On the next configuration change, the configuration from the Administration Server is pushed to the Managed Servers. On each Managed Server restart, the configuration is retrieved from the Administration Server.

15.9.7 Recovering a Managed Server

For many Oracle Fusion Applications, you recover the Managed Server in which the application is deployed. In addition, for some components, you may need to take additional steps, which are described in [Section 15.9.10](#) and [Section 15.9.11](#).

In this scenario, the Managed Server does not operate properly or cannot be started because the configuration has been deleted or corrupted or the configuration was mistakenly changed and you cannot ascertain what was changed.

To recover a Managed Server:

1. If the Administration Server is not reachable, recover the Administration Server, as described in [Section 15.9.6](#).
2. If the Managed Server fails to start or if the file system is lost:
 - a. Recover the Middleware home from the backup file, if required. For example:

```
(UNIX) tar -xf mw_home_backup_030212.tar
(Windows) jar xtf mw_home_backup_030212.jar
```

- b. Create a domain template jar file for the Administration Server, using the pack utility.

For example, on UNIX:

```
pack.sh -domain=APPLICATIONS_CONFIG/instance/domains/domain_name
        -template=/scratch/temp.jar -template_name=domain1
        -template_author=myname -log=/scratch/logs/my.log -managed=true
```

For example, on Windows:

```
pack.cmd -domain=APPLICATIONS_CONFIG\instance\domains\domain_name
          -template=C:\temp\temp.jar -template_name=domain1
          -template_author=myname -log=C:\temp\logs\my.log -managed=true
```

Specifying the `-managed=true` option packs up only the Managed Servers. If you want to pack the entire domain, omit this option.

- c. Unpack the domain template JAR file, using the unpack utility.

For example, on UNIX:

```
unpack.sh -template=/scratch/temp.jar
          -domain=APPLICATIONS_CONFIG/instance/domains/domain_name
          -log=/scratch/logs/new.log -log_priority=info
```

For example, on Windows:

```
unpack.cmd -template=C:\temp\temp.jar
           -domain=APPLICATIONS_CONFIG\instance\domains\domain_name
           -log=C:\temp\logs\new.log -log_priority=info
```

- d. Ensure that the application artifacts are accessible from the Managed Server host. That is, if the application artifacts are not on the same server as the Managed Server, they must be in a location accessible by the Managed Server.

Note:

- For stage mode applications, the Administration server copies the application bits to the staged directories on the Managed Server hosts.
- For nostage and external-stage mode applications, ensure that application files are available in the stage directories of the Managed Server.

See *Oracle Fusion Middleware Deploying Applications to Oracle WebLogic Server* for information about stage, nostage, and external-stage mode applications.

- e. Start the Managed Server, as described in [Task 4, "Start the Managed Servers"](#) in [Section 3.3.2.1](#).

The Managed Server connects to the Administration Server and updates its configuration changes.

15.9.8 Recovering the Databases

You can recover the databases used by Oracle Fusion Applications using RMAN. You can recover the databases at whatever level is appropriate by performing a restore and recover of the full database, a tablespace, or a data file.

Oracle Fusion Customer Relationship Management uses the following databases: Oracle Fusion Applications database, LDAP database, Oracle WebCenter Content, and Oracle Identity Manager database. You must maintain consistency among the databases. If any of these databases are recovered to a different point in time, you may need to reconcile the databases.

For information about recovering databases, see the *Oracle Database Backup and Recovery User's Guide*.

For information about reconciling the databases, see [Section 15.9.9](#).

15.9.9 Reconciling the Data

The following topics describe how to recover and reconcile the databases to maintain consistency among the Oracle Fusion Applications database, LDAP database, and Oracle Identity Manager database.

- [Recovering the Oracle Identity Manager Database and Reconciling It with the LDAP Database](#)
- [Recovering the Oracle Fusion Applications Database and Reconciling It with the LDAP Database](#)
- [Recovering the LDAP Database Using Multimaster Replication](#)

15.9.9.1 Recovering the Oracle Identity Manager Database and Reconciling It with the LDAP Database

Oracle Identity Manager users, role categories, role hierarchies, and role memberships are stored in the Oracle Identity Manager database. When a change in the information about users or roles takes place in Oracle Identity Manager, this information is propagated to the LDAP identity store. If the change takes place in the LDAP identity store directly, these changes are synchronized into Oracle Identity Manager. The LDAP

identity store can be Oracle Internet Directory or any third party solutions such as Active Directory.

To recover the Oracle Identity Manager database, use RMAN to perform a point-in-time recovery of the Oracle Identity Manager database. See *Oracle Database Backup and Recovery User's Guide* for information on recovering a database.

If you restore the Oracle Identity Manager database to a different point in time than the LDAP store, the reconciliation engine checks the change logs and reapplies all the changes that happened in the time period between the restore of the LDAP store and the Oracle Identity Manager database. For example, if the Oracle Identity Manager database is restored so that is 10 hours behind the LDAP store, the reconciliation engine checks the change logs and reapplies all the changes that happened in the last 10 hours in the LDAP store to the Oracle Identity Manager database.

You do not need to explicitly trigger the reconciliation. LDAP synchronization is set up as a scheduled task to submit reconciliation events periodically. You can also start the reconciliation process manually and monitor the reconciliation events from the Oracle Identity Manager console. See "Reconciliation Configuration" in *Oracle Fusion Middleware User's Guide for Oracle Identity Manager*.

Note: Oracle recommends that you make sure that the Oracle Identity Manager application is unavailable to the end users when a bulk reconciliation is occurring (as in the above recovery scenario). When the bulk reconciliation is complete, make sure that the Oracle Identity Manager application is again available to the end users. You can monitor the reconciliation with the Oracle Identity Manager console.

15.9.9.2 Recovering the Oracle Fusion Applications Database and Reconciling It with the LDAP Database

If the Oracle Fusion Applications database fails, and you restore it to a different point in time than the Oracle Identity Manager database, you may lose some user and roles data. For example, if you have created some users and roles in the Oracle Identity Manager database before the Oracle Fusion Applications database is restored, those users and roles will also exist in the LDAP database, but they will not exist in the Oracle Fusion Applications database.

To reconcile the Oracle Fusion Applications database with the Oracle Identity Manager and LDAP databases:

1. Restore the Oracle Fusion Applications database using RMAN to perform a point-in-time recovery. See *Oracle Database Backup and Recovery User's Guide* for information about recovering a database.
2. Synchronize the user and role information:
 - a. Log in to the HcmCore application as the HR_SPEC_ALL user.
 - b. Click **EssLink**.
 - c. Click **Retrieve Latest LDAP Changes**.
 - d. Click **Submit**.

This is an asynchronous process that runs in the background.

3. When the process completes, find the users that were created in Oracle Identity Manager before the Oracle Fusion Applications database was restored. Use the

following SQL command, where *creation_date* is the time you started running the synchronization process in Step 2:

```
select * from per_users where creation_date > creation_date
```

4. For each user returned in the previous step, you create the user in Oracle Fusion Human Capital Management. For information, see "Creating and Updating Person and Employment Records" in the HCM Foundation help.
5. For each user, create the user's roles, as described in "User and Role Provisioning" in the HCM Foundation help.

15.9.9.3 Recovering the LDAP Database Using Multimaster Replication

Oracle recommends that you use multi-master replication for your LDAP database, as described in the "Setting up Multimaster Replication" section in the *Oracle Fusion Middleware High Availability Guide*.

When you use multimaster replication, if one LDAP node fails, the LDAP traffic is automatically routed to another node. You do not need to take any further action.

15.9.10 Recovering Oracle Fusion Applications

In most cases, to recover Oracle Fusion Applications, you recover a Middleware home, a domain, a server, an Oracle home, or an Oracle instance, depending on the extent of the failure. However, you may need to take additional steps for particular components.

This section includes the following topics:

- [Recovering Oracle Fusion Customer Relationship Management](#)
- [Recovering Oracle Fusion Procurement](#)

15.9.10.1 Recovering Oracle Fusion Customer Relationship Management

To recover Oracle Fusion Customer Relationship Management:

1. Recover the Managed Server to which Oracle Fusion Customer Relationship Management is deployed, as described in [Section 15.9.7](#).
2. Recover the Java servers for Oracle E-Mail and Web Marketing, as described in [Section 15.9.11.5](#).
3. Note that the Oracle Product Data Quality repository (the FUSION_DQ schema) and the customer source must be kept synchronized. If you restore one, restore the other to the same point in time. In addition, the Data Quality engine server artifacts in the file system, the Data Quality Admin Configuration data in the customer master (fusion), and the deployment topology must be kept synchronized.

For a complete case study for recovering Oracle Fusion Customer Relationship Management in different recovery scenarios, see [Section 15.11](#).

15.9.10.2 Recovering Oracle Fusion Procurement

To recover Oracle Fusion Procurement:

1. Recover the Managed Server to which the application is deployed, as described in [Section 15.9.7](#).

2. If you use a certificate file to make SSL connections to supplier's web sites, make sure that the certificate file exists in the location that you specified. If it does not, recover it from the backup file.

15.9.11 Recovering Components Related to Oracle Fusion Applications

You may need to recover components related to Oracle Fusion Applications.

This section includes the following topics:

- [Recovering Oracle HTTP Server](#)
- [Recovering Oracle Enterprise Scheduler](#)
- [Recovering Oracle Enterprise Crawl and Search Framework](#)
- [Recovering Oracle Authorization Policy Manager](#)
- [Recovering Oracle E-Mail and Web Marketing](#)

15.9.11.1 Recovering Oracle HTTP Server

To recover Oracle HTTP Server, you recover the Oracle instance that contains Oracle HTTP Server, as described in [Section 15.9.5](#).

15.9.11.2 Recovering Oracle Enterprise Scheduler

To recover Oracle Enterprise Scheduler:

1. Recover the domain directory from the backup file, as described in [Section 15.9.3](#).
2. Recover the Oracle home to the original directory from the backup file, as described in [Section 15.9.4](#).
3. Recover the database containing the Oracle Fusion applications and MDS schemas to the most recent point in time, if needed, as described in [Section 15.9.8](#).

15.9.11.3 Recovering Oracle Enterprise Crawl and Search Framework

To recover Oracle Enterprise Crawl and Search Framework:

1. Recover the domain directory from the backup file, as described in [Section 15.9.3](#).
2. Recover the database containing schemas related to Oracle Enterprise Crawl and Search Framework, as described in [Section 15.9.8](#).

15.9.11.4 Recovering Oracle Authorization Policy Manager

To recover Oracle Authorization Policy Manager:

1. Stop all relevant processes. That is, stop all processes that are related to the domain, such as the Administration Server and Managed Servers, as described in [Section 3.3.2.2](#).
2. Recover the domain directory from the backup file, as described in [Section 15.9.3](#).
3. Recover the Oracle home to the original directory from the backup file, as described in [Section 15.9.4](#).
4. Recover the LDAP store, if necessary.

See the "Recovering Components" section in the *Oracle Fusion Middleware Administrator's Guide*.

5. Recover the database, if necessary. See [Section 15.9.8](#).

6. Start all relevant processes. That is, start all processes that are related to the domain, as described in [Section 3.3.2.1](#).

15.9.11.5 Recovering Oracle E-Mail and Web Marketing

To recover Oracle E-Mail and Web Marketing:

1. Recover the Managed Servers to which the Email Sending Daemon and Click Thru Daemon are deployed, as described in [Section 15.9.7](#).
2. Recover the installation directory for the Bounce Handling Daemon.

15.10 Recovering After Loss of Host

This section describes recovery strategies after losing the original operating environment. For example, you could have a serious system malfunction or loss of media.

Depending on the extent of the failure, you can recover the Applications base directory, the Administration Server, a Managed Server, or the database. Some Oracle Fusion Applications components require additional steps, which are described in subsequent sections.

This section includes the following topics:

- [Recovering the Applications Base Directory After Loss of Host](#)
- [Recovering After Loss of Administration Server Host](#)
- [Recovering After Loss of Managed Server Host](#)
- [Recovering the Databases After Loss of Host](#)
- [Additional Actions for Recovering Entities After Loss of Host](#)

This section also includes the following topics, which describe additional considerations you may need to take for Oracle Fusion Applications components:

- [Recovering Oracle Fusion Applications After Loss of Host](#)
- [Recovering Components Related to Oracle Fusion Applications](#)

15.10.1 Recovering the Applications Base Directory After Loss of Host

You can recover the Applications base directory if you lose the host that contains the directory.

To recover the Applications base directory:

1. Recover the Applications base directory from the backup file. For example:

```
cd APPLICATIONS_BASE
(UNIX) tar -xf ApplBase_backup_030212.tar
(Windows) jar xtf ApplBase_backup_030212.jar
```

2. Start all relevant processes. That is, start all processes that run in the Applications base home, as described in [Section 3.3.2.1](#).

15.10.2 Recovering After Loss of Administration Server Host

If you lose a host that contains the Administration Server, you can recover it to the same host or a different host.

This section includes the following topics:

- [Recovering the Administration Server to the Same Host](#)
- [Recovering the Administration Server to a Different Host](#)

15.10.2.1 Recovering the Administration Server to the Same Host

In this scenario, you recover the Administration Server either to the same host after the operating system has been reinstalled or to a new host that has the same host name. For example, the Administration Server is running on Host A and the Managed Server is running on Host B. Host A has failed for some reason and the Administration Server needs to be recovered.

To recover the Administration Server:

1. Recover the file system. For example, recover the domain containing the Administration Server, as described in [Section 15.9.3](#).
2. Attempt to start the Administration Server, as described in [Task 3, "Start the Administration Servers"](#) in [Section 3.3.2.1](#).

If the Administration Server starts, you do not need to take any further steps.

3. If the Administration Server fails to start, take the following steps on Host A:
 - a. Stop all relevant processes. That is, stop all processes that are related to the domain, such as the Administration Server and Managed Servers, as described in [Section 3.3.2.2](#).

- b. Recover the domain directory from the backup file:

```
cd DOMAIN_HOME
(UNIX) tar -xf domain_backup_030212.tar
(Windows) jar xtf domain_backup_030212.jar
```

This restores the Administration Server as well as the Managed Servers in the domain.

- c. Start the Administration Server, as described in [Task 3, "Start the Administration Servers"](#) in [Section 3.3.2.1](#).
- d. Start the Managed Servers, as described in [Task 4, "Start the Managed Servers"](#) in [Section 3.3.2.1](#).
- e. Start the Node Manager:

```
java weblogic.WLST
wls:/offline> startNodeManager()
```

15.10.2.2 Recovering the Administration Server to a Different Host

In this scenario, the Administration Server is running on Host A and the Managed Server is running on Host B. Host A has failed for some reason and the Administration Server needs to be moved to Host C.

Note: Note that this scenario assumes that the shared location containing the binary files and the Administration Server configuration is intact. If it is not, follow the steps in [Section 15.9.2](#) through [Section 15.9.5](#).

To recover the Administration Server to a different host:

1. Because you have included a new host, mount the file system to the new host, Host C.
2. If the Administration Server has a Listen address, create a new machine with the new host name, as described in [Section 15.10.5.3](#).
3. When you move .jks files to another host, you may receive warnings from the host name verification. You can regenerate those SSL certificates in the new host to set up SSL and configure a custom host name verifier. See the "Configure a Custom Host Name Verifier" section in the Oracle WebLogic Server Administration Console Help, which is located at:

http://download.oracle.com/docs/cd/E21764_01/apirefs.1111/e13952/taskhelp/security/ConfigureACustomHostNameVerifier.html

Also, see the "Configuring SSL" chapter in *Oracle Fusion Middleware Securing Oracle WebLogic Server*.

The .jks files are located in the following directory:

```
(UNIX) APPLICATIONS_BASE/fusionapps/wlserver_10.3/server/lib
(Windows) APPLICATIONS_BASE\fusionapps\wlserver_10.3\server\lib
```

4. Start the Node Manager on Host C if it was configured on the original host:


```
java weblogic.WLST
wls:/offline> startNodeManager()
```
5. Start the Administration Server, as described in [Task 3, "Start the Administration Servers"](#) in [Section 3.3.2.1](#).
6. If the Managed Servers are on the same failed host as the Administration Server, restore the Managed Servers, as described in [Section 15.10.3.2](#).
7. Start the Managed Servers. The "Restarting a Failed Administration Server" section in the *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server* describes different ways to restart them, depending on how they were configured.

One option is described in [Task 4, "Start the Managed Servers"](#) in [Section 3.3.2.1](#).
8. Ensure that additional application artifacts are available. For example, if the deployment mode is nostage or external-stage, applications may reside in directories outside of the domain directory. Make your application files available to the new Administration Server by copying them from backups or by using a shared disk. Your application files should be available in the same relative location on the new file system as on the file system of the original Administration Server.

If the application is staged, the Administration Server copies the application bits to the staged directories on the Managed Server hosts.
9. If your environment contains Oracle HTTP Server, modify the FusionVirtualHost_x.conf file, as described in [Section 15.10.5.2](#).
10. Update Oracle Inventory, as described in [Section 15.10.5.4](#).
11. Edit the targets.xml file for Fusion Middleware Control, as described in [Section 15.10.5.1](#).

15.10.3 Recovering After Loss of Managed Server Host

If you lose a host that contains a Managed Server, you can recover it to the same host or a different host.

This section includes the following topics:

- [Recovering a Managed Server to the Same Host](#)
- [Recovering a Managed Server to a Different Host](#)

15.10.3.1 Recovering a Managed Server to the Same Host

In this scenario, you recover a Managed Server to the same host after the operating system has been reinstalled or to a new host that has the same host name. The Administration Server is running on Host A and the Managed Server is running on Host B. Host B failed for some reason and the Managed Server needs to be recovered to Host B.

1. Start the Node Manager on Host B:

```
java weblogic.WLST
wls:/offline> startNodeManager()
```

2. Start the Managed Server, as described in [Task 4, "Start the Managed Servers"](#) in [Section 3.3.2.1](#).

If the Managed Server starts, it connects to the Administration Server and updates its configuration changes. You do not need to take any further steps.

3. If the Managed Server fails to start or if the file system is lost:

- a. Stop the Node Manager:

```
java weblogic.WLST
wls:/offline> stopNodeManager()
```

- b. Recover the Middleware home to Host B from the backup file, if required:

```
cd MW_HOME
(UNIX) tar -xf mw_home_backup_030212.tar
(Windows) jar xtf mw_home_backup_030212.jar
```

- c. Create a domain template jar file for the Administration Server running in Host A, using the pack utility.

For example, on UNIX:

```
pack.sh -domain=APPLICATIONS_CONFIG/instance/domains/domain_name
-template=/scratch/temp.jar -template_name=domain1
-template_author=myname -log=/scratch/logs/my.log -managed=true
```

For example, on Windows:

```
pack.cmd -domain=APPLICATIONS_CONFIG\instance\domains\domain_name
-template=C:\temp\temp.jar -template_name=domain1
-template_author=myname -log=C:\temp\logs\my.log -managed=true
```

Specifying the `-managed=true` option packs up only the Managed Servers. If you want to pack the entire domain, omit this option.

- d. Unpack the domain template jar file in Host B, using the unpack utility.

For example, on UNIX:

```
unpack.sh -template=/scratch/temp.jar
```

```
-domain=APPLICATIONS_CONFIG/instance/domains/domain_name
-log=/scratch/logs/new.log -log_priority=info
```

For example, on Windows:

```
unpack.cmd -template=C:\temp\temp.jar
-domain=APPLICATIONS_CONFIG\instance\domains\domain_name
-log=C:\temp\logs\new.log -log_priority=info
```

- e. Ensure that the application artifacts are accessible from the Managed Server host. That is, if the application artifacts are not on the same server as the Managed Server, they must be in a location accessible by the Managed Server.

Note:

- For applications that are deployed in nostage or external-stage mode, copy the application artifacts from the Administration Server host directory.
- For applications that are deployed in stage mode, the Administration server copies the application bits to the staged directories on the Managed Server hosts.

See *Oracle Fusion Middleware Deploying Applications to Oracle WebLogic Server* for information about deploying applications.

- f. If the Node Manager is not started, start it:

```
java weblogic.WLST
wls:/offline> startNodeManager()
```

- g. Start the Managed Server, as described in [Task 4, "Start the Managed Servers"](#) in [Section 3.3.2.1](#).

The Managed Server connects to the Administration Server and updates its configuration changes.

15.10.3.2 Recovering a Managed Server to a Different Host

In this scenario, the Administration Server is running on Host A and the Managed Server is running on Host B. Host B failed for some reason and the Managed Server needs to be recovered to Host C.

Important: Recover the Middleware home to the same location as the original.

To recover a Managed Server to a different host:

1. If the Managed Server configuration files are local to the host (that is, they are not on a shared file system):
 - a. Recover the Middleware home for the Managed Server to Host C.

```
cd MW_HOME
(UNIX) tar -xf mw_home_backup_030212.tar
(Windows) jar xtf mw_home_backup_030212.jar
```

Note that when you restore the Middleware home, you are restoring all of the domains because they are on a shared file system.

- b. Create a domain template jar file from the Administration Server running in Host A, using the pack utility.

For example, on UNIX:

```
pack.sh -domain=APPLICATIONS_CONFIG/instance/domains/domain_name
        -template=/scratch/temp.jar -template_name=domain1
        -template_author=myname -log=/scratch/logs/my.log -managed=true
```

For example, on Windows:

```
pack.cmd -domain=APPLICATIONS_CONFIG\instance\domains\domain_name
         -template=C:\temp\temp.jar -template_name=domain1
         -template_author=myname -log=C:\temp\logs\my.log -managed=true
```

Specifying the `-managed=true` option packs up only the Managed Servers. If you want to pack the entire domain, omit this option.

- c. Unpack the domain template jar file on Host C, using the unpack utility.

For example, on UNIX:

```
unpack.sh -template=/scratch/temp.jar
         -domain=APPLICATIONS_CONFIG/instance/domains/domain_name
         -log=/scratch/logs/new.log -log_priority=info
```

For example, on Windows:

```
unpack.cmd -template=C:\temp\temp.jar
          -domain=APPLICATIONS_CONFIG\instance\domains\domain_name
          -log=C:\temp\logs\new.log -log_priority=info
```

If you are recovering to a different domain home, use the `-app_dir` switch in the `unpack` command.

2. Ensure that the application artifacts are accessible from the Managed Server host. That is, if the application artifacts are not on the same server as the Managed Server, they must be in a location accessible by the Managed Server.

Note:

- For applications that are deployed in nostage or external-stage mode, copy the application artifacts from the Administration Server host directory.
- For applications that are deployed in stage mode, the Administration server copies the application bits to the staged directories on the Managed Server hosts.

See *Oracle Fusion Middleware Deploying Applications to Oracle WebLogic Server* for information about deploying applications.

3. If the Managed Server is *not* co-located with the Administration Server, take the following steps. (These steps are not needed if the Managed Server is co-located with the Administration Server.)
 - a. Edit the `nodemanager.domains` file, specifying the domain names and domain directories. Use the following format:


```
domain_name=domain_directory
```
 - b. Start the Node Manager on Host C, if it is not started:

```
java weblogic.WLST
wls:/offline> startNodeManager()
```

- c. Using WLST, connect to the Administration Server and then enroll the Node Manager running in the new host with the Administration Server:

```
connect('username', 'password', 'host:port')
nmEnroll('APPLICATIONS_CONFIG/instance/domains/domain_name',
        'MW_HOME/wlserver_n/common/nodemanager/instance_name')
```

- d. Change the Managed Server configuration to point to the new host:

- In the WebLogic Server Administration Console, create a machine, which is a logical representation of the computer that hosts one or more WebLogic Servers, and point it to the new host. (From the Home page, select **Machines**. Then, click **New**.) Follow the directions in the Administration Console help.

If you identify the Listen Address by IP address, you must disable Host Name Verification on the Administration Servers that access Node Manager. For more information, see the "Using Host Name Verification" section in *Oracle Fusion Middleware Securing Oracle WebLogic Server*.

- Change the Managed Server configuration to point to the new machine. (From the left pane of the Console, expand **Environment** and then **Servers**. Then, select the name of the server. Select the **Configuration** tab, then the **General** tab. In the **Machine** field, select the machine to which you want to assign the server.)

Change **Listen Address** to the new host. (If the listening address was set to blank, you do not need to change it.)

You only need to take these steps once for all Managed Servers on the same host.

4. When you move .jks files to another host, you may receive warnings from the host name verification. You can regenerate those SSL certificates in the new host to set up SSL and configure a custom host name verifier. See the "Configure a Custom Host Name Verifier" section in the Oracle WebLogic Server Administration Console Help. Also, see the "Configuring SSL" chapter in *Oracle Fusion Middleware Securing Oracle WebLogic Server*.

The .jks files are located in the following directory:

```
(UNIX) APPLICATIONS_BASE/fusionapps/wlserver_10.3/server/lib
(Windows) APPLICATIONS_BASE\fusionapps\wlserver_10.3\server\lib
```

5. Start the Managed Server, as described in [Task 4, "Start the Managed Servers"](#) in [Section 3.3.2.1](#).

The Managed Server connects to the Administration Server and updates its configuration changes.

6. If your environment contains Oracle HTTP Server, modify the FusionVirtualHost_x.conf file, as described in [Section 15.10.5.2](#).
7. Edit the targets.xml file for Fusion Middleware Control, as described in [Section 15.10.5.1](#).

Now you can start and stop the Managed Server on Host C using the Administration Server running on Host A.

15.10.4 Recovering the Databases After Loss of Host

If the physical host where your database resides is lost, you can recover the database using RMAN.

See "Restoring a Database on a New Host" in the *Oracle Database Backup and Recovery User's Guide* to learn how to use RMAN to recover your database in the event of a complete failure on the primary database host computer.

Oracle Fusion Customer Relationship Management uses the following databases: Oracle Fusion Applications database, LDAP database, Oracle WebCenter Content, and Oracle Identity Manager database. You must maintain consistency among the databases. If any of these databases are recovered to a different point in time, you may need to reconcile the databases.

For information about reconciling the LDAP database, the Oracle Fusion Applications database, and the Oracle Identity Manager database, see [Section 15.9.9](#).

15.10.5 Additional Actions for Recovering Entities After Loss of Host

Depending on the entity that you are recovering, you may need to take additional actions after loss of host. The sections about each entity may require you to follow one or more of the following procedures. If so, that is noted in the section describing how to recover the entity.

This section includes the following topics:

- [Changing the Host Name in the targets.xml File for Fusion Middleware Control](#)
- [Modifying the FusionVirtualHost_x.conf File](#)
- [Creating a New Machine for the New Host Name](#)
- [Updating Oracle Inventory](#)

15.10.5.1 Changing the Host Name in the targets.xml File for Fusion Middleware Control

When you recover a component to a different host, you must update the targets.xml file for Fusion Middleware Control. The file is located at:

```
(UNIX) APPLICATIONS_CONFIG/instance/domains/hostname/domain_name/sysman/state/targets.xml
(Windows) APPLICATIONS_CONFIG\instance\domains\hostname\domain_name\sysman\state\targets.xml
```

In the file, change the host name to the new host name for components that are recovered to a different host.

15.10.5.2 Modifying the FusionVirtualHost_x.conf File

When you recover an Administration Server or a Managed Server to a different host and your environment includes Oracle HTTP Server, you must modify the FusionVirtualHost_x.conf file on the new host. There is a separate file for each domain, for example FusionVirtualHost_fin.conf. The files are located in:

```
(UNIX) APPLICATIONS_CONFIG/instance/CommonDomain_webtier/config/OHS/ohs_name/moduleconf/FusionVirtualHost_x.conf
(Windows) APPLICATIONS_CONFIG\instance\CommonDomain_webtier\config\OHS\ohs_name\moduleconf\FusionVirtualHost_x.conf
```

Modify all of the instances of the host name and clusters (elements such as WebLogicHost and WebLogicCluster) entries in that file. For example:

```
<Location /console>
```

```

    SetHandler weblogic-handler
    WebLogicHost Admin_Host
    WebLogicPort Admin_Port
    WLProxySSL ON
    WLProxySSLPassThrough ON
</Location>
.
.
.
<Location /soa-infra>
    SetHandler weblogic-handler
    WebLogicCluster SOAHOST1:8001,*SOAHOST2*:8001*
    WLProxySSL ON
    WLProxySSLPassThrough ON
</Location>

```

15.10.5.3 Creating a New Machine for the New Host Name

If the Administration Server has an Listen address, you must create a new machine with the new host name and set the Listen address, before you start the Administration Server. A machine is a logical representation of the computer that hosts one or more WebLogic Servers.

To create a new machine:

1. Create a new machine with the new host name. Use the following WLST commands, in offline mode:

```

readDomain('Domain_Home')
machine = create('newhostname', 'Machine')
cd('/Machine/newhostname')
nm = create('newhostname', 'NodeManager')

cd('/Machine/newhostname/NodeManager/newhostname')
set('ListenAddress', 'newhostname')
updateDomain()

```

2. For the Administration Server, set the machine with the new host name, using the following WLST command, in offline mode:

```

readDomain('DomainHome')
cd('/Machine/newhostname')
machine = cmo
cd('/Server/AdminServer')
set('Machine', machine)
updateDomain()

```

3. Set the Listen address for the Administration Server:

```

readDomain("DomainHome")
cd("servers/AdminServer")
cmo.setListenPort(8001)
updateDomain()
exit()

```

15.10.5.4 Updating Oracle Inventory

For many components, when you recover to a different host, as in the case of loss of host, you must update the Oracle Inventory on UNIX. To do so, execute the following script:

```
ORACLE_COMMON_HOME/oui/bin/attachHome.sh
```

In addition, you must update `beahomelist` to edit the location of a Middleware home. Edit the following file to update the Middleware home information:

```
(UNIX) user_home/boa/beahomelist  
(Windows) C:\boa\beahomelist
```

15.10.6 Recovering Oracle Fusion Applications After Loss of Host

In most cases, to recover Oracle Fusion Applications, you recover the entire Middleware home, a domain, a server, an Oracle home, or an Oracle instance, depending on the extent of the failure. However, you may need to take additional steps you for particular components.

This section includes the following topics:

- [Recovering Oracle Fusion Supply Chain Management After Loss of Host](#)
- [Recovering Oracle Fusion Procurement After Loss of Host](#)

15.10.6.1 Recovering Oracle Fusion Supply Chain Management After Loss of Host

If you lose a host that contains Oracle Fusion Supply Chain Management, you can recover it to the same host or a different host.

To recover Oracle Fusion Supply Chain Management to the same host, recover the Managed Server to which the application is deployed, as described in [Section 15.10.3](#).

To recover Oracle Fusion Supply Chain Management to a different host:

1. Recover the Managed Server to which the application is deployed, as described in [Section 15.10.3](#).
2. Recover the Oracle instance for Global Order Promising, as described in [Section 15.9.5](#).

15.10.6.2 Recovering Oracle Fusion Procurement After Loss of Host

If you lose a host that contains Oracle Fusion Procurement, you can recover it to the same host or a different host.

To recover Oracle Fusion Procurement to the same host, recover the Managed Server to which the application is deployed, as described in [Section 15.10.3](#).

To recover Oracle Fusion Procurement to a different host:

1. Recover the Managed Server to which the application is deployed, as described in [Section 15.10.3](#).
2. Ensure that Oracle Business Intelligence Publisher and CUPS IP address and port number reflects the different host.

15.10.7 Recovering Components Related to Oracle Fusion Applications

In most cases, to recover components related to Oracle Fusion Applications, you recover the entire Middleware home, a domain, a server, an Oracle home, or an Oracle instance, depending on the extent of the failure. You may need to take additional steps for particular components related to Oracle Fusion Applications.

This section includes the following topics:

- [Recovering Oracle HTTP Server After Loss of Host](#)
- [Recovering Oracle Enterprise Scheduler After Loss of Host](#)

- [Recovering Oracle Enterprise Crawl and Search Framework After Loss of Host](#)
- [Recovering Oracle Authorization Policy Manager After Loss of Host](#)
- [Recovering Oracle E-Mail and Web Marketing After Loss of Host](#)
- [Recovering Oracle Essbase In Clustered Environment After Loss of Host](#)

15.10.7.1 Recovering Oracle HTTP Server After Loss of Host

To recover Oracle HTTP Server to the same host, recover the Oracle instance, as described in [Section 15.9.5](#).

To recover Oracle HTTP Server to a different host:

1. Recover the Middleware home, as described in [Section 15.9.2](#).
2. Start all relevant processes, as described in [Section 3.3.1](#).
3. Update the registration of the Oracle instance with the Administration Server, using the `opmnctl updateinstanceregistration` command on the new host. For example:

```
opmnctl updateinstanceregistration -adminHost admin_server_host
```

This command updates the OPMN `instance.properties` file.

4. Update the registration of the component with the Administration Server, using the `opmnctl updatecomponentregistration` command on the new host. For example, to update the registration for Oracle HTTP Server, use the following command:

```
opmnctl updatecomponentregistration -Host new_host -Port nonSSLPort
  -componentName ohs1 -componentType OHS
```

5. Edit the `targets.xml` file for Fusion Middleware Control, as described in [Section 15.10.5.1](#).
6. Modify the `ServerName` entry in the following file to have the new host name:

```
(UNIX) ORACLE_INSTANCE/config/OHS/ohs_name/httpd.conf
(Windows) ORACLE_INSTANCE\config\OHS\ohs_name\httpd.conf
```

7. Modify the `FusionVirtualHost_x.conf` file, as described in [Section 15.10.5.2](#).
8. If the Managed Server is in a cluster, modify the following files (there is more than one file, with the name `ear_name.ear`):

```
(UNIX) APPLICATIONS_CONFIG/fusionapps/applications/domain_name/deploy/ear_
name.ear/APP-INF/classes/wf_client_config.xml
(Windows) APPLICATIONS_CONFIG\fusionapps\applications\domain_name\deploy\ear_
name.ear\APP-INF\classes\wf_client_config.xml
```

Change the `rootEndPointURL` element so that it points to the Oracle HTTP Server or Load Balancer for the cluster on the new environment.

15.10.7.2 Recovering Oracle Enterprise Scheduler After Loss of Host

To recover Oracle Enterprise Scheduler, follow the procedure in [Section 15.9.11.2](#).

15.10.7.3 Recovering Oracle Enterprise Crawl and Search Framework After Loss of Host

To recover Oracle Enterprise Crawl and Search Framework to a different host, follow the procedure in [Section 15.9.11.3](#).

15.10.7.4 Recovering Oracle Authorization Policy Manager After Loss of Host

To recover Oracle Authorization Policy Manager, follow the procedure in [Section 15.9.11.4](#).

15.10.7.5 Recovering Oracle E-Mail and Web Marketing After Loss of Host

To recover Oracle E-Mail and Web Marketing, follow the procedure in [Section 15.9.11.5](#).

15.10.7.6 Recovering Oracle Essbase In Clustered Environment After Loss of Host

To recover Oracle Essbase after loss of host, you recover the entire Middleware home, a domain, a server, an Oracle home, or an Oracle instance, depending on the extent of the failure.

However, if the failed node contained Oracle Essbase and was clustered, you must create the Oracle Essbase instance and run the `essfoconfig.sh` script to set up the EPM registry. For more information, see "Configuring Essbase Clustering Using the Essbase Failover Automation Tool" in the *Oracle Fusion Applications Customer Relationship Management Enterprise Deployment Guide*.

15.11 A Case Study: Recovering Oracle Fusion Customer Relationship Management

This section provides a case study of recovering Oracle Fusion Customer Relationship Management and all of the other Oracle Fusion Applications offerings and Oracle Fusion Middleware components that are installed when you install Oracle Fusion Customer Relationship Management.

You can recover all or part of your Oracle Fusion Applications environment, as described in [Section 15.3](#).

This section includes the following topics:

- [The Recovery Case Study Scenario](#)
- [Recovering the Middleware Home Containing Java EE Components](#)
- [Recovering the Web Tier Middleware Home](#)
- [Recovering the Oracle Fusion Customer Relationship Management Domain](#)
- [Recovering Servers When the Installation Directory Is Shared Between Hosts](#)
- [Recovering Servers When the Managed Server Configuration Is Local](#)
- [Recovering an Oracle Instance in an Oracle Fusion Customer Relationship Management Installation](#)
- [Recovering an Oracle Fusion Customer Relationship Management Cluster](#)
- [Recovering Databases for Oracle Fusion Customer Relationship Management](#)

15.11.1 The Recovery Case Study Scenario

In this case study, the following domains were created when you installed and provisioned Oracle Fusion Customer Relationship Management:

Oracle Fusion Applications Component	Domain Name
Oracle Fusion Customer Relationship Management	CRMDomain

Oracle Fusion Applications Component	Domain Name
Oracle Fusion Financials	FinancialDomain
Oracle Fusion Human Capital Management	HCMDomain
Oracle Fusion Supply Chain Management	SCMDomain
Oracle Fusion Setup	CommonDomain

In addition, Oracle Fusion Customer Relationship Management uses the following Oracle Fusion Middleware components:

- Oracle WebCenter Portal
- Oracle Business Intelligence
- Oracle Essbase
- Oracle Enterprise Scheduler
- Oracle SOA Suite
- Oracle WebCenter Content
- Oracle Application Development Framework

Oracle Fusion Customer Relationship Management is installed in the following directory:

```
/scratch/oracle/APPLICATIONS_BASE
```

The *APPLICATIONS_BASE* directory contains:

- The Oracle Fusion Applications Middleware home directory, which is located at:
/scratch/oracle/APPLICATIONS_BASE/fusionapps
- One Middleware home that hosts the Java EE components, including Oracle SOA Suite and Oracle WebCenter Portal, and is located at:
/scratch/oracle/APPLICATIONS_BASE/fusionapps
- A second Middleware home that hosts system components such as Oracle HTTP Server and is located at:
/scratch/oracle/APPLICATIONS_BASE/webtier_mw_home

Note: In this case study, the Oracle Fusion application files are installed in shared storage, which is mounted on all hosts that run the applications. A single copy of binary files is shared across multiple hosts. The Administration Server configuration is also shared. However, the Managed Server configurations are not shared.

15.11.2 Recovering the Middleware Home Containing Java EE Components

In this scenario, you recover the Oracle Fusion Applications Middleware home, which contains Oracle Fusion Applications instances, and the Java EE components, Oracle SOA Suite and Oracle WebCenter Portal, after data loss or corruption.

To recover the Middleware home:

1. Stop all relevant processes using Fusion Applications Control. That is, stop all Managed Servers that are related to Oracle SOA Suite and Oracle WebCenter

Portal. For example, stop the FINSOAServer_1 (in the FinancialDomain), CRMSOAServer_1 (in the CRMDomain), SCMSOAServer_1 (in the SCMDomain), and FSSOAServer (in the HCMDomain). Stop the Administration Server, as described in [Section 3.3.2.2](#).

2. From the original Middleware home directory, move the files to a backup directory:

```
mv /scratch/oracle/APPLICATIONS_BASE/fusionapps
   /scratch/oracle/APPLICATIONS_BASE/fusionapps_backup
```

This ensures that any needed files are not overwritten.

3. Create a staging directory and restore the Applications base directory to that directory:

```
cd /scratch/stage/
tar -xvf crm_ApplBase_backup_030212.tar
```

4. From that stage directory, copy the Middleware home to the original location:

```
cp /scratch/stage/APPLICATIONS_BASE/fusionapps
   /scratch/oracle/APPLICATIONS_BASE/fusionapps
```

5. Start all relevant processes, as described in [Section 3.3.2.1](#).

15.11.3 Recovering the Web Tier Middleware Home

In this scenario, you recover the Web tier Middleware home that contains system components, such as Oracle HTTP Server, after data loss or corruption.

To recover the Middleware home:

1. Stop all relevant processes, such as all processes running in the Oracle instance, using the following command:

```
opmnctl stopall
```

2. From the original Middleware home directory, move the files to a backup directory:

```
mv /scratch/oracle/APPLICATIONS_BASE/webtier_mw_home
   /scratch/oracle/APPLICATIONS_BASE/webtier_mw_home_backup
```

This ensures that any needed files are not overwritten.

3. Create a staging directory and restore the Applications base directory to that directory:

```
cd /scratch/stage
tar -xvf crm_ApplBase_backup_030212.tar
```

4. From that stage directory, copy the Middleware home to the original location:

```
cp /scratch/stage/APPLICATIONS_BASE/webtier_mwhome
   /scratch/oracle/APPLICATIONS_BASE/webtier_mwhome
```

5. Start all relevant processes, using the following command:

```
opmnctl startall
```

15.11.4 Recovering the Oracle Fusion Customer Relationship Management Domain

In this scenario, you recover the domain that contains Oracle Fusion Customer Relationship Management after data loss or corruption. This procedure shows recovering the domain CRMDomain, but you can use the same procedure for other Oracle Fusion Applications domains.

To recover the domain:

1. Stop all Managed Servers in the CRMDomain, as described in [Task 5, "Stop the Managed Servers"](#) in [Section 3.3.2.2](#).
2. Stop the Administration Server for the CRMDomain, as described in [Task 4, "Stop the Administration Servers"](#) in [Section 3.3.2.2](#).

3. From the original domain directory, move the files to a backup directory:

```
mv /scratch/oracle/APPLICATIONS_CONFIG/instance/domains/CRMDomain
   /scratch/oracle/APPLICATIONS_CONFIG/instance/domains/CRMDomain_backup
```

This ensures that any needed files are not overwritten.

4. Create a staging directory and restore the Applications base directory to that directory:

```
cd /scratch/stage
tar -xvf crm_ApplBase_backup_030212.tar
```

5. From that stage directory, copy the domain to the original location:

```
cp /scratch/stage/APPLICATIONS_CONFIG/instance/domains/CRMDomain
   /scratch/oracle/APPLICATIONS_CONFIG/instance/domains/CRMDomain
```

6. Start the Administration Server, as described in [Task 3, "Start the Administration Servers"](#) in [Section 3.3.2.1](#).
7. Start all Managed Servers, as described in [Task 4, "Start the Managed Servers"](#) in [Section 3.3.2.1](#).

15.11.5 Recovering Servers When the Installation Directory Is Shared Between Hosts

When the applications directory is shared between hosts, and you need to recover an Administration Server or a Managed Server, you must recover the entire domain.

The following table shows a sample domain mapping and provides guidelines for recovering an Oracle WebLogic Server Administration Server or any Managed Server running in the domain. The list of servers in the table is not complete. Refer to the particular Oracle Fusion application offering for the complete list of servers.

Domain Name	Servers Running in Domain	What to Recover
CRMDomain	AdminServer CRMCommonServer_1 CRMPPerformanceServer_1 CRMANalyticsServer_1 CRMSearchServer_1 CRMODIServer_1 CRMSOAServer_1 CustomerServer_1 SalesServer_1 MarketingServer_1 OrderCaptureServer_1	The entire CRMDomain. For example: <i>APPLICATIONS_</i> <i>CONFIG/instance/domains/CRMDomain</i>
FinancialDomain	AdminServer FinCommonServer_1 FinAnalyticsServer_1 FinSOAServer_1 LedgerServer_1 ESSServer_1	The entire FinancialDomain. For example: <i>APPLICATIONS_</i> <i>CONFIG/instance/domains/FinancialDomain</i>
SCMDomain	AdminServer ProductManagementServer_1 SCMSOAServer_1	The entire SCMDomain. For example: <i>APPLICATIONS_</i> <i>CONFIG/instance/domains/SCMDomain</i>
HCMDomain	AdminServer CoreSetupServer_1 FINSOAServer_1 TalentServer_1 ESSServer_1 HCMAnalyticsServer_1	The entire HCMDomain. For example: <i>APPLICATIONS_</i> <i>CONFIG/instance/domains/HCMDomain</i>
SetupDomain	AdminServer HomePageServer_1 HelpPortalServer_1 FunctionalSetupServer_1 WCServices_1 SpacesServer_1	The entire SetupDomain. For example: <i>APPLICATIONS_</i> <i>CONFIG/instance/domains/SETupDomain</i>

15.11.6 Recovering Servers When the Managed Server Configuration Is Local

If domain directory of the Managed Server is different than the Administration Server domain directory, and you need to recover a Managed Server, you use `pack` and `unpack` to recover the Managed Server.

For example, to pack and unpack the Managed Server `FinCommonServer_1`:

1. If the Administration Server is not reachable, recover the Administration Server, as described in [Section 15.9.6](#).
2. If the Managed Server fails to start or if the file system is lost:
 - a. Recover the Middleware home from the backup file, if required. For example:

```
cd MW_HOME
tar -xf mw_home_backup_030212.tar
```

- b. Create a domain template jar file for the FinancialDomain Administration Server, using the pack utility. For example:

```
pack.sh -domain=APPLICATIONS_CONFIG/instance/domains/FinancialDomain
-template=/scratch/temp.jar -template_name=domain1
-template_author=myname -log=/scratch/logs/my.log -managed=true
```

Specifying the `-managed=true` option packs up only the Managed Servers. If you want to pack the entire domain, omit this option.

- c. Unpack the domain template JAR file, using the unpack utility. For example:

```
unpack.sh -template=/scratch/temp.jar
-domain=APPLICATIONS_CONFIG/instance/domains/FinancialDomain
-log=/scratch/logs/my.lognew.log -log_priority=info
```

- d. Ensure that the application artifacts are accessible from the Managed Server host. That is, if the application artifacts are not on the same server as the Managed Server, they must be in a location accessible by the Managed Server.

Note:

- For stage mode applications, the Administration server copies the application bits to the staged directories on the Managed Server hosts.
- For nostage and external-stage mode applications, ensure that application files are available in the stage directories of the Managed Server.

See *Oracle Fusion Middleware Deploying Applications to Oracle WebLogic Server* for information about stage, nostage, and external-stage mode applications.

- e. Start the Managed Server, as described in [Task 4, "Start the Managed Servers"](#) in [Section 3.3.2.1](#).

The Managed Server connects to the Administration Server and updates its configuration changes.

15.11.7 Recovering an Oracle Instance in an Oracle Fusion Customer Relationship Management Installation

To recover the Oracle instance `ohs` in an Oracle Fusion Customer Relationship Management installation:

1. Stop all relevant processes using `opmnctl` or kill all processes that are related to that Oracle instance. For example:

```
opmnctl stopall
```

2. Copy the original Oracle instance to a backup location in case you need to retrieve some configuration data:

```
cp /scratch/CRM/APPLICATIONS_CONFIG/webtier_mw_home/webtier/ohs
/scratch/CRM/APPLICATIONS_CONFIG/webtier_mw_home/webtier/ohs_backup
```

3. Restore the Oracle instance home from the backup file to a stage directory by restoring the Applications configuration directory:

```
mkdir stage
cd stage
tar -xvf crm_ApplConfig_backup_030212.tar
```

4. Copy the restored Oracle instance from step 3 to the original location:

```
cp /scratch/stage/APPLICATIONS_CONFIG/webtier_mw_home/webtier/ohs/
   /scratch/CRM/APPLICATIONS_CONFIG/webtier_mw_home/webtier/ohs
```

5. Start all relevant processes. That is, start all processes that are related to that Oracle instance:

```
opmnctl startall
```

Note: If you are recovering the Oracle instance containing Oracle HTTP Server to a new host after loss of host, see [Section 15.10.7.1](#).

15.11.8 Recovering an Oracle Fusion Customer Relationship Management Cluster

In this scenario, the cluster CRMPPerformanceCluster has been erroneously deleted or the cluster-level configuration, such as the JMS configuration or container-level data sources, was mistakenly changed and committed. The server CRMPPerformanceServer_1 cannot be started or does not operate properly or the services running inside the server are not starting. You cannot ascertain what was changed.

Caution: Performing a domain-level recovery can impact other aspects of a running system and all of the configuration changes performed after the backup was taken will be lost.

If the configuration changes are few, then the easiest way is to redo the configuration changes. If that is not feasible, restore the domain to which the cluster belongs. (See the table in [Section 15.11.5](#) for Managed Servers to domain mappings.) Take the following steps:

1. Stop the cluster:

```
stop('CRMPPerformanceCluster', 'Cluster')
```

2. Stop the Administration Server, as described in [Task 4, "Stop the Administration Servers"](#) in [Section 3.3.2.2](#).

3. Copy the original CRMDomain to a backup location in case you need to retrieve some configuration data:

```
cp /scratch/CRM/APPLICATIONS_CONFIG/instance/domains/CRMDomain
   /scratch/CRM/APPLICATIONS_CONFIG/instance/domains/CRMDomain_backup
```

4. Restore the CRMDomain home from the backup file to a stage directory by restoring the Applications configuration directory:

```
mkdir stage
cd stage
tar -xvf crm_ApplConfig_backup_030212.tar
```

5. Copy the restored CRMDomain from step 4 to the original location:

```
cp /scratch/stage/APPLICATIONS_CONFIG/instance/domains/CRMDomain  
/scratch/CRM/APPLICATIONS_CONFIG/instance/domains/CRMDomain
```

6. Start the Administration Server, as described in [Task 3, "Start the Administration Servers"](#) in [Section 3.3.2.1](#).
7. Start the cluster. You can use the Oracle WebLogic Server Administration Console or WLST. For example, to use the WLST start command:

```
start('CRMPeformanceCluster', 'Cluster')
```

Note: You can use this procedure to recover a cluster if the membership is mistakenly deleted.

15.11.9 Recovering Databases for Oracle Fusion Customer Relationship Management

You can recover the databases using RMAN. You can recover the databases at whatever level is appropriate by performing a restore and recovery of either the full databases, a tablespace, or a data file.

Oracle Fusion Customer Relationship Management uses the following databases: Oracle Fusion Applications database, LDAP database, Oracle WebCenter Content, and Oracle Identity Manager database. You must maintain consistency among the databases. If any of these databases are recovered to a different point in time, you may need to reconcile the databases, as described in [Section 15.9.9](#).

For more information, see the following topics:

- For data loss, corruption, or media failure loss, see [Section 15.9.8](#).
- For loss of host, see [Section 15.10.4](#).

Configuring High Availability and Scaling Out Oracle Fusion Applications

This chapter describes how to configure Oracle Fusion Applications for high availability and how to scale out an Oracle Fusion Applications environment.

This chapter contains the following topics:

- [Introduction to Oracle Fusion Applications High Availability Architecture](#)
- [Oracle WebLogic Cluster Configuration for Oracle Fusion Applications](#)
- [Oracle Fusion Applications High Availability](#)
- [Oracle RAC Database Configuration for Oracle Fusion Applications Repositories](#)
- [Scaling Out Oracle Fusion Applications](#)

The following appendixes describe additional high availability topics:

- [Appendix A, "High Availability for Oracle Fusion Middleware Extensions for Applications"](#)
- [Appendix B, "High Availability for Oracle Enterprise Scheduler"](#)

For information about configuring high availability for Oracle Fusion Middleware products such as Oracle WebLogic Server, Oracle SOA Suite, Oracle Identity Management, and Oracle Business Intelligence, see the *Oracle Fusion Middleware High Availability Guide*.

16.1 Introduction to Oracle Fusion Applications High Availability Architecture

Providing high availability for Oracle Fusion Applications involves configuring an Oracle WebLogic cluster for high availability of the middle tiers and configuring Oracle Real Applications Clusters (Oracle RAC) for high availability of the Oracle Database. It also involves scaling out Oracle Fusion Applications and Oracle RAC database instances.

Oracle Fusion Applications is built on standards-based Oracle Fusion Middleware. Therefore, it benefits from the high availability solutions for Oracle Fusion Middleware.

Characteristics of an Oracle Fusion Applications high availability configuration are:

- Oracle Fusion Applications are Oracle Application Development Framework (Oracle ADF) applications running on Oracle WebLogic Server Clusters.

- Oracle Fusion Applications connect to Oracle RAC databases using multi data sources.
- Oracle Fusion Applications use Enterprise JavaBeans (EJBs) for inter-cluster communication within the same domain.
- Oracle Fusion Applications use Web services for inter-domain communication.
- Oracle Fusion Applications use the Java Message Service (JMS) operational interface (shared interfaces) to Oracle Streams Advanced Queueing (AQ), known as AQ JMS, for inter-domain communication.

This chapter assumes that you have a basic understanding of high availability concepts such as active-active deployments, active-passive deployments, and disaster recovery deployments. For information about these concepts, see the "Oracle Fusion Middleware High Availability Framework" chapter in the *Oracle Fusion Middleware High Availability Guide*.

See the following documentation resources to learn more about high availability for Oracle Fusion Middleware and Oracle Database:

- *Oracle Fusion Middleware High Availability Guide*
- *Oracle Database High Availability Overview*
- *Oracle Database High Availability Best Practices*

16.2 Oracle WebLogic Cluster Configuration for Oracle Fusion Applications

Oracle Fusion Applications runs on Oracle WebLogic Server. During an Oracle Fusion Applications installation, each individual product family installs into its own WebLogic Server domain. At the time of provisioning, a WebLogic cluster (with one Oracle WebLogic Server) is created for each Java EE application that is part of the product family. If you want to scale out any of the Java EE applications, new Managed Servers can be added to the application's WebLogic cluster.

For more information about Oracle WebLogic Server clustering, see the "High Availability for WebLogic Server" chapter in the *Oracle Fusion Middleware High Availability Guide*.

16.3 Oracle Fusion Applications High Availability

This section provides a description of the single instance architecture of Oracle Fusion Applications, and then provides information that helps you deploy Oracle Fusion Applications in a high availability configuration.

If you have not already done so, read [Chapter 1](#) for an introduction to:

- Oracle Fusion Applications
- Oracle Fusion Applications product families and components
- Key Oracle Fusion Applications concepts

16.3.1 Oracle Fusion Applications Single Instance Architecture

[Figure 1–1](#) shows an Oracle Fusion Applications single instance architecture. These are stateful applications that maintain session state.

Oracle Fusion applications are built using the following technologies:

- ADF, Oracle Middleware Extensions for Applications (Applications Core), and Oracle WebCenter framework, which run as JavaEE applications on WebLogic Server
- Clustered EJBs. Inter-cluster communication within one domain is done using EJBs. Inter-cluster communication across different domains is done using web services (SOAP).
- Event Delivery Network (EDN) queues are used for event propagation. EDN uses Advanced Queueing (AQ) database queues.

This section contains the following topics that describe high availability details:

- [Oracle Fusion Applications External Dependencies](#)
- [Oracle Fusion Applications Run-Time Processes](#)
- [Oracle Fusion Applications Request Flow](#)
- [Oracle Fusion Applications Configuration Artifacts](#)
- [Oracle Fusion Applications Deployment Artifacts](#)
- [Oracle Fusion Applications Log File Location](#)

16.3.1.1 Oracle Fusion Applications External Dependencies

The following lists shows some of the external entities that Oracle Fusion Applications depends on that are relevant from a high availability perspective:

- Oracle Fusion Applications repository database
- Oracle Metadata Service (MDS) repository database
- Oracle Enterprise Scheduler Service (ESS) repository database and queues
- EDN data source (AQ queues)
- Identity Management infrastructure (for authentication/authorization)
- LDAP store for users, roles, and policies.
- Oracle Business Intelligence Foundation, which includes Essbase
- Web Services invocation (mostly asynchronous)
- Java Message Service (JMS), for Asynchronous Web Services invocation
- Oracle Data Integrator (for replication framework)
- Oracle Universal Content Management system

Note: When you configure Oracle Identity Manager for Fusion Applications in a high availability environment, Oracle recommends that you use Multimaster Replication. See [Setting up Multimaster Replication](#).

16.3.1.2 Oracle Fusion Applications Run-Time Processes

All Oracle Fusion Applications components run as JavaEE applications in WebLogic Server and do not start a process of their own. Standard WebLogic tools such as the Administration Console, Oracle WebLogic Scripting Tool (WLST), and Oracle Enterprise Manager Fusion Middleware Control start, stop, and manage the applications.

Oracle Fusion Applications also has C components such as Oracle Global Order Promising (GOP) and the Bounce Handling Daemon. Oracle Process Manager and Notification Server (OPMN) manage these.

16.3.1.3 Oracle Fusion Applications Request Flow

The typical client HTTP request goes from the web browser or Web Services client to Oracle HTTP Server to JDBC/RMI/WS (HTTP). [Table 16-1](#) shows the Oracle Fusion Applications clients, the protocols they use to connect, and their high availability configuration.

Table 16-1 Oracle Fusion Applications Client Connection Protocols for High Availability

Client	Connection Protocol	High Availability Configuration
Web browser	HTTP	Load balancer
ESS	HTTP (WS SOAP ¹ over HTTP)	Load balancer
Web Services invocation	HTTP (WS SOAP over HTTP)	Load balancer
Inter-cluster communication	EJB (RMI ² over T3)	Clustered provider URL
Oracle Data Integrator	JDBC ³	Oracle RAC/multi data source

¹ Web Services Simple Object Access Protocol

² Remote Method Invocation

³ Java Database Connectivity

16.3.1.4 Oracle Fusion Applications Configuration Artifacts

All configuration for Oracle Fusion Applications is stored in the Oracle Fusion Applications repository database or MDS repository.

Configuration for Oracle WebLogic is stored in the domain configuration.

16.3.1.5 Oracle Fusion Applications Deployment Artifacts

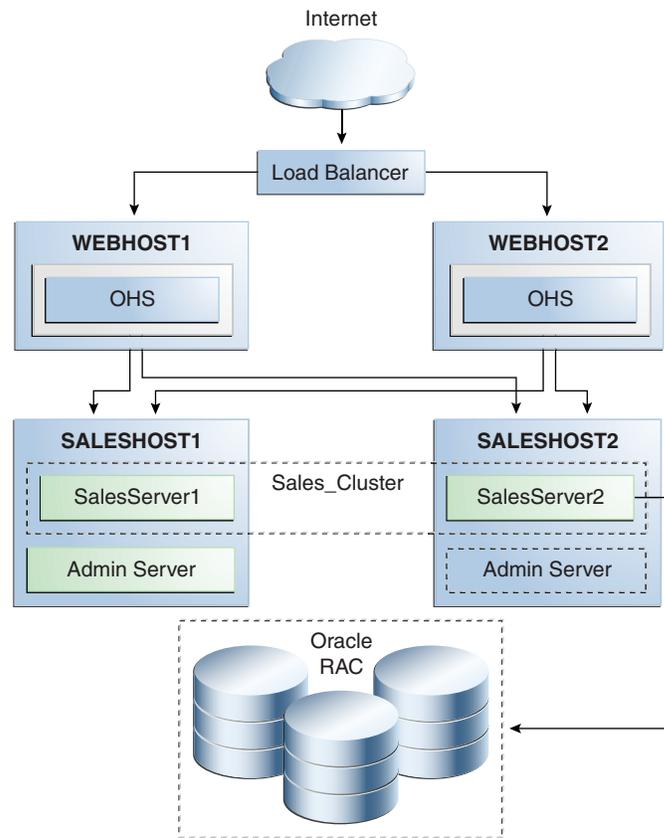
The Oracle Fusion applications are deployed using the nostaged deployment model. This allows you to install a single copy of the Oracle Fusion Applications binaries on a NAS or SAN shared storage and then run Fusion Applications from multiple computers using that single copy on the shared storage. Typically, application extensions get deployed on a different Managed Server than the application itself.

16.3.1.6 Oracle Fusion Applications Log File Location

For detailed information about Oracle Fusion Applications log file locations, see [Section 13.6.4](#).

16.3.2 Oracle Fusion Applications High Availability Architecture and Failover Considerations

[Figure 16-1](#) shows the high availability architecture of an Oracle Fusion Applications cluster deployment.

Figure 16–1 Oracle Fusion Applications High Availability Cluster Deployment

Some of the key characteristics of this architecture are:

- It is an active-active clustered architecture.
- It stores product binaries on shared storage.
- Each computer in the cluster has its own copy of the domain configuration on local storage.
- The application tier includes hosts SALESHOST1 and SALESHOST2, which comprise a WebLogic cluster with two Managed Servers running the Sales application.
- The WebLogic cluster is front-ended in the web tier by hosts WEBHOST1 and WEBHOST2 running Oracle HTTP Server. Oracle HTTP Server provides sticky routing by default.
- It has a hardware load balancer in front of the two Oracle HTTP Servers configured with sticky routing.
- The Oracle RAC repository in the data tier has multi data source configuration.
- ADF session state replication is enabled by default. The following high availability session state replication settings are enabled out of the box:
 - In `weblogic.xml`, the `replicate_if_clustered` option is specified:


```
<session-descriptor>
<persistent-store-type>replicated_if_clustered</persistent-store-type>
</session-descriptor>
```
 - In `adf-config.xml`, the `adf-scope-ha` parameter is set to `true`:

```
<adf-controller-config
xmlns="http://xmlns.oracle.com/adf/controller/config">
<adf-scope-ha-support>true</adf-scope-ha-support>
</adf-controller-config>
```

- In `adf-config.xml`, the `adf-scope-ha-support` parameter is set to `true`:

```
<adfc:adf-controller-config>
<adfc:adf-scope-ha-support>true</adfc:adf-scope-ha-support>
</adfc:adf-controller-config>
```

- The `jbo.dofailover` parameter is set to `true`, either in the startup environment as `-D` or in the ADF Business Component configuration file (`bc4j.xcfg`):

```
<AppModuleConfig ...
<AM-Pooling jbo.dofailover="true" />
</AppModuleConfig>
```

- In the `setDomainEnv.sh` file, change `JAVA_PROPERTIES` to the following value and then export `JAVA_PROPERTIES`:

```
JAVA_PROPERTIES="-Dplatform.home=${WL_HOME} -Dwls.home=${WLS_HOME}
-Dweblogic.home=${WLS_HOME} -Djbo.ampool.doampooling=true"
```

- EJBs used for inter-cluster communication within the same domain deploy uniformly to the WebLogic cluster and leverage WebLogic Server clustering capabilities.
- Web services inter-domain communication is routed through the load balancer.
- It uses the external dependency high availability configuration that [Table 16–1](#) describes.

This section contains the following topics about Oracle Fusion Applications changes and failover considerations:

- [Starting and Stopping the Oracle Fusion Applications Cluster](#)
- [Oracle Fusion Applications Cluster-Wide Configuration Changes](#)
- [Oracle Fusion Applications Failures and Expected Behaviors](#)

16.3.2.1 Starting and Stopping the Oracle Fusion Applications Cluster

See [Section 2.5](#) for descriptions of the various scenarios for starting and stopping the Oracle Fusion Applications environment.

16.3.2.2 Oracle Fusion Applications Cluster-Wide Configuration Changes

All Oracle Fusion Applications configuration changes can be done at the cluster level. You do not need to make configuration changes to individual instances in a cluster.

16.3.2.3 Oracle Fusion Applications Failures and Expected Behaviors

This section describes the different types of failures that can occur in an Oracle Fusion Applications high availability deployment, and the expected behaviors when these failures occur. This section contains the following topics:

- [Process Failure](#)
- [Machine Failure](#)
- [External Dependency Failures \(Oracle RAC Instance, Web Services Provider\)](#)

- [Troubleshooting Oracle Fusion Applications](#)

16.3.2.3.1 Process Failure Node Manager detects Managed Server failures and restarts the Managed Server automatically. If repeated restart attempts fail, then in case of a simple active-active deployment, surviving cluster members continue to service the request. The `ORACLE_INSTANCE/config/OHS/ohsName/mod_wl_ohs.conf` file for Oracle HTTP Server redirects user requests to a server which has the replica of the session state. The Node Manager `RestartInterval` parameter specifies the amount of time Node Manager attempts to start a failed server. The `RestartMax` parameter specifies the number of times Node Manager attempts to start a failed server. See the Node Manager Properties table in the *Oracle Fusion Middleware Node Manager Administrator's Guide for Oracle WebLogic Server* for more information about these parameters.

The server promotes the state to primary, creates a replica elsewhere in the cluster, and services the user request.

WebLogic Server clustering provides EJB failover for inter-cluster communication within the domain.

Web services inter-domain communication is routed through the load balancer, which also provides failover.

If a cluster member fails, messages continue to be available to surviving cluster members because AQ JMS is used for inter-domain communication.

When system components fail, OPMN attempts to restart the process in place.

16.3.2.3.2 Machine Failure When a computer fails, the expected behavior is the same as for process failure that [Section 16.3.2.3.1](#) describes except that Node Manager and OPMN do not attempt restarts.

16.3.2.3.3 External Dependency Failures (Oracle RAC Instance, Web Services Provider) This section describes external dependency failures.

- Oracle RAC Instance Failure. When an Oracle RAC instance fails, multi data sources and built-in application retry logic provide failover. Under certain circumstances, an Oracle RAC database outage generates the following message on your Web browser:

```
Encountered database recoverable exception. Please retry your request.
```

If this message appears, retry the ongoing session after logging in again (if required).

- Web Services Provider Failure. If the failure occurs before the Web Services stack is invoked, the application receives an error and can retry. If failure occurs after the acknowledgement is received, then the Asynchronous Web Services infrastructure guarantees a successful response.

16.3.2.3.4 Troubleshooting Oracle Fusion Applications If you experience issues with your Oracle Fusion Applications high availability deployment, check the WebLogic Server log. See [Section 16.3.1.6](#) for details.

16.4 Oracle RAC Database Configuration for Oracle Fusion Applications Repositories

Oracle Real Application Clusters (Oracle RAC) enables you to cluster servers to provide resilient and scalable access to a database. A **cluster** comprises multiple interconnected computers or servers that appear as if they are one server to end users

and applications. Oracle RAC simultaneously provides a highly available and scalable database for Oracle Fusion Applications.

This section provides an overview of Oracle RAC and information about setting up a highly available deployment for Oracle Fusion Applications configured with an Oracle RAC database as the persistent store for data.

Note: This chapter discusses configuring Oracle RAC databases only. Other persistent repositories that connect with Oracle Fusion Applications, such as Oracle Hyperion, are configured implicitly and transparently.

For complete information about Oracle RAC, see the *Oracle Real Application Clusters Administration and Deployment Guide*.

This section contains the following topics:

- [Oracle Real Application Clusters](#)
- [Providing High Availability for Oracle Fusion Applications Repositories](#)
- [Best Practices for Deploying JDBC Multi Data Sources on Servers and Clusters](#)

16.4.1 Oracle Real Application Clusters

Oracle RAC and Oracle Clusterware¹ allow the Oracle Database to run any packaged or custom application across a set of clustered servers. This capability provides the highest levels of availability and the most flexible scalability. If a clustered server fails, Oracle Database continues running on the surviving servers. When more processing power is needed, you can add another server without interrupting access to data.

Oracle RAC enables multiple database instances that are linked by an interconnect to share access to an Oracle database. In an Oracle RAC environment, the Oracle Database runs on two or more systems (nodes) in a cluster while concurrently accessing a single shared database. The result is a single database system that spans multiple hardware systems and enables Oracle RAC to provide high availability and redundancy during failures in the cluster. Oracle RAC accommodates all system types, from read-only data warehouse (DSS) systems to update-intensive online transaction processing (OLTP) systems.

The clusters that are typical of Oracle RAC environments can provide continuous service for both planned and unplanned outages. Oracle RAC builds higher levels of availability on top of the standard Oracle features. All single-instance high availability features, such as the Oracle Flashback technologies and online reorganization, also apply to Oracle RAC. Applications scale in an Oracle RAC environment to meet increasing data processing demands without changing the application code. In addition, allowing maintenance operations to occur on a subset of components in the cluster while the application continues to run on the rest of the cluster can reduce planned downtime.

For information about Oracle RAC design and deployment techniques, see the "Design and Deployment Techniques" chapter of the *Oracle Real Application Clusters Administration and Deployment Guide*.

¹ Oracle Clusterware is software that enables the servers to be bound together to operate as if they are one server and manages the availability of user applications and Oracle databases.

For information about online application maintenance and upgrades, see the "Online Application and Maintenance and Upgrades" section of *Oracle Database High Availability Overview*.

16.4.2 Providing High Availability for Oracle Fusion Applications Repositories

Providing high availability requires detailed planning to ensure there are no single points of failure throughout the infrastructure. Even though Oracle RAC makes your database highly available, if a critical application becomes unavailable, then your business can be negatively affected.

For example, because Oracle Fusion Applications uses the Lightweight Directory Access Protocol (LDAP) for authentication, the best practice is to make the LDAP server highly available. If the database is up but the LDAP server is down, users cannot connect to the applications and the entire system appears to be down.

The following sections describe deploying Oracle RAC for high availability in an Oracle Fusion Applications environment:

- [Preparing the System to Use Oracle RAC](#)
- [Providing High Availability Through Multi Data Sources](#)
- [Failover Processing and Recovery for Oracle RAC Node or Instance Failures](#)
- [Load Balancing Across Oracle RAC Nodes](#)
- [Retry Logic to Protect Read-Only Operations](#)

16.4.2.1 Preparing the System to Use Oracle RAC

Use Oracle Fusion Applications provisioning to orchestrate the integration process, including the specification of your Oracle RAC database.

For more information about the provisioning and installation process and installing Oracle Database, see [Section 1.2.1.2](#).

16.4.2.2 Providing High Availability Through Multi Data Sources

Oracle Fusion Applications support only multi data source configurations for Oracle RAC. Multi data sources encapsulate individual data sources that pool connections to specific instances of Oracle RAC.

A **data source** is a Java object that application components use to obtain connections to a relational database. Specific connection information, such as the URL or user name and password, are set on a data source object as properties and do not need to be explicitly defined in an application's code. This abstraction allows applications to be built in a portable manner, because the application is not tied to a specific back-end database. The database can change without affecting the application code.

A **multi data source** is like a pool of data sources that provides failover processing and recovery. It also provides load balancing between nodes of a highly available database system, such as an Oracle RAC database. When an Oracle RAC instance fails, a multi data source determines which data source to use to satisfy application requests. If you are using an Oracle RAC database, you must configure multi data sources.

The data source member list for a multi data source supports dynamic updates. This allows Oracle RAC environments to add and remove database nodes and corresponding data sources without redeployment, and it provides the ability to:

- Grow (scale) and shrink Oracle RAC clusters in response to throughput
- Shut down Oracle RAC nodes temporarily to perform planned maintenance

See [Section 16.4.3](#) for best practices when configuring multi data sources for Oracle RAC and Oracle Fusion Applications. For more information about using multi data sources with Oracle RAC, see the "Using Multi Data Sources with Oracle RAC" section in the *Oracle Fusion Middleware High Availability Guide*.

16.4.2.3 Failover Processing and Recovery for Oracle RAC Node or Instance Failures

WebLogic Server periodically checks the status of data sources in a multi data source. If an Oracle RAC node or instance fails, application connection requests are managed as follows:

- Existing connections

There is no failover of existing connections. In-flight transactions are typically rolled back when the database is the transaction manager. When the WebLogic Server is the Transaction Manager, in-flight transactions are failed over; they are driven to completion or rolled back based on the transaction state at the time of failure.
- New connection requests

New session requests are redirected to a working Oracle RAC instance in the cluster, either by Oracle WebLogic Server or by the Oracle Thin driver.

For information about multi data source failover limitations and requirements, see the "Configuring JDBC Multi Data Sources" chapter in *Oracle Fusion Middleware Configuring and Managing JDBC Data Sources for Oracle WebLogic Server*.

16.4.2.4 Load Balancing Across Oracle RAC Nodes

Load balancing involves distributing requests among two or more processes. Oracle WebLogic Server provides a load balancing algorithm for multi data sources. If an application requires load balancing across Oracle RAC nodes, then WebLogic Server supports this capability through the use of JDBC multi data sources that have been configured for load balancing.

WebLogic accesses the data sources that form a multi data source using a round-robin scheme. When switching connections, WebLogic Server selects a connection from the next data source in the order listed.

See [Section 16.4.3](#) for recommendations when configuring load balancing with WebLogic Server.

16.4.2.5 Retry Logic to Protect Read-Only Operations

Under certain circumstances, an Oracle RAC database outage generates the following message:

```
Encountered database recoverable exception. Please retry your request.
```

If this message appears, retry the ongoing session after logging in again (if required).

See [Section 16.4.3](#) to configure retry properties with WebLogic Server.

16.4.3 Best Practices for Deploying JDBC Multi Data Sources on Servers and Clusters

When you create an Oracle Fusion Applications environment using an Oracle RAC database, the provisioning tool configures automatically a multi data source to use each Oracle RAC database instance.

The default configuration uses these best practices that provide optimal availability:

- Deploys a multi data source to a cluster or server by selecting the server or cluster as a deployment target in the WebLogic Server Administration Console.

When a multi data source is deployed on a server, WebLogic Server creates an instance of the multi data source on the server. When you deploy a multi data source to a cluster, WebLogic Server creates an instance of the multi data source on each server in the cluster.

- Deploys all data sources that are used by a multi data source to satisfy connection requests on the same servers and clusters as the multi data source.
Multi data sources do not route connection requests to other servers in a cluster or in a domain.
- Sets the load-balancing properties as shown in [Table 16-2](#).

Table 16-2 Recommended Property Settings for the Load-Balancing

Property Name	Recommended Value
test-frequency-seconds	120
algorithm-type	Load-Balancing

- If the data sources require support for non-XA drivers, then the properties for the individual data sources are set as shown in [Table 16-3](#).

Table 16-3 Recommended Data Source Property Settings for Non-XA Drivers

Property Name	Recommended Value
Driver	oracle.jdbc.OracleDriver
Property command	<property> <name>oracle.net.CONNECT_TIMEOUT</name> <value>10000</value> </property>
connection-creation-retry-frequency-seconds	60
test-frequency-seconds	120
test-connections-on-reserve	true
test-table-name	SQL SELECT 1 FROM DUAL
seconds-to-trust-an-idle-pool-connection	0
global-transactions-protocol	None

For more information about configuring JDBC multi data sources, see "Configuring JDBC Multi Data Sources" in *Oracle Fusion Middleware Configuring and Managing JDBC Data Sources for Oracle WebLogic Server*.

16.5 Scaling Out Oracle Fusion Applications

Scale-out enables organizations to improve application performance, scalability, and availability on an incremental, as-needed basis by adding multiple replicated servers.

Oracle Fusion Applications is always deployed in a cluster. Even in a single instance deployment, the Oracle Fusion Applications instance is part of a cluster with one member.

You can use a scale out operation to move from a non-high availability deployment to a high availability deployment. You can scale up or scale out the Oracle Fusion Applications topology. When you *scale up* the topology, you add a new Managed Server to computers that are already running one or more Managed Servers. When you *scale out* the topology, you add a new Managed Server to new computers.

Before you run scale up or scale out steps, check that you meet these requirements:

- There must be at least one existing computer running at least one Managed Server configured with Oracle Fusion Applications within the topology. The computer contains an `APPLICATIONS_BASE` home directory (`/oracle` in [Figure 1-6](#)) and an Oracle Fusion Applications Middleware home directory (`/oracle/fusionapps` in [Figure 1-6](#)) in shared storage.
- The computer on which the new Managed Server is deployed can access the existing home directories for Oracle Fusion Applications. Use the existing installations in shared storage for creating a new Managed Server.

For *scale up*, you do not need to install Oracle Fusion Applications binaries in a new location, and you do not need to run pack and unpack to bootstrap the domain configuration.

For *scale out*, you must run pack and unpack to bootstrap the domain configuration in the new computer.

Note: The steps below assume that Oracle Fusion Applications clusters such as the Sales cluster do not use WebLogic JMS or XA.

The scale up steps for Oracle Fusion Middleware components such as Oracle SOA Suite and Oracle WebCenter are in the scale up section of the component's chapter in the *Oracle Fusion Middleware High Availability Guide*. Scale out steps for Oracle Fusion Middleware components such as Oracle SOA and Oracle WebCenter are in the scale out section of the component's chapter in the *Oracle Fusion Middleware High Availability Guide*. To make Oracle HTTP Server highly available, see the Oracle HTTP Server chapter in the *Oracle Fusion Middleware High Availability Guide*.

This section includes the following topics:

- [Section 16.5.1, "Scaling Up the Topology \(Adding Managed Servers to Existing Machines\)"](#)
- [Section 16.5.2, "Scaling Out the Topology \(Adding Managed Servers to New Machines\)"](#)
- [Section 16.5.3, "Scaling Out Oracle RAC Databases"](#)

16.5.1 Scaling Up the Topology (Adding Managed Servers to Existing Machines)

In this case, you have a computer that runs a Managed Server configured with Oracle Fusion Applications components and you want to add another Managed Server to that computer. Use the steps in this section to add a second Managed Server to a computer and to add subsequent Managed Servers on that computer.

Follow these steps to scale up the topology. This example uses the Oracle Fusion Customer Relationship Management application; you can also use these steps to scale up other Oracle Fusion Applications clusters:

1. Using the Administration Console, clone the source Managed Server (for example, `SalesServer1`) into a new Managed Server. The source Managed Server to clone should be one that already exists on the computer where you want to run the new Managed Server.

To clone a Managed Server:

- a. Select **Environment** -> **Servers** from the Oracle WebLogic Server Administration Console.
- b. Select the Managed Server that you want to clone (for example, `SalesServer1`).
- c. Select **Clone**.

Name the new Managed Server `server_nameN` (for example, `SalesServer2`), where `N` is a number to identify the new Managed Server.

2. For the listen address, assign the host name to use for this new Managed Server.

Ensure that the port number (9001, in this example, so as not to conflict with the current port 8001 that the existing Managed Server, `SalesServer1`, uses) for the new Managed Server is not used on this computer.

3. Update the cluster address to include the new server:

- a. Select **Environment** -> **Cluster** from the Administration Console.
- b. Click on the `Sales_Cluster` server.
- c. In the Change Center, click **Lock & Edit**.
- d. Add the new server's address and port to the **Cluster Address** field. For example:

```
SALESHOST1:8001,SALESHOST1:9001
```

4. Disable host name verification for the new Managed Server. Before starting and verifying the `SalesServern` Managed Server, you must disable host name verification. You can re-enable it after you configure server certificates for the communication between the Oracle WebLogic Administration Server and the Node Manager in `SALESHOSTn`. If the source server from which the new one has been cloned had already disabled hostname verification, these steps are not required; the hostname verification setting propagates to the cloned server.

To disable host name verification:

- a. Start the Oracle WebLogic Administration Console.
- b. Expand the **Environment** node in the Domain Structure window.
- c. Click **Servers**.
- d. Select `SalesServern` in the Names column of the table in the Summary of Servers page.
- e. Click the **SSL** tab in the Settings page for the server.
- f. Click **Advanced**.
- g. Set Hostname Verification to **None**.
- h. Click **Save**.

5. Change JDBC LLR Table Name:
 - a. In the Oracle EM, expand the Environment node in the Domain Structure window
 - b. Click **Servers**.
 - c. Select the scaled out server.
 - d. In the General tab, click **Advanced**.
 - e. Change the JDBC LLR Table Name to the next higher number so that it doesn't conflict with the number of the original server.
 - f. Click **Save**.

6. Ensure that the Node Manager is running on the computer. To start the Node Manager, use the installation in shared storage from the existing computer and start Node Manager by passing the host name of the computer as a parameter as follows:

```
SALESHOST1> WL_HOME/server/bin/startNodeManager computer_ip
```

7. Start and test the new Managed Server from the Oracle WebLogic Server Administration Console:
 - a. Ensure that the newly created Managed Server (SalesServer2 in this example) is running.
 - b. Access the application on the newly created Managed Server (<http://SALESHOST1:9001/sales/faces/index>). The application should be functional.

8. Configure Oracle HTTP Server to route to the Administration Server that contains the new Managed Server; set the `WebLogicCluster` parameter to the list of computers in the cluster. Follow these steps for all Oracle HTTP Server hosts in your deployment:

- a. `cd` to the following directory in the Web Tier configuration home:

```
(UNIX) APPLICATIONS_BASE/instance/SetupDomain_webtier/config/OHS/ohs1/moduleconf
```

```
(Windows) APPLICATIONS_BASE\instance\CommonDomain_webtier\config\OHS\ohs1\moduleconf
```

- b. Create a backup of the `FusionVirtualHost_crm.conf` file and then make these edits to the file (the example is for the Sales server):

```
<Location /sales>
  SetHandler weblogic-handler
  WebLogicCluster SALESHOST1:8001,SALESHOST1:9001
</Location>
```

Where `portn` is a different port number than the port used for the first Managed Server.

- c. Restart the Oracle HTTP Server by following these steps:

```
(UNIX) cd APPLICATIONS_BASE/instance/SetupDomain_webtier/bin
opmnctl stopall
opmnctl startall
```

```
(Windows) cd APPLICATIONS_BASE\instance\SetupDomain_webtier\bin
opmnctl stopall
```

```
opmnctl startall
```

- d. Check that both the SALESHOST1:8001/sales/faces/index and SALESHOST1:9001/sales/faces/index are accessible.

- e. Check the Sales servers from Oracle HTTP Server:

```
WEBHOST1:7777/sales/faces/index
```

This example assumes 7777 is being used for the Oracle HTTP Server port.

- f. Login to the WebLogic Server Administration Console and navigate to **CRMDomain > Environments > Servers**.

16.5.1.1 Testing the Routing from Oracle HTTP Server to the New Managed Server

To test the routing from Oracle HTTP Server to the new Managed Server.

1. If desired, shut down the existing Managed Server (in this case, SalesServer1).
2. Access the application using WEBHOST1:7777/sales/faces/index through the web server, which should access the application on the new Managed Server (in this case, SalesServer2).

This step should show that Oracle HTTP Server routed the request to the new Managed Server (SalesServer2) while SalesServer1 was down, and that SalesServern serviced the request.

3. If you have not already done so, restart the Managed Server that you shut down.

16.5.2 Scaling Out the Topology (Adding Managed Servers to New Machines)

When you *scale out* the topology, you add new Managed Servers configured with Oracle Fusion Applications to new computers where no Oracle Fusion Applications Managed Servers are configured.

You scale out the topology using Enterprise Manager Cloud Control. Benefits of Cloud Control include:

- **Ease of use.** Cloud Control automatically completes the provisioning and deployment processes. It gives you all the information you need about a Fusion Instance, Fusion Product Family, Fusion Product, and Fusion Application instances.
- **Increased reliability.** The Cloud Control process uses a job system to automatically and regularly rediscover WebLogic domains.
- **Enhanced interface.** Home pages enable you to monitor instances and product families at a glance. The interface makes transparent the relationship between the server and application deployment details. Cloud Control shows cluster applications from a product perspective.

Enterprise Manager Cloud Control enables you scale out a Fusion Product or a specific Fusion Application by:

- Provisioning a new WebLogic Server instance within the cluster
- Deploying an instance of the Fusion Application on the new WebLogic Server instance

There are two ways to scale out Fusion Applications:

- To scale out a Fusion Product Family, Product, or Fusion Cluster Application target, see ["Scaling Out a Fusion Product Family, Product, or Fusion Cluster Application Target"](#).
- To scale out a Fusion Instance target, see the topic "Fusion Application Instance Scale Out - Scale Out Selections" in the Cloud Control online help.

Scaling Out a Fusion Product Family, Product, or Fusion Cluster Application Target

After you provide the information required to provision the new WebLogic Server instance, Enterprise Manager invokes a deployment procedure to automatically complete the provisioning and deployment processes.

To scale out a Fusion Product Family, Product, or Fusion Cluster Application target:

1. From the **Targets** menu, choose **Fusion Applications** to open the Fusion Applications target home page.
2. Select one of the following: **Fusion Product Family**, **Product**, or **Fusion Cluster Application**.
3. Choose **Provisioning > Scale Out** from the **Product Family**, **Product**, or **Fusion Cluster Application** menu for the selected target.

4. Click on the Fusion Product or Fusion Application you want to scale out in the topology view in the region on the left.

Note that some Fusion Products may be associated with up to two Fusion Applications in the topology view.

The WebLogic Server cluster(s) that the selected Fusion Product or Fusion Application is currently deployed appears in the region on the right.

5. Click **Scale Out** in the row for a cluster you want to add an instance to.
6. Enter the credentials for the Administration Server user in the **WebLogic Domain Administrator Username** and **Password** fields.

The Administration Server is a WebLogic Server instance that manages the WebLogic Server domain the cluster belongs to.

7. Specify a directory to store temporary data files during the scale out operation in the **Administration Host Server Working Directory** field; for example, `/tmp`.
8. Enter the credentials required to access the Administration Server host.
9. Specify the following for the new WebLogic Server instance:
 - The hostname for the host that the instance will be provisioned to.
 - A unique identifier for the new instance.
 - The port that the instance will use.

10. Click **Configure More Details**.

11. Enter any additional information required to access the host machine that the new WebLogic Server instance will be created on. Required information includes:
 - The credentials required to access the host.
 - The directory to store temporary data files during the scale out operation; for example, `/tmp`
 - The port that the WebLogic Server instance will use.

12. Click **OK**.
13. Click **Scale Out**.

Enterprise Manager creates a Deployment Procedure job to provision the new WebLogic Service instance, add it to the cluster, and deploy the Fusion Application instance to it.

Note: Follow the preceding steps to scale out infrastructure applications (SOA, ESS).

16.5.3 Scaling Out Oracle RAC Databases

This section provides guidance on how to add nodes to existing Oracle RAC environments by using Oracle cloning, and then configuring the multi data source to route database requests to the new Oracle RAC instance.

This section contains the following topics:

- [Cloning Oracle RAC Nodes to Scale Out Oracle Fusion Applications](#)
- [Configuring the Multi Data Source to Include New Oracle RAC Nodes](#)

16.5.3.1 Cloning Oracle RAC Nodes to Scale Out Oracle Fusion Applications

The Oracle RAC cloning procedures assume that you successfully installed and configured an Oracle RAC environment that you want to add nodes and instances to. To add nodes to an Oracle RAC environment using cloning, you must first extend the Oracle Clusterware configuration, then extend the Oracle Database software with Oracle RAC, and then add the listeners and instances by running the Oracle assistants.

For more information about cloning nodes within the same cluster, see the chapter about "Using Cloning to Extend Oracle RAC to Nodes in the Same Cluster" in *Oracle Real Application Clusters Administration and Deployment Guide*.

For Oracle cloning that is specific to UNIX system environments, see the chapter about "Adding and Deleting Oracle RAC from Nodes on Linux and UNIX Systems" in *Oracle Real Application Clusters Administration and Deployment Guide*.

For Oracle cloning that is specific to Windows system environments, see the chapter about "Adding and Deleting Oracle RAC from Nodes on Windows Systems" in *Oracle Real Application Clusters Administration and Deployment Guide*.

16.5.3.2 Configuring the Multi Data Source to Include New Oracle RAC Nodes

After adding an Oracle RAC node to the database cluster as a part of scaling out the database, you must configure the multi data source to route database requests to the new Oracle RAC instance. You do this by adding a data source to each of the multi data sources that will route database requests to the newly added instance of the database.

Note: Verify that the property values of the data source that you add match the property values in the other peer data source of the multi data source. See [Table 16–3, "Recommended Data Source Property Settings for Non-XA Drivers"](#) for a list of property values.

The example configures a multi data source with the properties shown in [Table 16–4](#) but you can substitute the names and properties to scale any Oracle Fusion Applications cluster.

Table 16–4 Multi Data Source Example to Scale Out Oracle RAC with Oracle Fusion Applications

Property Name	Property Value
name	ApplicationDB
Jndi-name	jdbc/ApplicationDBDS
Data Sources	ApplicationDB01, ApplicationDB02
Target	Sales_Cluster

To configure the new Oracle RAC instance in a multi data source:

1. Log in to the Oracle WebLogic Administration Console.
2. Create a new data source, as follows:
 - a. Navigate to *domain_name* > **Services** > **Data Sources**
For example, navigate to **CRM_Domain** > **Services** > **Data Sources**.
 - b. In the Change Center, click **Lock & Edit** (if needed).
 - c. On the Summary of JDBC Data Sources page, click **New** to create a new data source. (Choose **Generic Data Source**.)
 - d. In the **Name** field, enter the name of the computer (ApplicationDB03) and enter the JNDI Name. The Database Type is **Oracle**.
 - e. Click **Next**.
 - f. For the Database Driver, choose the driver of type **Thin** and **RAC Server-Instance Connection**.
 - g. On the next screen, add details for this data source by providing Oracle RAC instance details for the new Oracle RAC instance, plus other properties matching its peer datasources in the Multi Data Source.
 - h. Click **Finish**.
 - i. In the Change Center, click **Activate Changes**.
 - j. In the Change Center, click **Lock & Edit**.
 - k. Navigate to *domain_name* > **Services** > **Data Sources**.
 - l. On the Summary of JDBC Data Sources page, choose the multi data source ApplicationDB.
 - m. On the right-hand pane for Settings for ApplicationDB, click the **Data Sources** tab.
 - n. In the Available list, find the newly created data source (ApplicationDB03) and click it to add it to the multi data source.
 - o. In the Change Center, click **Save and Activate Changes**.

Moving Components for Oracle Fusion Applications Across Environments

This chapter describes how to move Oracle Fusion Applications components from one environment to another.

This chapter contains the following topics:

- [Introduction to Moving Oracle Fusion Applications](#)
- [Understanding the Full-Movement Tasks](#)
- [Getting Started with Moving Components](#)
- [Moving the Oracle Identity Management Domain Component Artifacts](#)
- [Moving Binary Files to a Production Target Environment](#)
- [Moving Configurations to a Production Target Environment](#)
- [Moving Binary and Configuration Files to a Non-Distributed Target Environment](#)
- [Moving Functional Setup Data](#)
- [Completing Post-Movement Tasks](#)
- [A Case Study: Moving Oracle Fusion Supply Chain Management Components to a Target Distributed Environment](#)

17.1 Introduction to Moving Oracle Fusion Applications

Replicating an Oracle Fusion Applications environment requires moving Oracle Fusion Applications components from one environment to another. The task of moving Oracle Fusion Applications components from one environment to another is simplified by movement tools (scripts for moving binary and configuration information). These tools minimize the amount of work that would otherwise be required to reapply all the customization and configuration changes made in one environment to another.

Typically, Oracle Fusion Applications is installed, configured, customized, and validated in a source environment, such as a test environment. Once the system is stable and performs as desired, the target environment, such as a production environment, would then be created by moving all the components from the source environment, instead of redoing all the changes that were incorporated into the source environment.

The movement of Oracle Fusion Applications components from one environment to another can be achieved in several different movement scenarios, but this chapter focuses on the full-movement scenario.

In a **full-movement scenario**, the target environment does not exist. First, the source environment is created, configured, customized, and tested. Then, the target environment is created by moving all the components along with their configurations from the source environment.

When moving Oracle Fusion Applications components from one environment to another, you must adhere to the following guidelines:

- The source and target environments must share the same operating system and the same platform architecture (in terms of number of bits). For example, you cannot move components and their configurations from a source environment that is running Microsoft Windows to a target environment that is running UNIX. Similarly, you cannot move components and their configurations from a source environment with a 32-bit platform to a target environment with a 64-bit platform.
- The source and target environments must have identical topologies. (Any changes to the target environment topology can be done post movement).
- The Oracle Fusion Applications domain name must not be changed in the provisioned environment or while performing the full-movement tasks.
- The path relative to `APPLICATIONS_CONFIG` that is used while moving the domain and system component should be the same as the path used in the provisioned source environment.
- The Common domain should be moved to the target environment first because Metadata Services (MDS) operations for Oracle Fusion Applications partitioning is handled in the Common domain.

With the completion of the full-movement tasks, the following artifacts are moved from the source environment to the target environment:

- Seed data (created with Oracle Fusion Middleware Repository Creation Utility (RCU) in the target environment)
- Oracle WebLogic Server domain configuration, stored in the file system
- System component configuration, stored in the file system
- Configuration and metadata stored in MDS, such as Oracle Application Development Framework (Oracle ADF) connections, service-oriented architecture (SOA) composites, and so on
- Configuration and metadata stored in component-specific schemas outside of MDS
- Non-user layer customizations, such as Site and Enterprise Layer, in MDS
- Security artifacts created by the Oracle Fusion Applications Provisioning framework, such as application IDs, policies, and so on
- Functional setup data

Following are some of the artifacts that are not moved to the target environment:

- Transactional data, which is assumed to be test transactions
- Transient data, such as job status, process status, and so on
- Content data, such as Oracle WebCenter Content data, wiki pages, discussion forums, and so on (which can be moved using data movement tools if necessary)

Note: In the case of Oracle HTTP Server (OHS), the contents of document root are moved as part of OHS movement, as they are part of application data.

- Any configuration or metadata associated with individual users on the source, such as:
 - Users in the identity store
 - User layer customizations in MDS
 - Component configurations associated with specific users, such as Oracle Human Workflow work groups associated with users

Users are assumed to be test users, and they may not be present in the target environment.

- MDS sandboxes

17.2 Understanding the Full-Movement Tasks

Moving Oracle Fusion Applications components from one environment to another is a complex process involving multiple tasks.

[Table 17–1](#) lists the high level set of tasks for moving artifacts from the source environment to the target environment in a full-movement scenario. It also provides references to detailed information for each task.

In any given Oracle Fusion Applications environment, there are a certain number of Oracle WebLogic Server domains and system component instances. If the Administration Servers of all these domains and all the system component instances run on the same host, then the environment is **non-distributed**. In all other cases, the environment is **distributed**. If you have a non-distributed environment, which is typical of a QA environment, you can reduce the number of steps required for Tasks 3 and 4 by performing the procedure in [Section 17.7](#). If you have a distributed environment, you must use [Section 17.5](#) and [Section 17.6](#) to complete Tasks 3 and 4.

Table 17–1 Full-Movement Tasks

Task	Description	Documentation Reference
1	Complete the prerequisite tasks.	Section 17.3, "Getting Started with Moving Components"
2	Move the Oracle Identity Management domain component artifacts.	Section 17.4, "Moving the Oracle Identity Management Domain Component Artifacts"
3	<p>Copy the binaries (that is, the installed applications or components along with any applied patches) in the source environment, and apply the archive to the target environment to move the Oracle Fusion Applications binary files.</p> <ol style="list-style-type: none"> 1. Run the <code>copyBinary</code> script for each Middleware home in the source environment to create the binary archives. 2. Apply the copied binary files from the source environment to the target environment by running the <code>pasteBinary</code> script in the target environment. 3. Move the <code>dbclient</code> installation. 	<p>Section 17.5, "Moving Binary Files to a Production Target Environment"</p> <p>Note: If you have a non-distributed environment, you can alternatively use Section 17.7, "Moving Binary and Configuration Files to a Non-Distributed Target Environment," to reduce the number of steps.</p>

Table 17–1 (Cont.) Full-Movement Tasks

Task	Description	Documentation Reference
4	<p>Move the configurations of each Oracle WebLogic Server domain, OHS instance, and Node Manager in Oracle Fusion Applications from the source environment to the target environment.</p> <ol style="list-style-type: none"> 1. Run the <code>copyConfig</code> script for each Oracle WebLogic Server domain, OHS instance, and Node Manager in the source environment to create the configuration archives. 2. Extract the move plan by running the <code>extractMovePlan</code> script on the list of configuration archive locations, separated by comma. 3. Modify the move plan to specify properties for the target environment. 4. Apply the copied configurations from the source environment to the target environment by running the <code>pasteConfig</code> script for each Oracle WebLogic Server domain, OHS instance, and Node Manager separately, using the same move plan. 5. If there are multiple machines for the domain, recreate the local domain directory using the Oracle WebLogic Server pack and unpack mechanisms. 6. Complete the component-specific configuration move by performing additional movement tasks on the Oracle Fusion Middleware components. 	<p>Section 17.6, "Moving Configurations to a Production Target Environment"</p> <p>Note: If you have a non-distributed environment, you can alternatively use Section 17.7, "Moving Binary and Configuration Files to a Non-Distributed Target Environment," to reduce the number of steps.</p>
5	Apply the setup configuration data of the source environment to the target environment.	Section 17.8, "Moving Functional Setup Data"
6	Complete post-movement tasks.	Section 17.9, "Completing Post-Movement Tasks"

Note: Oracle Fusion Applications movement has a default logging level of `INFO`. To ensure `copyBinary`, `copyConfig`, `pasteBinary`, `pasteConfig`, and `extractMovePlan` scripts, set the `T2P_JAVA_OPTIONS` variable to include the following:

```
-Dt2p.logging.level=log_level
```

where `log_level` is one of the following (ordered from least verbose to most verbose):

```
SEVERE
WARNING
CONFIG
FINE
FINER
FINEST
ALL
```

For an example of moving Oracle Fusion Applications product family components from a source environment to a target environment using these full-movement tasks, see [Section 17.10](#).

17.3 Getting Started with Moving Components

Before you begin to move Oracle Fusion Applications components across environments, you must complete the following tasks:

- Prepare the source environment. See [Section 17.3.1](#).
- Prepare the target environment. See [Section 17.3.2](#).

17.3.1 Preparing the Source Environment

Before performing the full-movement tasks, you must make sure that the source environment is completely functional for the following requirements:

1. The database server is installed and the required schemas are created. The seed data must also be loaded.

See the "Installing a Transaction Database" chapter in the *Oracle Fusion Applications Installation Guide*.

2. The Oracle Identity Management environment is set up.

See the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Identity Management (Oracle Fusion Applications Edition)*.

3. All the necessary software has been installed through Oracle Fusion Applications provisioning.

See the *Oracle Fusion Applications Installation Guide*.

4. The setup data has been created using Oracle Fusion Functional Setup Manager (Functional Setup Manager).

See the *Oracle Fusion Applications Information Technology Management, Implement Applications Guide*.

5. All needed customizations and extensions have been applied to the source environment.

See the *Oracle Fusion Applications Extensibility Guide*.

6. Make sure your environment is fully ready to be moved. For example, if you have performed any edits of the Oracle BI repository, you must upload the final repository file through Fusion Applications Control before performing the movement.

See the "Configuring Repositories" chapter in the *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

7. The environment has been tested, and any necessary patches have been applied.

See the *Oracle Fusion Applications Patching Guide*.

17.3.2 Preparing the Target Environment

Before performing the full-movement tasks, you must complete the following preparations in the target environment:

1. Ensure the database server is installed. You do not need to run the Oracle Fusion Applications Repository Creation Utility (Applications RCU) to create schemas.

See the "Installing a Transaction Database" chapter in the *Oracle Fusion Applications Installation Guide*.

- For either UNIX or Windows, at an operating system command prompt, navigate to the location of the T2PDataUtil executable on both the source and the target:

```
(UNIX) FA_ORACLE_HOME/applications/lcm/t2p/bin
(Windows) FA_ORACLE_HOME\applications\lcm\t2p\bin
```

FA_ORACLE_HOME is located under the fusionapps Oracle Fusion Middleware home directory.

- Run the T2PDataUtil script from the source:

```
(UNIX) T2PDataUtil.sh -dbhost host -dbsid sid -dbport port -mdir move_directory
-command generate_move_data
(Windows) T2PDataUtil.cmd -dbhost host -dbsid sid -dbport port mdir move_
directory -command generate_move_data
```

The following table describes the options for the T2PDataUtil command for the source environment:

Option	Description
-dbhost	The Oracle Database host name.
-dbsid	The SID for the Oracle Database.
-dbport	The listen port number for the Oracle Database.
-mdir	The directory where the move data will be placed. This data is used in Step 5. If this option is not specified, the T2PDataUtil script defaults to the following locations: (UNIX) FA_ORACLE_HOME/applications/lcm/t2p/movedata (Windows) FA_ORACLE_HOME\applications\lcm\t2p\movedata
-command	The command to execute. Enter generate_move_data.

- Clone the production Oracle Database from the test Oracle Database.

See the "Duplicating a Database" section in the *Oracle Database Backup and Recovery User's Guide*.

- Run the T2PDataUtil script on the target. Ensure the move data specified in mdir option in Step 2 is accessible by the script.

```
(UNIX) T2PDataUtil.sh -dbhost host -dbsid sid -dbport port mdir move_directory
-command apply_move_data -arh /tmp/apps_rcu_home
(Windows) T2PDataUtil.cmd -dbhost host -dbsid sid -dbport port mdir move_
directory -command apply_move_data -arh /tmp/apps_rcu_home
```

The following table describes the options for the T2PDataUtil command for the target environment:

Option	Description
-dbhost	The Oracle Database host name.
-dbsid	The SID for the Oracle Database.
-dbport	The listen port number for the Oracle Database.

Option	Description
-mdir	The directory where the move data will be placed. This data is used in Step 5. If this option is not specified, the T2PDataUtil script defaults to the following locations: (UNIX) <i>FA_ORACLE_HOME</i> /applications/lcm/t2p/movedata (Windows) <i>FA_ORACLE_HOME</i> \applications\lcm\t2p\movedata
-command	The command to execute. Enter <i>apply_move_data</i> .
-arh	The Oracle Fusion Applications Repository Creation Utility (Applications RCU) location.

- If the partition feature is activated in the source environment and the partition configuration is not using preexisting tablespaces owned by the *SEARCHSYS* schema, then you must create the same tablespace in the Oracle Secure Enterprise Search (Oracle SES) database of the target environment before invoking *pasteConfig*.

See the "Partitioning for Parallel Query" section in the *Oracle Secure Enterprise Search Administration API Guide*.

- If SSL is required for communication with the Managed Servers, create the identity keystores and trust keystores for the middle tier Managed Servers and Node Manager.

See the "Enabling Node Manager to Run in SSL Mode for CRMHOST2" section in the *Oracle Fusion Applications Customer Relationship Management Enterprise Deployment Guide*.

- Set up the Oracle Identity Management environment, if it is not already set up.

All the required Oracle Identity Management components should be integrated and working together. The components required in the Oracle Identity Management domain are Oracle Internet Directory, Oracle Virtual Directory, Oracle Access Manager, Oracle Identity Manager, and optionally Oracle Fusion Applications. See the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Identity Management (Oracle Fusion Applications Edition)*.

17.4 Moving the Oracle Identity Management Domain Component Artifacts

Prior to moving Oracle Fusion Applications components across environments, you must move the Oracle Fusion Applications related artifacts for Oracle Identity Management from the source domain to the target domain. The artifacts created by Oracle Fusion Applications provisioning in the source domain are distributed across the following Oracle Identity Management products:

- Oracle Internet Directory (used as the identity store and policy store)
- Oracle Virtual Directory (used as the identity store)
- Oracle Access Manager
- Oracle Identity Manager
- Oracle Identity Federation

Movement of the Oracle Identity Management artifacts assumes that a source Oracle Identity Management environment has been provisioned for Oracle Fusion Applications and that a target Oracle Identity Management environment is configured for Oracle Fusion Applications. Move to the target environment any customizations that were applied to the source environment on an as needed basis.

To move the artifacts of Oracle Identity Management domain components for Oracle Fusion Applications from a source environment to a target environment, complete the following procedures:

1. Move the security artifacts in the identity store. See [Section 17.4.1](#).
2. Move the policy and credential store by seeding the Oracle ADF credentials. See [Section 17.4.2](#).
3. Move the data security policies. See [Section 17.4.3](#).
4. Modify the Credential Store Framework (CSF) entries in the target environment. See [Section 17.4.4](#).
5. Move Oracle Access Manager artifacts. See [Section 17.4.5](#).
6. Configure Oracle Identity Manager for Oracle Fusion Applications. See [Section 17.4.6](#).
7. Move Oracle Identity Federation artifacts. See [Section 17.4.7](#).

17.4.1 Moving the Security Artifacts in the Identity Store

Movement of the security artifacts in the identity store from the source to target environment is based on the following assumptions:

- The source and target Lightweight Directory Access Protocol (LDAP) instances are based on the same product, such as Oracle Internet Directory, Active Directory, and so on.
- The schemas related to LDAP in the target environment are compatible with that of the source environment. Using the same version of the directory server guarantees schema compatibility between the source and destination directory instances.
- Only the system users, roles, and groups are moved to the target LDAP server. The users in the Enterprise Users container in the source environment are not moved. The temporary enterprise groups that are created are also not moved.

To move the security artifacts in the identity store from the source to target environment, follow the procedures in the following tasks and execute them in the Oracle Identity Management domain:

- [Task 1, "Create the System Identities"](#)
- [Task 2, "Move the App ID Users"](#)
- [Task 3, "Move the App ID Groups"](#)
- [Task 4, "Move the Enterprise Roles"](#)
- [Task 5, "Move the Roles"](#)

Task 1 Create the System Identities

Create all the required system identities (`systemid`) at the target LDAP, if they do not exist. The minimal list of users required for configuring the Oracle Fusion Applications domain include the following users:

- Oracle WebLogic Server administration user

Note: The user ID of the Oracle WebLogic Server administration user in the target environment must be the same as the user ID of the Oracle WebLogic Server administration user in the source environment. However, the passwords can be different.

- Oracle Access Manager administration user
- Oracle Identity Manager administration user and administration group
- Other users for configuring the authentication providers in Oracle WebLogic Server

If the users were created using `idmConfigTool` in the source environment, you can use the tool with the same parameters in the target environment. For information, see the "Integrating Oracle Access Manager with Oracle Identity Manager by Using `idmConfigTool`" section in the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Identity Management (Oracle Fusion Applications Edition)*.

Running the command creates the following complete set of artifacts:

- Creates the following groups:
 - Group with read privileges to the Users Container
(`orclFAUserReadPrivilegeGroup`)
 - Group with read and write privileges to the Users Container
(`orclFAUserWritePrivilegeGroup`)
 - Group with read privileges to the Groups Container
(`orclFAGroupReadPrivilegeGroup`)
 - Group with write privileges to the Groups Container
(`orclFAGroupWritePrivilegeGroup`)
 - Group with write privileges to a partial set of attributes
(`orclFAUserWritePrefsPrivilegeGroup`)

It also provides the access control information (ACI) to the Users and Groups Container in Oracle Internet Directory.

Note: If the target identity store already exists, then the users and groups containers might be protected by Access Control Policies. To provide privileges to the above groups, you need to manually modify the policies. For information, see the "Managing Directory Access Control" chapter in the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.

- Creates the user specified by the `IDSTORE_READONLYUSER` parameter (this user is configured as the proxy user in the Authenticator configuration of Oracle WebLogic Server, so this user should be the same as the user in the source environment), and assigns it to the `orclFAWritePrefsPrivilegeGroup` and `orclFAGroupReadPrivilegeGroup` groups.
- Creates the user specified by the `IDSTORE_READWRITEUSER` parameter and assigns it to the `orclFAUserWritePrivilegeGroup` and `orclFAGroupWritePrivilegeGroup` groups.

- Creates the user specified by `SUPER_USER` (this is the Oracle WebLogic Server administration user, and needs to be the same as the user configured in the source environment), and assigns it to the `orclFAUserReadPrivilegeGroup` group.
- Creates the System ID container.
- Creates Oracle Access Manager users.
- Creates the Oracle Identity Management user and sets the corresponding ACIs.

Task 2 Move the App ID Users

App IDs are stored under `cn=appidUsers, cn=users, dc=mycompany, dc=com` in the Users Container. To move these entries, use the following Lightweight Directory Interchange Format (LDIF) export and import procedure:

1. Dump the contents of the source App ID container to a file.

```
ldapsearch -h oid_hostname -p ldap_port -D "binddn" -w password -b
'cn=appidUsers,cn=users,dc=mycompany,dc=com' -L -s sub 'objectclass=*' cn sn
uid givenname displayname objectclass >& appidUsersDump.ldif
```

2. Search for the naming context and replace it with the new naming context.

```
perl -pi -e 's/oldnamingcontext/newnamingcontext/g' appidUsersDump.ldif
```

3. Upload the LDIF file to the target environment.

```
ldapadd -h oid_hostname -p ldap_port -D "binddn" -w password -a -c -f
appidUsersDump.ldif
```

Task 3 Move the App ID Groups

App ID groups are stored under `cn=appidGroups, cn=groups, dc=mycompany, dc=com` in the Groups Container.

To move these entries, use the following LDIF export and import procedure:

1. Dump the source entries to a file.

```
ldapsearch -h oid_hostname -p ldap_port -D "binddn" -w password -b
'cn=appidGroups,cn=users,dc=mycompany,dc=com' -L -s sub 'objectclass=*' "*"
orclguid >& appidGroupsDump.ldif
```

2. Search for the naming context and replace it with the new naming context.

```
perl -pi -e 's/oldnamingcontext/newnamingcontext/g' appidGroupsDump.ldif
```

3. Upload the LDIF file to the target environment.

```
ldapmodify -h oid_hostname -p ldap_port -D "binddn" -w password -a -f
appidGroupsDump.ldif
```

Task 4 Move the Enterprise Roles

Enterprise roles are stored under `cn=fusionGroups, cn=groups, dc=mycompany, dc=com` in the LDAP directory. The enterprise roles can be associated with App ID users, App ID groups, and other enterprise users. Since the enterprise users in the source environment are not moved from the source directory to the target directory, the memberships are not moved between the source and target directories. To move these entries, use the following LDIF export, clean up, and import procedure:

1. Dump the source entries to a file.

```
ldapsearch -h oid_hostname -p ldap_port -D "binddn" -w password -b
'cn=fusionGroups,cn=users,dc=mycompany,dc=com' -L -s sub 'objectclas=*' dn cn
```

```
objectclass orclguid >fusionGroupsDump.ldif
```

2. Search for the naming context and replace it with the new naming context.

```
perl -pi -e 's/oldnamingcontext/newnamingcontext/g' fusiongroupsDump.ldif
```

3. Upload the LDIF file to the target environment.

```
ldapadd -h oid_hostname -p ldap_port -D "binddn" -w password -a -c -f fusiongroupsDump.ldif
```

Task 5 Move the Roles

Create the Administrators, Monitors, and Operators roles at the target identity store, then assign the required users to these roles. If these groups already exist, then you can ignore the error message.

1. Move the Administrators group from the source environment.

```
ldapsearch -h oid_hostname -p ldap_port -D "binddn" -w password -b 'cn=Administrators,cn=groups,dc=mycompany,dc=com' >& administrators.ldif
```

2. Remove any source environment users that should not be in the group.

```
perl -pi -e 's/oldnamingcontext/newnamingcontext/g' administrators.ldif
```

3. Upload the LDIF file to the target environment.

```
ldapadd -h oid_hostname -p ldap_port -D "binddn" -w password -a -c -f administrators.ldif
```

4. Move the Monitors group from the source environment.

```
ldapsearch -h oid_hostname -p ldap_port -D "binddn" -w password -b 'cn=Monitors,cn=groups,dc=mycompany,dc=com' >& monitors.ldif
```

5. Remove any source environment users that should not be in the group.

```
perl -pi -e 's/oldnamingcontext/newnamingcontext/g' monitors.ldif
```

6. Upload the LDIF file to the target environment.

```
ldapadd -h oid_hostname -p ldap_port -D "binddn" -w password -a -c -f monitors.ldif
```

7. Move the Operators group from the source environment.

```
ldapsearch -h oid_hostname -p ldap_port -D "binddn" -w password -b 'cn=Operators,cn=groups,dc=mycompany,dc=com' >& operators.ldif
```

8. Remove any source environment users that should not be in the group.

```
perl -pi -e 's/oldnamingcontext/newnamingcontext/g' operators.ldif
```

9. Upload the LDIF file to the target environment.

```
ldapadd -h oid_hostname -p ldap_port -D "binddn" -w password -a -c -f operators.ldif
```

17.4.2 Moving the Policy and Credential Store by Seeding the Oracle ADF Credentials

There are two policy domains created in the Oracle Fusion Applications deployments, one for Oracle Identity Management domain components and one for Oracle Fusion Applications components. Movement of the policy store from source to target involves moving contents of both policy domains from the source to target environment. It is assumed that the Oracle Identity Management domain related entities in the policy and credential store are moved as part of the Oracle Identity Management domain move.

To move the policy store that is related to Oracle Fusion Applications from the source environment to the target environment, follow the procedures in the following tasks:

- [Task 1, "Move the Policy and Credential Store Artifacts"](#)
- [Task 2, "Seed Oracle ADF Credentials"](#)

Task 1 Move the Policy and Credential Store Artifacts

To move the policy and credential store artifacts, you must first identify the policy domain created for Oracle Fusion applications. The policy domain is an LDAP container that is seeded by the Oracle Fusion Applications provisioning in the source environment. It is very important that the policy store container name is the same in both the source and target environments. Then, move the entire content of the policy store from the source environment to the target environment. For information, see the "Migrating Large Volume Policy and Credential Stores" section in *Oracle Fusion Middleware Application Security Guide*.

Task 2 Seed Oracle ADF Credentials

Because there is one map for all the domains, you only need to seed the Oracle ADF credentials from one domain. The credentials are then automatically propagated to the other domains. For example, you perform this procedure from the `CommonDomain` domain.

To seed the Oracle ADF credentials:

1. At an operating system command prompt, navigate to the location of the WLST script:

```
(UNIX) FA_MW_HOME/oracle_common/common/bin
(Windows) FA_MW_HOME\oracle_common\common\bin
```

FA_MW_HOME is named `fusionapps` and located under the *APPLICATIONS_BASE*.

2. Run the `wlst` script:

```
(UNIX) wlst.sh
(Windows) wlst.cmd
```

3. Connect to Oracle WebLogic Server with the following WLST command:

```
connect([options])
```

See the "connect" section in the *Oracle Fusion Middleware WebLogic Scripting Tool Command Reference*.

4. Obtain the password for a credential in the domain credential store with the following WLST command:

```
listCred(map="oracle.apps.security", key="appidname-KEY")
```

where *appidname* is one of the following values:

```
FUSION_APPS_PRJ_ADF_APPID
FUSION_APPS_ATK_BI_APPID
FUSION_APPS_CRM_BI_APPID
FUSION_APPS_FIN_BI_APPID
FUSION_APPS_HCM_BI_APPID
FUSION_APPS_IC_BI_APPID
FUSION_APPS_PRC_BI_APPID
FUSION_APPS_PRJ_BI_APPID
FUSION_APPS_SCM_BI_APPID
```

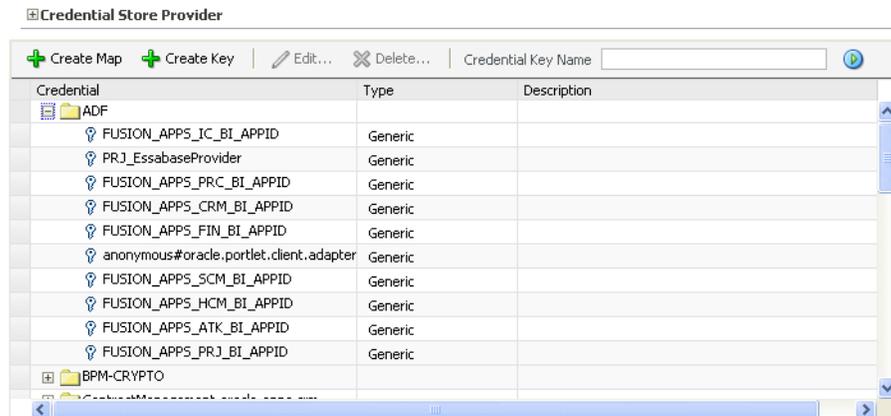
See the "listCred" section in the *Oracle Fusion Middleware WebLogic Scripting Tool Command Reference*.

These credentials are automatically populated with values from the source environment in Fusion Applications Control. In Step 5 you will modify these values to reflect the target environment.

5. Update the credentials with Fusion Applications Control using the same map and key values as used with the `listCred` command in Step 4:
 - a. From the navigation pane, expand the farm, **WebLogic Domain**, and select the domain.
 - b. From the **WebLogic Domain** menu, choose **Security > Credentials**.
The Credentials page displays.
 - c. In the **Credential Store Provider** section, expand **ADF**.

Credentials

A credential store is the repository of security data that certify the authority of entities used by Java 2, J2EE, and ADF applications. Applications can use the Credential Store, a single, consolidated service provider to store and manage their credentials securely.



These credentials are automatically populated with values from the source environment.

- d. Select a credential and click **Edit**.

The Edit Key dialog displays.

- e. View and modify the values for a key in the following **Property Name** fields:

- **un**: This value is set to the `appidname` for the key. You should not have to edit this value. This property is not available for the `PRJ_EssbaseProvider` key.

- **pw**: Enter the password provided by the `listCred` command from Step 4. This property is not available for the `PRJ_EssbaseProvider` key.

- **password:** Enter the password provided by the listCred command from Step 4 for FUSION_APPS_PRJ_ADF_APPID.

- f. Click **OK**.
- g. Perform sub-steps d through f for each of the remaining keys.

17.4.3 Moving Data Security Policies

Data security policies created in the simplified Data Security management UI or in Oracle Authorization Policy Manager should be migrated using Functional Setup Manager. The business object name associated for data security is Application Data Security Policy.

For more information on moving data security policies, see the following documentation:

- [Section 17.8](#) for further information on how to move data Oracle Fusion Applications security policies using Functional Setup Manager export and import functions.
- "Moving Common Reference Objects" section in the "Importing and Exporting Setup Data" chapter in the *Oracle Fusion Applications Common Implementation Guide* to understand data security policy migration
- "Importing and Exporting Setup Data" chapter in the *Oracle Fusion Applications Information Technology Management, Implement Applications Guide* for information on moving setup data

17.4.4 Modifying the CSF Entries in the Target Environment

You must review the credentials in the credential store and update them for the target environment in the target environment. The credentials are described in [Table 17-2](#).

Table 17-2 CSF Entries to Modify

Map	Key	Description
oracle.apps.security	FUSION_APPS_ECSF_SES_ADMIN-KEY	This entry stores the SEARCHSYS schema password for ECSF and Oracle SES for Oracle SES 11.2.1.
oracle.apps.security	FUSION-DB-KEY	Oracle Fusion database credentials
oracle.bi.enterprise	scheduler.schema	This entry stores the schema password for Oracle BI Scheduler schema.
oracle.patching	FUSION_APPS_CRM_MDS_SOA_SCHEMA-KEY	User credentials for the patch MDS SOA schema
oracle.patching	FUSION_APPS_FIN_MDS_SOA_SCHEMA-KEY	User credentials for the patch MDS SOA schema
oracle.patching	FUSION_APPS_HCM_MDS_SOA_SCHEMA-KEY	User credentials for the patch MDS SOA schema
oracle.patching	FUSION_APPS_MDS_ESS-KEY	User credentials for the patch MDS Oracle Enterprise Scheduler schema
oracle.patching	FUSION_APPS_MDS-KEY	User credentials for the patch MDS schema

Table 17-2 (Cont.) CSF Entries to Modify

Map	Key	Description
oracle.patching	FUSION_APPS_MDS-SPACES	User credentials for the patch MDS Oracle WebCenter Portal Spaces schema
oracle.patching	FUSION_APPS_ODI_SCHEMA-KEY	Oracle Data Integrator schema credentials
oracle.patching	FUSION_APPS_OIC_MDS_SOA_SCHEMA-KEY	User credentials for the patch MDS SOA schema
oracle.patching	FUSION_APPS_PATCH_FUSION_DYNAMIC_SCHEMA-KEY	The credential for the FUSION_DYNAMIC schema in the Oracle Fusion transactional database
oracle.patching	FUSION_APPS_PATCH_FUSION_RUNTIME_SCHEMA-KEY	The credential for the FUSION_RUNTIME schema in the Oracle Fusion transactional database
oracle.patching	FUSION_APPS_PATCH_FUSION_SCHEMA-KEY	User credentials for patch
oracle.patching	FUSION_APPS_PATCH_ID_STORE-KEY	User credentials for patching to connect to the Identity Store. Currently, this credential is a read-write user supplied to provisioning.
oracle.patching	FUSION_APPS_PATCH_POLICY_STORE-KEY	User credentials for patching to connect the Policy Store
oracle.patching	FUSION_APPS_PATCH_WLS_ADMIN	User credentials for Oracle WebLogic Server administrator
oracle.patching	FUSION_APPS_PATCH_WLS_NODE_MANAGER-KEY	User credentials for Oracle WebLogic Server Node Manager
oracle.patching	FUSION_APPS_PRC_MDS_SOA_SCHEMA-KEY	User credentials for the patch MDS SOA schema
oracle.patching	FUSION_APPS_PRJ_MDS_SOA_SCHEMA-KEY	User credentials for the patch MDS SOA schema
oracle.patching	FUSION_APPS_SCM_MDS_SOA_SCHEMA-KEY	User credentials for the patch MDS SOA schema
oracle.patching	FUSION_APPS_SETUP_MDS_SOA_SCHEMA-KEY	User credentials for the patch MDS SOA schema
oracle.search	SEARCH_DATABASE	<p>Oracle SES data source credentials.</p> <p>Provide user name and password.</p> <p>Use the database connection string of the target environment for the user name in this format:</p> <p><i>host:port:database_service_name</i></p> <p>The password is the dummy value SEARCH.</p>
oracle.t2p	FUSION_APPS_ODI_SCHEMA-KEY	Oracle Data Integrator schema credentials

Table 17-2 (Cont.) CSF Entries to Modify

Map	Key	Description
oracle.wsm.security	basic.credentials	Oracle WebLogic Server user credentials key. This credential corresponds to the user name and password of bootstrap credential, FUSION_APPS_PROV_PATCH_APPID, which comes from the provisioning plan.
oracle.wsm.security	enc-csf-key	User credential for specifying which keystore alias to use for encrypting the message. Both enc-csf-key and sign-csf-key aliases will be the same.
oracle.wsm.security	FUSION_APPS_HCM_HR2HR_APPLOGIN-KEY	The user credentials for the PeopleSoft Enterprise Application web service SOAP interaction. The user credentials corresponds to the PeopleSoft Enterprise Application application user with access privileges to invoke the PeopleSoft Application web service FT_FUSION_TO_PSFT_MSG.1.wsdl.
oracle.wsm.security	FUSION_APPS_PRC_SUPP_SYNC_SVC-KEY	The credentials of the external service for the supplier synchronization service (integration)
oracle.wsm.security	FUSION_APPS_PRJ_P6INT_ADMIN-KEY	The Primavera P6 administration super user credentials designated for use with the integration between Oracle Fusion Projects Portfolio Management and Oracle Primavera P6 Enterprise Project Portfolio Management (P6 EPPM).
oracle.wsm.security	keystore-csf-key	User credential is for specifying the password of the default_keystore.jks keystore.
oracle.wsm.security	sign-csf-key	User credential is for specifying which keystore alias to use for signing the message. Both enc-csf-key and sign-csf-key aliases will be the same.

To modify the CSF entries:

1. For either UNIX or Windows, at an operating system command prompt, navigate to the following directories on the target:

(UNIX) `FA_ORACLE_HOME/applications/lcm/t2p/bin`

(Windows) `FA_ORACLE_HOME\applications\lcm\t2p\bin`

`FA_ORACLE_HOME` is located under the fusionapps Oracle Fusion Middleware home directory.

2. Use the following command to update the CSF entries in the target environment:

```
java
-classpath T2PAutomation.jar:cloningclient.jar:obfuscatepassword.jar:ojmisc.jar
oracle.t2p.automation.T2PAutomationMain
-task MoveIDM -props -props property_file -workdir working_dir
```

[Table 17–3](#) describes the options for the `MoveIDM` command.

Table 17–3 Options for the `MoveIDM` Command

Options	Description
<code>-classpath</code>	The dependent library files containing classes required to run the command. Use the following value for the <code>MoveIDM</code> command: T2PAutomation.jar:cloningclient.jar:obfuscatepassword.jar:ojmisc.jar oracle.t2p.automation.T2PAutomationMain
<code>-task</code>	The command to execute. Enter <code>MoveIDM</code> .

Table 17-3 (Cont.) Options for the MoveIDM Command

Options	Description
-props	<p>The absolute path to the properties file name. The properties file must have following properties:</p> <pre>#Mandatory properties. Connection info for IDM domain. AdminUser=FAAdmin AdminPassword=fusionfal AdminURL=t3://hostname:port #Script location, needs to be the absolute path to scripts folder, inclusive. scripts-folder=/scratch/aime/work/t2p/scripts #Optional if set as environment variable. WLST_ PROPERTIES=-Dweblogic.security.SSL.trustedCAKeyStore=\$APPL TOP/fusionapps/wlserver_10.3/server/lib/fusion_trust.jks #Optional, if not specified defaults to PLAIN NMConnectionType=PLAIN</pre> <p>Change the sample value for your target environment.</p> <p>The MoveIDM task expects the following properties files in the scripts directory:</p> <ul style="list-style-type: none"> ■ csf_maps.properties: <pre>#Key and map pairs for CSF entries FUSION_APPS_ECSF_SES_ADMIN-KEY=oracle.apps.security scheduler.schema=oracle.bi.enterprise FUSION-DB-KEY=oracle.apps.security FUSION_APPS_CRM_MDS_SOA_SCHEMA-KEY=oracle.patching FUSION_APPS_FIN_MDS_SOA_SCHEMA-KEY=oracle.patching ...</pre> ■ csf_passwords.properties: <pre>#Replace the user and password for each of the key below. Make sure that they are ':' separated FUSION_APPS_ECSF_SES_ADMIN-KEY=user:password scheduler.schema=user:password FUSION-DB-KEY=user:password FUSION_APPS_CRM_MDS_SOA_SCHEMA-KEY=user:password FUSION_APPS_FIN_MDS_SOA_SCHEMA-KEY=user:password FUSION_APPS_HCM_MDS_SOA_SCHEMA-KEY=user:password</pre> <p>Replace the user and password values for all the keys.</p>
-workdir	<p>The directory used for log files and any other temporary files generated by the automation framework. The log files are located at <i>working_dir</i>\logs (UNIX) or <i>working_dir</i>\logs folder. This command is not needed if the T2P_DIR environment variable is set.</p>

17.4.5 Moving Oracle Access Manager Artifacts

To move Oracle Access Manager, follow the steps in task "Move Oracle Access Manager 11g to a New Target Environment" in the *Oracle Fusion Middleware Administrator's Guide* to move Oracle Access Manager artifacts.

Oracle Access Manager artifacts are moved from a test environment to the production environment when you move the Webgate installation and change the Webgate

identifier to reflect the target end points. You perform this task prior to moving the configuration files to the production environment, as described in [Section 17.6.3](#).

17.4.6 Configure Oracle Identity Manager for Oracle Fusion Applications

To configure Oracle Identity Manager for Oracle Fusion Applications, follow the procedures in the following tasks:

- [Task 1, "Move the Callback Configuration File"](#)
- [Task 2, "Seed the Oracle Fusion Applications Artifacts"](#)
- [Task 3, "Reconcile Records from the Identity Store"](#)

Note: The Administration Server and Oracle Identity Manager and SOA Managed Servers must be running when you perform the procedures.

Movement of the Governance Risk and Controls client configuration is performed when you move the configurations in the Oracle Fusion Applications environment. See [Section 17.6.8.3](#).

Task 1 Move the Callback Configuration File

The `CallbackConfiguration.xml` file is an MDS document that contains information on the web services that need to be invoked for completion and postprocessing callbacks. For information, see the "MDS Utilities and User Modifiable Metadata Files" chapter in the *Oracle Fusion Middleware Developer's Guide for Oracle Identity Manager*.

To move `CallbackConfiguration.xml` from the source to the target environment:

1. In the source environment, use the `weblogicExportMetadata.sh` utility, located in `IDENTITY_MANAGER_HOME/bin/`, to export `CallbackConfiguration.xml`, located in `/metadata/iam-features-callbacks/sample_data/`.
2. Edit the `CallbackConfiguration.xml` file to specify the host and port of the target environment in the web service URLs.
3. In the target environment, import the `CallbackConfiguration.xml` file into MDS using the `weblogicImportMetadata.sh` utility located in `IDENTITY_MANAGER_HOME/bin/`.

Task 2 Seed the Oracle Fusion Applications Artifacts

To seed the Oracle Fusion Applications artifacts:

Note: Use JConsole to invoke the managed beans (MBean), and use the Oracle WebLogic Server domain administrator credentials for performing all MBean operations.

For more information, see the *Oracle Fusion Middleware Oracle WebLogic Server MBean Reference*.

1. Invoke the `MLSLanguageSeedingMBean.seedAllLanguages` method to seed all the multi-language support (MLS) and multi-representation (MR) languages in the target system. This MBean method does not take any parameters.
2. Invoke the `APPIDSeedingMBean.seedFusionAPPID` method once for each App ID to create the App IDs for Oracle Fusion Applications in Oracle Identity Manager.

(You can get the App IDs from the source environment where they are already seeded.) This MBean method takes a single parameter that is the actual App ID itself and will seed this App ID in the Oracle Identity Manager data store. For information about App IDs, see [Table 3-4](#).

Task 3 Reconcile Records from the Identity Store

Use Oracle Identity Manager to reconcile records from the identity store.

To reconcile records from the identity store:

1. Sign in to Oracle Identity Manager.
2. Click **Advanced**, then the **System Management** tab, and then **Scheduler**. Alternatively, you can click the **Search Scheduled Jobs** link on the Welcome screen.
3. Run each of the following scheduled tasks by selecting it and clicking **Run Now**. You must run one job at a time in the following sequence and wait for each job to complete successfully before running the next one:
 - a. **Fusion Applications Role Category Seeding**
 - b. **LDAP Role Create and Update Full Reconciliation**
 - c. **LDAP Role Hierarchy Full Reconciliation**
 - d. **LDAP User Create and Update Full Reconciliation**
 - e. **LDAP Role Membership Full Reconciliation**

17.4.7 Moving Oracle Identity Federation Artifacts

Oracle Identity Federation artifacts are moved from a source environment to target environment when you move Oracle Identity Management components across environments.

Oracle Identity Federation has already been moved from the source environment to target environment. For information, see the "Moving Identity Management to an Existing Production Environment" section in the *Oracle Fusion Middleware Administrator's Guide*.

17.5 Moving Binary Files to a Production Target Environment

Moving installed binaries and patches in the Middleware home of an Oracle Fusion Applications source environment to a production target environment requires the use of a set of scripts. For more information, see the "Movement Scripts" section in the *Oracle Fusion Middleware Administrator's Guide*.

If you have a non-distributed, QA environment, whereby the Administration Servers of all the domains and all the system component instances run on the same host, you can automate moving the binary files with a different set of scripts than those described in this section. For more information, see [Section 17.2](#) and [Section 17.7](#).

Notes:

- If you want to create obfuscated password files for use with these scripts, use the following scripts:

```
(UNIX) FA_MW_HOME/oracle_common/bin/obfuscatePassword.sh
```

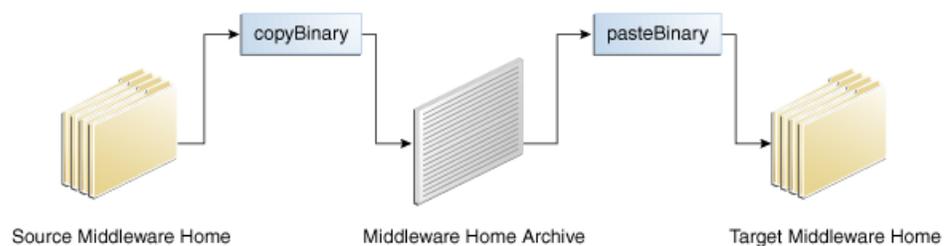
```
(Windows) FA_MW_HOME\oracle_common\bin\obfuscatePassword.cmd
```

These interactive scripts prompt you to enter your password and the location of the password file to be written.

- The Java version should be at least 1.6.04.

The `copyBinary` and `pasteBinary` scripts allow you to create a copy of the source Middleware home into an archive file, then apply the archive to the target environment, as shown in [Figure 17-1](#).

Figure 17-1 Movement of Binaries



Since the binaries are in shared storage and are shared by all Oracle Fusion Applications domains, `copyBinary` needs to be performed only once from the primordial machine and `pasteBinary` only once on the target primordial machine.

To move binary files in the Oracle Fusion Applications environment, you must complete the following steps:

1. Create the binary archives by running the `copyBinary` script for each Middleware home in the source environment. See [Section 17.5.1](#).
2. Apply the copied binary files from the source environment to the target environment by running the `pasteBinary` script in the target environment. See [Section 17.5.2](#).
3. Manually move the Oracle Database installation from the source environment to the target environment. See [Section 17.5.3](#).

17.5.1 Creating Binary Archives

Run the `copyBinary` script for each Middleware home in the source environment. This copies the binary files of each Middleware home, including all of its Oracle homes and its Oracle WebLogic Server home, into an archive file.

The `copyBinary` script is available in the following `fusionapps` Middleware directories:

```
(UNIX) FA_MW_HOME/oracle_common/bin
```

```
(Windows) FA_MW_HOME\oracle_common\bin
```

Note: The Oracle homes in the Middleware home must share the same platform architecture (in terms of number of bits). The `copyBinary` script does not support a mix of 32-bit and 64-bit Oracle homes.

When you run the script, you must specify a Java home of the same number of bits as the Oracle homes. For example, for 64-bit Oracle homes, you must specify a 64-bit Java home.

Before creating the binary archive of a Middleware home, make sure that the Oracle WebLogic Server product directories are installed in the Middleware home. Oracle WebLogic Server installed outside of Middleware home is not supported.

In a Windows environment, before creating an archive of a Middleware home, make sure that no Java or Oracle WebLogic Server processes are running from that Middleware home.

Use the following syntax to run the `copyBinary` script:

```
copyBinary -javaHome path_of_jdk
           -archiveLoc archive_location
           -sourceMWHomeLoc MW_HOME
           [-invPtrLoc Oracle_InventoryLocation]
           [-logDirLoc log_dir_path]
           [-silent {true | false}]
           [-ignoreDiskWarning {true | false}]
```

The following example shows how to create an archive of a Middleware home on UNIX:

```
copyBinary.sh -javaHome USER_HOME/jrockit_160_04
              -archiveLoc /FIN_T2P/FIN_clone1.jar
              -sourceMWHomeLoc /net/mount1/appbase/fusionapps
              -invPtrLoc /etc/oraInst.loc
              -logDirLoc /FIN_T2P/logs
              -silent true
              -ignoreDiskWarning true
```

`/net/mount1/appbase` represents the top level applications base directory for binaries.

[Table 17-4](#) describes the options for the `copyBinary` script.

Table 17–4 Options for the copyBinary Script

Option	Shortcut	Description	Mandatory or Optional?
-javaHome	None	The absolute path of the Java Developer's Kit.	Mandatory
-d64	None	If the operating system is 64-bit UNIX, pass the -d64 option to the scripts in the command line. To set the runtime property, you can specify the -d64 option in the T2P_JAVA_OPTIONS environment variable. For example: <pre>setenv T2P_JAVA_OPTIONS "-d64 -Djava.io.tmpdir=/home/t2p/temp"</pre>	Optional
-archiveLoc	-al	The absolute path of the archive location. Use this option to specify the location of the archive file to be created with the copyBinary script. The archive location must not already exist, but its parent directory must exist and have write permission. Ensure that the archive location is not within the Middleware home structure.	Mandatory
-invPtrLoc	-ipl	The absolute path to the Oracle Inventory pointer. Use this option if the inventory location is not in the default location, so that the operation can read the Oracle homes present in the inventory. You must have write permission to the inventory location. The default location is /etc/oraInst.loc on UNIX. This option is not available on Windows.	Available on UNIX only. Optional, if the inventory is in the default location. Otherwise, it is mandatory.
-logDirLoc	-ldl	The location of an existing directory. A new log file is created in the directory.	Optional
-silent	None	Specifies whether the operation operates silently. That is, it does not prompt for confirmation. The default is that the operation prompts for confirmation. To continue, you must type yes, which is not case sensitive. Typing anything other than yes causes the script to return an error. To specify that it not prompt for confirmation, specify this option with the value of true.	Optional
-sourceMWHomeLoc	-smw	The absolute path of the Middleware home to be archived. You can only specify one Middleware home.	Mandatory
-ignoreDiskWarning	-idw	Specifies whether the operation ignores a warning that there is not enough free space available. The default is false. You may need to use this flag if the target is NFS mounted or is on a different file system, such as Data ONTAP	Optional

17.5.2 Applying the Binary Files to the Target Environment

Run the `pasteBinary` script to apply the archive to the target environment. This pastes the binary files of the source Middleware home in the target environment. You must run the `pasteBinary` script for each Middleware home in the target environment.

The `pasteBinary` script is available from the following `fusionapps` Middleware directories:

(UNIX) `FA_MW_HOME/oracle_common/bin`
 (Windows) `FA_MW_HOME\oracle_common\bin`

Note: If you are applying the archive of a Middleware home on a host that does not yet have Oracle Fusion Middleware installed, then:

- The host must have JDK 1.6.04 or higher installed. In addition, ensure that the `PATH`, `CLASSPATH`, and `JAVA_HOME` environment variables point to the JDK.
- Copy the `pasteBinary` script from the following location on the source host to the target host:

(UNIX) `FA_MW_HOME/oracle_common/bin/pasteBinary.sh`
 (Windows) `FA_MW_HOME\oracle_common\bin\pasteBinary.cmd`

- Copy the following file from the following location on the source host to the target host:

(UNIX) `FA_MW_HOME/oracle_common/jlib/cloningclient.jar`
 (Windows) `FA_MW_HOME\oracle_common\jlib\cloningclient.jar`

If you run the `pasteBinary` script from a different location than `FA_MW_HOME/oracle_common/bin`, then you must copy the `cloningclient.jar` file to the same directory.

- In a UNIX environment, make sure to properly set the `inventory_loc` and `inst_group` properties in the inventory file. For example:

```
inventory_loc=/FIN_T2P/oraInventory
inst_group=dba
```

The inventory file location itself is specified through the `invPtrLoc` option of the `copyBinary` command. If the `invPtrLoc` option is not specified, the inventory pointer file is assumed to be in its default location of `/etc/oraInst.loc`.

- Make sure that the files have execute permission.
 - The Middleware home must be created on the shared storage of the target environment.
 - The `pasteBinary` script must be run from the primordial machine.
 - Even though the absolute path need not be the same as the one in the source environment, the directory structure relative to `APPLICATIONS_BASE` should be the same.
-

Use the following syntax to run the `pasteBinary` script:

```
pasteBinary -javaHome path_of_jdk
             -archiveLoc archive_location
             -targetMWHomeLoc target_MW_Home_location
             [-executeSysPrereqs {true | false}]
             [-invPtrLoc Oracle_InventoryLocation]
             [-logDirLoc log_dir_path]
             [-silent {true | false}]
             [-ignoreDiskWarning {true | false}]
```

The following example shows how to apply the archive to the directory `/net/mount1/appbase/fusionapps` on UNIX:

```
pasteBinary.sh -javaHome USER_HOME/jrockit_160_04
               -archiveLoc /FIN_T2P/FIN_clone1.jar
               -targetMWHomeLoc /net/mount1/appbase/fusionapps
               -logDirLoc /FIN_T2P/logs/
               -silent true
               -ignoreDiskWarning true
```

`/net/mount1/appbase` represents the top level applications base directory for binaries.

Table 17-5 describes the options for the `pasteBinary` script.

Table 17-5 Options for the `pasteBinary` Script

Option	Shortcut	Description	Mandatory or Optional?
<code>-javaHome</code>	None	The absolute path of the Java Developer's Kit.	Mandatory
<code>-d64</code>	None	If the operating system is 64-bit UNIX, pass the <code>-d64</code> option to the scripts in the command line. To set the runtime property, you can specify the <code>-d64</code> option in the <code>T2P_JAVA_OPTIONS</code> environment variable. For example: <code>setenv T2P_JAVA_OPTIONS "-d64</code> <code>-Djava.io.tmpdir=/home/t2p/temp"</code>	Optional
<code>-archiveLoc</code>	<code>-al</code>	The absolute path of the archive location. Use this option to specify the location of the archive file created with the <code>copyBinary</code> script. The location must exist. Ensure that the archive location is not within the Middleware home structure.	Mandatory
<code>-executeSysPrereqs</code>	<code>-esp</code>	Specifies whether the <code>pasteBinary</code> operation checks the prerequisites of the Oracle homes. The default is that it checks the prerequisites. To specify that it does not check the prerequisites, specify this option with the value <code>false</code> .	Optional
<code>-invPtrLoc</code>	<code>-ipl</code>	The absolute path to the Oracle Inventory pointer. Use this option if the inventory location is not in the default location, so that the operation can read the Oracle homes present in the inventory. You must have write permission to the inventory location. The default location is <code>/etc/oraInst.loc</code> on UNIX. This option is not available on Windows.	Available on UNIX only. Optional, if the inventory is in the default location. Otherwise, it is mandatory.
<code>-logDirLoc</code>	<code>-ldl</code>	The location of an existing directory. A new log file is created in the directory.	Optional

Table 17-5 (Cont.) Options for the pasteBinary Script

Option	Shortcut	Description	Mandatory or Optional?
-silent	None	Specifies whether the operation operates silently. That is, it does not prompt for confirmation. The default is that the operation prompts for confirmation. To continue, you must type <i>yes</i> , which is not case sensitive. Typing anything other than <i>yes</i> causes the script to return an error. To specify that it not prompt for confirmation, specify this option with the value of <i>true</i> .	Optional
-targetMWHomeLoc	-tmw	The absolute path of the target Middleware home. If a multihost target topology is being used, the path would need to be <i>/net/mount1/appbase</i> addresses in order to make it visible to other hosts. For example: <i>/net/mount1/appbase/fusionapps</i> . Make sure that the Middleware home's parent directory, for example, <i>fusionapps</i> exists. If it does not exist, the program will error out. The <i>targetMWHomeLoc</i> cannot be inside another Middleware home.	Mandatory
-ignoreDiskWarning	-idw	Specifies whether the operation ignores a warning that there is not enough free space available. The default is <i>false</i> . You may need to use this flag if the target is NFS mounted or is on a different file system, such as Data ONTAP.	Optional

17.5.3 Moving the Oracle Database Client Installation

The Oracle Database client home is in the following location:

```
(UNIX) APPLICATIONS_BASE/dbclient
(Windows) APPLICATIONS_BASE\dbclient
```

Its contents must be manually moved from the source environment to the target environment.

Perform the following steps to move the *dbclient* installation:

1. Copy or tar the */net/mount1/appbase/dbclient* directory to the target directly or from source.

/net/mount1/appbase represents the top level applications base directory for binaries.

2. Run the installer clone command, shown in the following example:

```
/net/mount1/appbase/dbclient/oui/bin/runInstaller -clone
ORACLE_BASE=/net/mount1/appbase ORACLE_HOME=/net
/net/mount1/appbase/dbclient ORACLE_HOME_NAME="dbclient" -jreLoc
/net/mount1/appbase
/fusionapps/jrockit_160_17_R28.0.0-679 -defaultHomeName
```

You will be prompted to run the *root.sh* command during the process.

3. Use an account with root access and run *net/oracle/dbclient/root.sh*.

Note: You will not be prompted for permission to copy to the local bin.

17.6 Moving Configurations to a Production Target Environment

For each Oracle Fusion application, you must move the configurations of each Oracle WebLogic Server domain, OHS instance, and Node Manager that is part of the source environment to the target environment. This configuration movement relies on the use of the configuration scripts. For information, see [Section 17.6.1](#).

If you have a non-distributed, QA environment, whereby the Administration Servers of all the domains and all the system component instances run on the same host, you can automate moving the configuration files with a different set of scripts than those described in this section. For more information, see [Section 17.2](#) and [Section 17.7](#).

Note: On Windows, before you move the configurations to the target environment, you must set the the `USER_MEM_ARG` variable. For information, see [Section 17.6.2](#).

You must also move the Webgate installation. For information, see [Section 17.6.3](#).

To move configurations, you must complete the following steps:

1. Create the configuration archives by running the `copyConfig` script for each domain, OHS instance, and Node Manager in the source environment. See [Section 17.6.4](#).
2. Extract the move plan by running the `extractMovePlan` script on the list of configuration archive locations, separated by comma. See [Section 17.6.5](#).
3. Modify the move plan to specify properties for the target environment. See [Section 17.6.6](#).
4. Apply the copied configurations from the source environment to the target environment by running the `pasteConfig` script for each Oracle WebLogic Server domain, OHS instance, and Node Manager separately, using the same move plan. See [Section 17.6.7](#).
5. Complete the component-specific configuration move by performing additional movement tasks on the Oracle Fusion Middleware components. See [Section 17.6.8](#).
6. If there are multiple machines for the domain, recreate the local domain directory using the Oracle WebLogic Server `pack` and `unpack` mechanisms. See [Section 17.6.9](#).

Notes:

- If the database is not tuned correctly, `copyConfig` and `pasteConfig` operations can result in performance issues. To avoid these performance issues, in addition to following standard database performance tuning guidelines, ensure that you have sufficient RAM allocated for your RDBMS for the import of the MDS tables. Also run statistics against the target database by executing the following procedure:

```
BEGIN
dbms_stats.gather_schema_stats(
ownname => 'FUSION_MDS',
METHOD_OPT => 'FOR ALL COLUMNS SIZE AUTO',
CASCADE => TRUE,
ESTIMATE_PERCENT => NULL);
END;
```

where `FUSION_MDS` is the MDS schema name for your Oracle Fusion Applications installation.

- If OHS is configured with Oracle Webgate, note that the `WebgateInstallDir` property and references to this path are updated in the `httpd.conf` file located in the following directories:

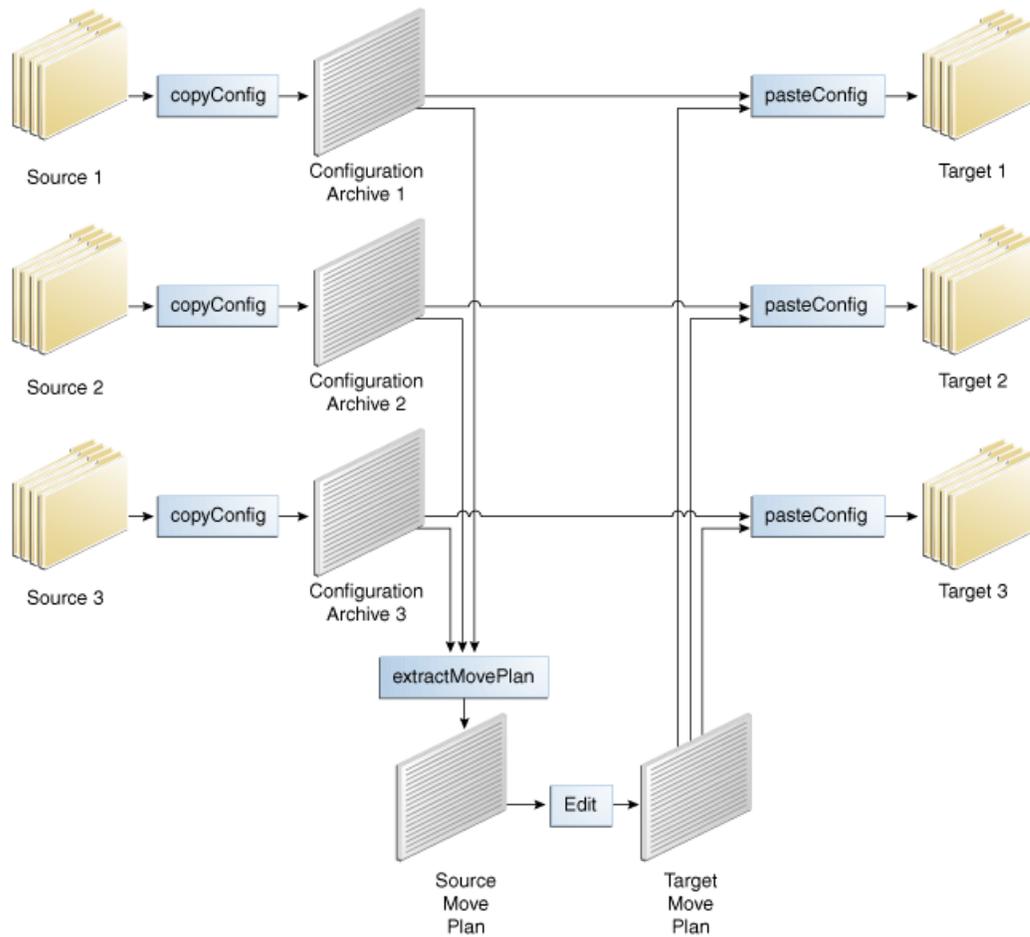
```
(UNIX)  INSTANCE_HOME/config/OHS/component_name/httpd.conf
(Windows) INSTANCE_HOME\config\OHS\component_name\httpd.conf
```

- The Java version should be at least 1.6.04.
-

17.6.1 Understanding the Configuration Scripts

Moving configurations in the Oracle Fusion Applications environment requires the use of a set of configuration scripts. These scripts are available in `FA_MW_HOME/oracle_common/bin`. For more information, see the "Movement Scripts" section in the *Oracle Fusion Middleware Administrator's Guide*.

Using the `copyConfig`, `extractMovePlan`, and `pasteConfig` scripts allows you to create a copy of the source Oracle WebLogic Server domain home or OHS instance home configuration files into an archive file, extract the configuration information from one or more configuration archive into a move plan so you can edit the properties for the target environment, and apply the copied configurations (along with the move plan) to the target environment. [Figure 17-2](#) illustrates the flow of the configuration movement from the source environment to the target environment.

Figure 17-2 Movement of Configurations

You must create a configuration archive for each source. The environment specific configuration information from each of the configurations archives are extracted to and consolidated in the source move plan that you must edit to specify the properties for the target environment. Then you must apply the edited move plan and the copied configuration archives from each source to each target.

Note: The configuration archive files should be copied to the target environment, preferably to the shared storage so they are visible to all the target domains. Similarly, the extracted move plan and the updated move plan must also be present in the shared storage.

17.6.1.1 copyConfig Script

The `copyConfig` script creates a configuration archive that contains the snapshot of the configuration of a source entity (for example, Oracle WebLogic Server domain, OHS instance, or Node Manager). The underlying components of a source entity persist their configuration information in different data stores, such as a file system, Oracle Metadata Services (MDS), Lightweight Directory Access Protocol (LDAP), database, and so on.

You must run the `copyConfig` script for each source entity in the source environment. A configuration archive is created for each source domain or component instance.

The Administration Server and all Managed Servers in the domain must be up and running when you run the script.

17.6.1.2 extractMovePlan Script

The `extractMovePlan` script extracts the environment specific configuration information from one or more configuration archives into a move plan (XML file).

The `extractMovePlan` script also extracts any J2EE component specific configuration file, such as Oracle SOA composite configuration plans, application deployment plans, adapter deployment plans, and so on.

A **move plan** contains the environment specific configuration information, such as data source definitions, host names, port numbers, and end-point URLs, from one or more configuration archives created for each source domain or component instance in the source environment. You must edit the move plan to specify the properties for the target environment.

17.6.1.3 pasteConfig Script

The `pasteConfig` script applies the copied configurations from the source environment to the target environment. Inputs for the script include the location of the configuration archive created with the `copyConfig` script for each source entity and the modified move plan. The `pasteConfig` script recreates the configuration information for the source entity in the target environment. It also merges the move plan property values for the target environment. You must run the `pasteConfig` script for each source entity in the target environment.

17.6.2 Setting the JVM System Variable on Windows

Before you run the `pasteConfig` script in Windows, you must set the `USER_MEM_ARG` variable as follows:

```
set USER_MEM_ARGS="-Xms256m -Xmx1024m -XX:CompileThreshold=8000 -XX:PermSize=128m  
-XX:MaxPermSize=1024m"
```

17.6.3 Moving the Webgate Installation

Webgate acts as the proxy for user authentication. It must be manually moved from the source environment to the target environment. To move the Webgate installation, follow the instructions below:

- UNIX. See [Section 17.6.3.1](#).
- Windows. See [Section 17.6.3.2](#).

Note: The Webgate directory must be in the following directory:

Oracle_Instance/config/OHS/ohs_component_name

Successfully moving the Webgate installation is based on the assumption that the web server and Webgate setup on the source machine is functioning correctly.

17.6.3.1 Moving the Webgate Installation on UNIX

Before you move the Webgate installation on UNIX, make sure that the file permissions of the directory to which the Webgate installation is moved are the same as those of the source.

To move the Webgate installation on UNIX:

1. Go to the Webgate installation directory on source machine. For example:

```
$ cd /net/mount1/appbase/webgate
```

2. Copy files recursively from the source machine to the target machine.

For example, if the source machine file system is mounted on the network, then the Webgate installation can be copied with the following command:

```
$ cp -rf * /net/mount1/appbase/webgate
```

If the source machine file system is not mounted on the network, then the Webgate installation can be copied with the following command:

```
$ tar cvf source_webgate.tar
```

Then you must copy the archive file to the target machine and use the following command to extract the files from the archive:

```
$ tar xvf source_webgate.tar -C target_Webgate_installation
```

3. If the Oracle Access Manager access servers to be used in the target environment are different than those used in the source environment:

Copy the artifacts generated for this Webgate to Webgates installation directory from the target Oracle Access Manager location in Oracle Identity Management. Artifacts generated should be copied as follows

- ObAccessClient.xml to *webgate_install_dir/access/oblix/lib*

- logout.html to *webgate_install_dir/access/oamsso*

If the security mode is Simple:

- password.xml to *webgate_install_dir/access/oblix/config*

- aaa_key.pem and aaa_cert.pem to *webgate_install_dir/access/oblix/config/simple*

If the security mode is Cert:

- password.xml to *webgate_install_dir/access/oblix/config*

- It is the user's responsibility to generate certificate request, get it signed by a third-party Certificate Authority, and copy it to *webgate_install_dir/access/oblix/config*

17.6.3.2 Moving the Webgate Installation on Windows

Before you move the Webgate installation on Windows, if the target machine is running Windows 64 make sure that the Microsoft Visual Studio C++ 2005 runtime libraries (MSVCR80.dll and MSVCP80.dll) are installed. You can copy the libraries from SYSVOL\windows\SysWOW64 on the source machine to the same directory on the target machine.

To move the Webgate installation on Windows:

1. Copy the *access* directory from the Webgate installation directory on source machine and paste the directory in the Webgate installation directory on the target machine.
 - a. Create an archive file by archiving the contents of the *source_webgate_install_dir\access* directory.

- b. Transfer the archive file to the target machine using File Transfer Protocol (FTP) or similar means.
 - c. On target machine, extract the archive file to the required `target_webgate_install_dir` directory.
2. If the Oracle Access Manager access servers to be used in the target environment are different than those used in the source environment:
- a. Create the Webgate profile, then modify it in the Oracle Access Manager Access System Console. For information, see the "Managing OAM 10g Webgates with OAM 11g" section in the *Oracle Fusion Middleware Administrator's Guide for Oracle Access Manager with Oracle Security Token Service*.
 - b. Create a Oracle Access Manager 10g Webgate profile using Oracle Access Manager Console. This will generate the artifacts for this Webgate under the following directories:

`FA_MW_HOME\user_projects\domains\base_domain\output\agent_ID`

For information, see the "Managing OAM 10g Webgates with OAM 11g" chapter in the *Oracle Fusion Middleware Administrator's Guide for Oracle Access Manager with Oracle Security Token Service*.

- c. Copy the artifacts generated for this Webgate to Webgates installation directory from the target Oracle Access Manager location in Oracle Identity Management. Artifacts generated should be copied as follows:

- ObAccessClient.xml to `webgate_install_dir\access\oblix\lib`

- logout.html to `webgate_install_dir/access/oamsso`

If the security mode is Simple:

- password.xml to `webgate_install_dir\access\oblix\config`

- aaa_key.pem and aaa_cert.pem to `webgate_install_dir\access\oblix\config\simple`

If the security mode is Cert:

- password.xml to `webgate_install_dir\access\oblix\config`

- It is the user's responsibility to generate certificate request and get it signed by a third-party Certificate Authority and copy to `webgate_install_dir\access\oblix\config`

17.6.4 Creating Configuration Archives

Run the `copyConfig` script for each source entity (for example, Oracle WebLogic Server domain, OHS instance, Node Manager, or Oracle BI Enterprise Edition) to create a copy of the source entity configuration files into an archive file.

This section covers the following topics:

- [Running copyConfig for Oracle WebLogic Server Domains](#)
- [Running copyConfig for System Components](#)
- [Running copyConfig for Node Manager](#)
- [Running copyConfig for Oracle BI Enterprise Edition](#)
- [Running copyConfig for the Oracle Global Order Promising Server Instance](#)

17.6.4.1 Running copyConfig for Oracle WebLogic Server Domains

J2EE domains reside in the Oracle WebLogic Server Domain configuration home, `APPLICATIONS_CONFIG/instance/domains`. `CRMDomain` and `FinancialDomain` are two examples of J2EE domains.

Running `copyConfig` creates a copy of the source Oracle WebLogic Server Java EE component into a configuration archive file that contains the Oracle WebLogic Server domain configuration and all configurations stored in MDS. It also packs the other configuration information specific to components such as SOA composites, human workflow, Oracle B2B agreements, and Oracle Application Development Framework (Oracle ADF) connection settings for all connections.

Note: `copyConfig` handles only global data sources defined in each Oracle WebLogic Server domain. For application level data sources, you must manually configure them on the target domain.

Use the following syntax to run the `copyConfig` script for each Oracle WebLogic Server domain:

```
copyConfig -javaHome absolute_location_of_JDK
           -archiveLoc archive_location
           -sourceDomainLoc source_WebLogic_domain_location
           -sourceMWHomeLoc source_Middleware_home_location
           -domainHostName source_domain_admin_server_host_name
           -domainPortNum source_domain_admin_server_port
           -domainAdminUserName source_domain_admin_server_user_name
           -domainAdminPassword source_domain_admin_server_user_password_file_
location
           [-additionalParams search.encrypt.key=pwdfile]
           [-logDirLoc existing_log_directory_location]
           [-silent {true | false}]
```

The following example copies the configuration of the `FinancialDomain` domain:

```
copyConfig.sh -javaHome USER_HOME/jrockit_160_04
             -archiveLoc /FIN_T2P/FIN_FinancialDomain1.jar
             -sourceDomainLoc /net/mount2/instance/applications/FinancialDomain
             -sourceMWHomeLoc /net/mount1/appbase/fusionapps
             -domainHostName server.us.example.com
             -domainPortNum 9001
             -domainAdminUserName faadmin
             -domainAdminPassword /FIN_T2P/domain_admin_server_pw.txt
             -logDirLoc /FIN_T2P/logs
             -silent true
```

The following example copies the configuration of the `CommonDomain: domain`. Notice the use of the `-additionalParams` option for Oracle SES.

```
configConfig.sh -javaHome USER_HOME/jrockit_160_20_D1.0.1-1705
               -archiveLoc /COMMON_T2P/COMMON_CommonDomain.jar
               -sourceDomainLoc /net/mount2/instance/applications/CommonDomain
               -sourceMWLoc /net/mount1/appbase/fusionapps
               -domainHostName server.us.example.com
               -domainPortNum 8001
               -domainAdminUserName faadmin
               -domainAdminPassword /COMMON_T2P/domain_admin_server_pw.txt
               -additionalParams search.encrypt.key=/SES_T2P/encrypt_key.txt
```

`/net/mount1/appbase` represents the top-level applications base directory for binaries, while `/net/mount2` represents the top level applications configuration directory for configuration files.

[Table 17-6](#) describes the options for the `copyConfig` script for Oracle WebLogic Server Java EE components.

Table 17-6 Options for the `copyConfig` Script for Oracle WebLogic Server Java EE Components

Option	Shortcut	Description	Mandatory or Optional?
<code>-javaHome</code>	None	The absolute location of the Java Developer's Kit.	Mandatory
<code>-d64</code>	None	If the operating system is 64-bit UNIX, pass the <code>-d64</code> option to the scripts in the command line. To set the runtime property, you can specify the <code>-d64</code> option in the <code>T2P_JAVA_OPTIONS</code> environment variable. For example: <pre>setenv T2P_JAVA_OPTIONS "-d64 -Djava.io.tmpdir=/home/t2p/temp"</pre>	Optional
<code>-archiveLoc</code>	<code>-al</code>	The absolute location of the archive. Use this option to specify the location of the archive file to be created by the <code>copyConfig</code> script.	Mandatory
<code>-sourceDomainLoc</code>	<code>-sdl</code>	The absolute location of the source domain containing the Java EE component.	Mandatory
<code>-sourceMWHomeLoc</code>	<code>-smw</code>	The absolute location of the source Middleware home.	Mandatory
<code>-domainHostName</code>	<code>-dhn</code>	The name of the host on which the source domain is configured.	Mandatory
<code>-domainPortNum</code>	<code>-dpn</code>	The port number of the source domain.	Mandatory
<code>-domainAdminUserName</code>	<code>-dau</code>	The name of the administrative user for the source domain.	Mandatory
<code>-domainAdminPassword</code>	<code>-dap</code>	The password file location for the administrative user for the domain. For example, <code>/FIN_T2P/domain_admin_server_pw.txt</code> .	Mandatory

Table 17–6 (Cont.) Options for the copyConfig Script for Oracle WebLogic Server Java EE Components

Option	Shortcut	Description	Mandatory or Optional?
-additionalParams	-ap	Absolute path of the Oracle SES password file for the encryption key. For example, <code>scratch/t2p/encrypt.txt</code> . The encryption key in the file must be at least 8 characters, must contain both letters and numbers, and must only contain ASCII characters. For example, <code>welcome1</code> . Oracle SES is available in the <code>CommonDomain</code> domain. Therefore, you do not need to use this option in the other domains.	Mandatory for the <code>CommonDomain</code> domain only
-logDirLoc	-ldl	The location of an existing log directory. A new log file is created in the directory.	Optional
-silent	None	Specifies whether the operation operates silently. That is, it does not prompt for confirmation. The default is that the operation prompts for confirmation. To continue, you must type <code>yes</code> , which is not case sensitive. Typing anything other than <code>yes</code> causes the script to return an error. To specify that it not prompt for confirmation, specify this option with the value of <code>true</code> .	Optional

17.6.4.2 Running copyConfig for System Components

Running `copyConfig` creates a copy of the source component instance for a system component, such as OHS, by copying the configuration files of that component instance into an archive file.

Use the following syntax to run the `copyConfig` script for system components:

```
copyConfig -javaHome absolute_location_JDK
           -archiveLoc archive_location
           -sourceInstanceHomeLoc source_instance_home_location
           -sourceComponentName source_component_name
           [-logDirLoc existing_log_directory_location]
           [-silent {true | false}]
```

The following example shows how to create a copy of the OHS instance named `ohs1` in the Oracle instance home located in `APPLICATIONS_CONFIG/instance/CommonDomain_webtier` on UNIX:

```
copyConfig.sh -javaHome USER_HOME/jrocket_160_17_R28.0.0-679/
              -archiveLoc /FIN_T2P/FIN_CommonDomain_webtier.jar
              -sourceInstanceHomeLoc /net/mount2/instance/CommonDomain_webtier
              -sourceComponentName ohs1
              -logDirLoc /FIN_T2P/logs
              -silent true
```

`/net/mount2` represents the top level applications configuration directory for configuration files.

[Table 17–7](#) describes the options for the `copyConfig` script for system components.

Table 17-7 Options for the copyConfig Script for System Components

Option	Shortcut	Description	Mandatory or Optional?
-javaHome	None	The absolute location of the Java Developer's Kit.	Mandatory
-d64	None	If the operating system is 64-bit UNIX, pass the -d64 option to the scripts in the command line. To set the runtime property, you can specify the -d64 option in the T2P_JAVA_OPTIONS environment variable. For example: setenv T2P_JAVA_OPTIONS "-d64 -Djava.io.tmpdir=/home/t2p/temp"	Optional
-archiveLoc	-al	The absolute location of the archive. Use this option to specify the location of the archive file to be created by the copyConfig script.	Mandatory
-logDirLoc	-ldl	The location of an existing log directory. A new log file is created in the directory.	Optional
-silent	None	Specifies whether the operation operates silently. That is, it does not prompt for confirmation. The default is that the operation prompts for confirmation. To continue, you must type yes, which is not case sensitive. Typing anything other than yes causes the script to return an error. To specify that it not prompt for confirmation, specify this option with the value of true.	Optional
-sourceComponentName	-scn	The name of the source component to be moved. For example, if your Oracle HTTP Server component is named ohs1, specify ohs1.	Mandatory
-sourceInstanceHomeLoc	-sih	The absolute location of the Oracle instance home for the source component.	Mandatory

17.6.4.3 Running copyConfig for Node Manager

Node Manager is an Oracle WebLogic Server utility that enables you to start, shut down, and restart the Administration Server and Managed Server instances from a remote location. The Node Manager for each host resides in Oracle WebLogic Server home, in `wlserver_version/common/nodemanager/host_name`.

Running copyConfig creates a copy of the source Node Manager configuration into an archive file.

Use the following syntax to run the copyConfig script for Node Manager on each machine within topology of the source environment:

```
copyConfig -javaHome absolute_location_of_JDK
           -archiveLoc archive_location
           -sourceNMHomeLoc source_Node_Manager_home_location
           [-logDirLoc existing_log_directory_location]
           [-silent {true | false}]
```

The following example shows how to create a copy of the source Node Manager configuration located in `/net/abcdef03/scratch/work/mw2903/wlserver_10.3/common/nodemanager/abcdef04.us.example.com` (with a folder corresponding to each Node Manager) into an archive file:

```
copyConfig.sh -javaHome USER_HOME/jrockit_160_17_R28.0.0-679/
              -archiveLoc /FIN_T2P/node_manager_abcdef04.jar
              -sourceNMHomeLoc /net/mount1/appbase/fusionapps/wlserver_
10.3/common/nodemanager/abcdef04.us.example.com
              -silent true
```

Table 17–8 describes the options for the `copyConfig` script for Node Manager.

Table 17–8 Options for the `copyConfig` Script for Node Manager

Option	Shortcut	Description	Mandatory or Optional?
<code>-javaHome</code>	None	The absolute location of the Java Developer's Kit.	Mandatory
<code>-d64</code>	None	If the operating system is 64-bit UNIX, pass the <code>-d64</code> option to the scripts in the command line. To set the runtime property, you can specify the <code>-d64</code> option in the <code>T2P_JAVA_OPTIONS</code> environment variable. For example: <pre>setenv T2P_JAVA_OPTIONS "-d64 -Djava.io.tmpdir=/home/t2p/temp"</pre>	Optional
<code>-archiveLoc</code>	<code>-al</code>	The absolute location of the archive. Use this option to specify the location of the archive file to be created by the <code>copyConfig</code> script.	Mandatory
<code>-sourceNMHomeLoc</code>	<code>-snh</code>	The absolute location of the source Node Manager Home.	Mandatory
<code>-logDirLoc</code>	<code>-ldl</code>	The location of an existing log directory. A new log file is created in the directory.	Optional
<code>-silent</code>	None	Specifies whether the operation operates silently. That is, it does not prompt for confirmation. The default is that the operation prompts for confirmation. To continue, you must type <code>yes</code> , which is not case sensitive. Typing anything other than <code>yes</code> causes the script to return an error. To specify that it not prompt for confirmation, specify this option with the value of <code>true</code> .	Optional

17.6.4.4 Running `copyConfig` for Oracle BI Enterprise Edition

You copy the configuration of the following Oracle BI Enterprise Edition components:

- Oracle BI Server
- Oracle BI Presentation Services
- Oracle BI Cluster Controller
- Oracle BI Scheduler
- JavaHost
- Oracle Essbase server, if it is installed in your environment

Running `copyConfig` creates a copy of the source Oracle BI Enterprise Edition configuration into an archive file. You do not need to use the `copyconfig` script for each component. Use the `copyconfig` script once for all the Oracle BI Enterprise Edition components.

Use the following syntax to run the `copyConfig` script to run a complete Oracle BI Enterprise Edition instance copy:

```
copyConfig -javaHome path_of_jdk
           -archiveLoc archive_location
           -sourceInstanceHomeLoc src_instance_path
           [-additionalParams
           essbaseServerUserName=xxx, essbaseServerPassword=pwdfile]
```

```
[-logDirLoc existing_log_directory_location]
[-silent {true | false}]
```

The following example shows how to apply the copy of Oracle BI Enterprise Edition to the Oracle BI Enterprise Edition home located in BIDomain. /net/mount2/ represents the top-level Applications configuration directory, where the configuration files reside.

```
copyConfig -javaHome USER_HOME/jrockit_160_17_R28.0.0-679/
           -archiveLoc /BI_T2P/BIInstance.jar
           -sourceInstanceHomeLoc /net/mount2/appbase/instance/BIInstance
           -additionalParams
           essbaseServerUserName=weblogic,essbaseServerPassword=/tmp/welcome1.txt
           -silent true
```

Table 17–9 describes the options for the copyConfig script for Oracle BI Enterprise Edition.

Table 17–9 Options for the copyConfig Script for Oracle BI Enterprise Edition

Option	Shortcut	Description	Mandatory or Optional?
-javaHome	None	The absolute path of the Java Developer's Kit.	Mandatory
-archiveLoc	-al	The absolute path of the archive location. Use this option to specify the location of the archive file created by the copyConfig script.	Mandatory
-sourceInstanceHomeLoc	-sih	The absolute path of the Oracle BI Enterprise Edition instance for the source component.	Mandatory
-additionalParams	-ap	Specifies Oracle Essbase administrative user and password.	Mandatory
-logDirLoc	-ldl	The location of an existing directory. A new log file is created in the directory.	Optional
-silent	None	Specifies whether the operation operates silently. That is, it does not prompt for confirmation. The default is that the operation prompts for confirmation. To specify that it not prompt for confirmation, specify this option with the value of true.	Optional

17.6.4.5 Running copyConfig for the Oracle Global Order Promising Server Instance

Running copyConfig creates a copy of the source Oracle Global Order Promising Server configuration into an archive file.

Use the following syntax to run the copyConfig script to run a complete Oracle Global Order Promising Server instance copy:

```
copyConfig -javaHome path_of_jdk
           -archiveLoc archive_location
           -sourceInstanceHomeLoc src_instance_path
           -sourceComponentName GlobalOrderPromisingServer1
           [-logDirLoc existing_log_directory_location]
           [-silent {true | false}]
```

The following example shows how to create a copy of Oracle Global Order Promising Server configuration located in /net/mount1/appbase/instance/gop_1 into an archive file.

```
copyConfig -javaHome USER_HOME/jrockit_160_17_R28.0.0-679/
```

```

-archiveLoc USER_HOME/gop_1.jar
-sourceInstanceHomeLoc /net/mount1/appbase/instance/gop_1
-sourceComponentName GlobalOrderPromisingServer1
-logDirLoc USER_HOME/clone_logs
-silent false

```

Table 17–10 describes the options for the `copyConfig` script for Oracle Global Order Promising Server.

Table 17–10 Options for the `copyConfig` Script for Oracle Global Order Promising Server

Option	Shortcut	Description	Mandatory or Optional?
<code>-javaHome</code>	None	The absolute path of the Java Developer's Kit.	Mandatory
<code>-archiveLoc</code>	<code>-al</code>	The absolute path of the archive location. Use this option to specify the location of the archive file created by the <code>copyConfig</code> script.	Mandatory
<code>-sourceComponentName</code>	<code>-sih</code>	The name of the existing component which will be archived, <code>GlobalOrderPromisingServer1</code> .	Mandatory
<code>-sourceInstanceHomeLoc</code>	<code>-sih</code>	The absolute path of the Oracle Global Order Promising Server instance for the source component.	Mandatory
<code>-logDirLoc</code>	<code>-ldl</code>	The location of an existing directory. A new log file is created in the directory.	Optional
<code>-silent</code>	None	Specifies whether the operation operates silently. That is, it does not prompt for confirmation. The default is that the operation prompts for confirmation. To specify that it not prompt for confirmation, specify this option with the value of <code>true</code> .	Optional

17.6.5 Extracting the Move Plan

Run the `extractMovePlan` script to extract the configuration settings from multiple configuration archives of the source environment into a single move plan (XML file). Extracting the move plan from multiple configuration archives consolidates and tries to remove duplicate configuration properties in the consolidated move plan. It also tries to remove the configuration properties that could be derived from other configuration properties. Running `extractMovePlan` also extracts other application specific standard configuration plans such as the configuration plan for SOA composite, the deployment plan for J2EE App, and so on.

Note: Extract the move plan in a location that is shared by all systems in the configuration topology. You can then use the move plan and underlying configuration or deployment plans from the same location to recreate multiple domains or system component instances.

Use the following syntax to run the `extractMovePlan` script:

```

extractMovePlan -javaHome absolute_location_of_JDK
                -archiveLoc archive_location [,archive_location...]
                -planDirLoc move_plan_directory
                [-optimizationHints fusionApps,sameSchemaNameSinglePassword]
                [-logDirLoc existing_log_directory_location]

```

The following example extracts configuration information from two configuration archives into a single move plan:

```
extractMovePlan.sh -javaHome USER_HOME/jrockit_160_20_D1.0.1-1705
                  -archiveLoc /FIN_T2P/FIN_CommonDomain1.jar,/FIN_T2P/FIN_
FinancialDomain1.jar
                  -planDirLoc /FIN_T2P/CommonDomain1_plan
                  -opth fa,ssnsp
                  -logDirLoc /FIN_T2P/logs
```

A list of configuration archive locations, separated by comma, is specified to create a single move plan containing the configuration settings from the listed configuration archives.

Note: Make sure there is no space before or after the comma in the list of configuration archive locations.

You can also extract the configuration settings from each individual configuration archive to create multiple move plans. However, this neither consolidates, nor optimizes the move plans.

Table 17–11 describes the options for the `extractMovePlan` script:

Table 17–11 Options for the `extractMovePlan` Script

Option	Shortcut	Description	Mandatory or Optional?
-javaHome	None	The absolute location of the Java Developer's Kit.	Mandatory
-d64	None	If the operating system is 64-bit UNIX, pass the <code>-d64</code> option to the scripts in the command line. To set the runtime property, you can specify the <code>-d64</code> option in the <code>T2P_JAVA_OPTIONS</code> environment variable. For example: <pre>setenv T2P_JAVA_OPTIONS "-d64 -Djava.io.tmpdir=/home/t2p/temp"</pre>	Optional
-archiveLoc	-al	The absolute location of one or more archives. Use this option to specify the locations of the archive files created by the <code>copyConfig</code> script. If you specify more than one archive, separate the locations with a comma with no spacing in between.	Mandatory

Table 17–11 (Cont.) Options for the extractMovePlan Script

Option	Shortcut	Description	Mandatory or Optional?
-planDirLoc	-pdl	The absolute path to a directory to which the move plan, along with any needed configuration plans, is to be extracted. The directory must not already exist.	Mandatory
-logDirLoc	-ldl	The location of an existing directory. A new log file is created in the directory.	Optional
-optimizationHints	-opth s	Specifies the configuration values to auto-populate based on the topology on the target environment. These values are omitted from the move plan. Use of the hints is recommended if they apply to your environment. The -optimizationHints takes the following arguments: <ul style="list-style-type: none"> ▪ fusionapps (fa). You need to provide this hint always. This simplifies the move plan for Oracle Fusion Applications. ▪ sameSchemaNameSinglePassword (ssnsp). The same password is used for all schemas associated with a particular database. However, the passwords for APPID users are different. You are prompted to provide the password only once for each database. Therefore, you need to pass a separate password for each of the APPID users. 	Optional

17.6.6 Modifying the Move Plan

While extracting the move plans, the framework internally derives the properties and optimizes. In those cases, you may not see all the properties in the move plan. If the framework does not find a matching value, it does not derive the property. In those cases, you will see the properties in the move plan.

After the move plan is extracted, you must modify one consolidated one with a text editor to specify the properties (such as data source definitions, host names, port numbers, and end point addresses) for the target environment before you apply the copied configurations from the source environment to the target environment.

The properties that you edit differ depending on the type of component. [Table 17–12](#) provides pointers to the appropriate list of properties for each component. Some of the references link to tables in the "Move Plans for Components" section of the *Oracle Fusion Middleware Administrator's Guide*; other links direct you to tables in this chapter. Tables referenced in the *Oracle Fusion Middleware Administrator's Guide* are generic tables that work for both Oracle Fusion Applications and Oracle Fusion Middleware environments. To avoid confusion, only use the tables referenced.

Table 17–12 Move Plan Properties for Components

Component	Where to find the list of properties:
Oracle Fusion Applications FusionAppsTopology	Section 17.6.6.1
Oracle Fusion Applications patching framework	Section 17.6.6.2
Oracle Fusion Applications additional environment properties	Section 17.6.6.3
Node Manager	"Move Plan Properties for Node Manager"

Table 17–12 (Cont.) Move Plan Properties for Components

Component	Where to find the list of properties:
All Java components	"Common Move Plan Properties for Java Components"
Oracle BI EE	Section 17.6.6.4
Oracle ADF connections	"Move Plan Properties for Oracle ADF Connections"
Oracle Data Integrator	Section 17.6.6.5
Oracle Enterprise Crawl and Search Framework	Section 17.6.6.6
Oracle Enterprise Scheduler Service (ESS)	Section 17.6.6.7
Oracle Global Order Processing	Section 17.6.6.8
Oracle HTTP Server	"Move Plan Properties for Oracle HTTP Server"
Oracle Secure Enterprise Search	Section 17.6.6.9
Oracle SOA Suite	"Move Plan Properties for Oracle SOA Suite"
Oracle Topology Manager	Section 17.6.6.10
Oracle Virtual Directory	"Move Plan Properties for Oracle Virtual Directory"
Oracle WebCenter Content: Inbound Refinery	"Move Plan Properties for WebCenter Content Server, Records, and Inbound Refinery"
Oracle WebCenter Content: Imaging	Section 17.6.6.11, "Editing Oracle WebCenter Content: Imaging Properties"

17.6.6.1 Editing Oracle Fusion Applications FusionAppsTopology Properties

Modify the move plan by editing the specific Oracle Fusion Applications entries within the `movableComponent` of type `ApplicationTopology` and name `FusionAppsTopology` according to the target environment.

[Table 17–13](#) describes the `FusionAppsTopology` properties to edit in the move plan, if they are present.

Table 17–13 Move Plan Properties for Oracle Fusion Applications FusionAppsTopology Properties

Configuration Group	Property	Description	Sample Value
WEBTIER_CONFIG	Listen Address	The listen address used by OHS on the webtier host. This is usually the hostname of the machine where OHS server is running.	ohs_host1.oracle.com
APPTIER_CONFIG	Listen Address	The listen address used by Managed Servers on the host. This is usually the host name of the machine where Managed Server is running.	app_host1.oracle.com
	Custom Identity Keystore File	The encrypted custom identity keystore's passphrase file	/scratch/identity_keystore.jks

Table 17–13 (Cont.) Move Plan Properties for Oracle Fusion Applications FusionAppsTopology

Configuration Group	Property	Description	Sample Value
	Custom Identity Private Key Alias	The alias used when loading the private key for Oracle WebLogic Server from the keystore.	app_host1.oracle.com_identity_private_key
	Custom Identity Keystore Passphrase File	Path to the password file which will be used while loading the private key for Oracle WebLogic Server into the keystore)	/scratch/ identity_private_key_pass.txt
	Custom Trust Keystore File	Path to the file containing custom trust keystore	/scratch/ trust.jks
	Custom Trust Keystore Passphrase File	JDBC connect string to be used as a J2EE datasource	/scratch/trust_pass.txt
DATATIER_CONFIG	Url	The JDBC URL in the format of jdbc:oracle:thin:@//DB_host_name.domain/sid to be used as a J2EE datasource	jdbc:oracle:thin:@db_host1.example.com:4119/orcl
	User and Password File	<p>The Oracle database schema name (user name) to login to the database.</p> <p>During execution of the extractMovePlan script, if the ssnsnp flag is used with the -optimizationHints option, there will be a single Password File property that will be used by all users associated with the datasource mentioned in the Url configuration property. However, the password for the APPID users will be different. You need to derive the passwords for those APPID users from target IDMDomain domain. For these APPID users, both the User and Password File configuration property will be populated. For more information, see Task 2, "Seed Oracle ADF Credentials".</p> <p>During execution of extractMovePlan command, if the ssnsnp flag is not used with the -optimizationHints option, there will be a User and Password File property for each database schema.</p>	fusion_runtime /private/ fusion_runtime_password.txt

17.6.6.2 Editing Oracle Fusion Applications Patching Framework Properties

You must manually update the Oracle Fusion Applications patching framework configuration information to specify the properties for the target domain. To update the Oracle Fusion Applications patching framework configuration information, edit the values of the patching framework properties in the move plan.

[Table 17–14](#) describes the move plan properties that you can change for the patching framework

Table 17–14 Move Plan Properties for Oracle Fusion Applications Patching Framework

Properties	Description	Example
Top directories for various components used by Patching Framework	The following properties are defined in <code>FUSION_prov.properties</code> :	
<code>APPTOP</code>	Applications base directory, <code>APPLICATIONS_BASE</code>	<code>/u01/APPTOP</code>
Database-Related Properties	The following properties are in the <code>ASOAppsTableSpaceCustomizations</code> config group. Each aggregate storage (ASO) application has a name and tablespace property, followed by the following properties:	
<code>TWO_TASK</code>	Database service as defined by the <code>SERVICE_NAME</code> parameter in the <code>tnsnames.ora</code> file	<code>prod</code>
Taxonomy Related Properties		
<code>TAXONOMY_URL</code>	URL to connect to the MBean to fetch taxonomy information. This URL is usually the same as the Administration Server URL. See Section 2.4 to find the URL for the Administration Server.	<code>adc6170361.us.oracle.com:7001</code>
Oracle Access Manager-related properties	The following properties are defined in <code>FUSION_prov.properties</code> :	
<code>OAM_ACCESS_SERVER_HOST</code>	Host name for Oracle Access Manager	<code>adc6260031.us.oracle.com</code>
<code>OAM_ACCESS_SERVER_PORT</code>	Port for Oracle Access Manager connection	<code>3060</code>
<code>OAM_ADMIN_SERVER_HOST</code>	Host for the Administration Server on the Oracle Access Manager	<code>adc6260031.us.oracle.com</code>
<code>OAM_ADMIN_SERVER_PORT</code>	Host for the Administration Server on the Oracle Access Manager connection	<code>3060</code>
<code>OAM_WEB_DOMAIN</code>	Oracle Access Manager web domain used for Oracle Fusion Applications policies	<code>OraFusionApp</code>
<code>OAM_ACCESS_SERVER_MODE</code>	Mode for the Oracle Access Manager server	<code>Simple</code> <code>Open</code>

17.6.6.3 Editing Additional Oracle Fusion Applications Environment Properties

This section describes the move plan properties for Oracle Fusion Applications, startup configuration, Oracle Human Workflow configuration, and Oracle Database client configuration.

Note: The Oracle Database client configuration is reset to the initial out-of-box configuration set by provisioning. Any additional properties added to the source environment are not moved to the target environment. Therefore, you need to copy any additional properties to the target environment.

Table 17–15 describes the Oracle Fusion Applications properties.

Table 17–15 Move Plan Properties for Oracle Fusion Applications Properties

Configuration Group	Property	Description	Sample Value
Fusion Applications Environment Properties	APPLICATIONS_BASE	The top-level directory for the Oracle Fusion Applications binaries	/scratch/aime/work/APPT OP
	APPLICATIONS_CONFIG	The top-level directory for the Oracle Fusion Applications configuration files (includes the instance directory)	/scratch/aime/work/APPT OP/instance
wf_client_config	rootEndPointURL	OHS or load balancer URL value for the SOA cluster configured in wf_client_config.xml.	http://abc.example.com: 10603
Database Client Configuration	Wallet Password File	Location of password file which contains password to create the database client wallet	/tmp/wallet_ password.txt
	ORACLE_SERVICE_NAME	Database service name of the tnsnames.ora file to connect to database	adbc1234
	DB_HOST_NAME	Database host name of the tnsnames.ora file to connect to database. (In the case of Oracle Real Application Clusters (Oracle RAC) database the property names are DB_HOST_NAME0, DB_HOST_NAME1, and so forth.)	adc.example.com
	DB_PORT	Database port number of the tnsnames.ora file to connect to database (In the case of Oracle RAC databases, the property names are: DB_PORT0, DB_PORT1, and so forth)	1521
	Fusion Runtime Schema	Database schema name (user name) to login to database	fusion_runtime

Table 17–15 (Cont.) Move Plan Properties for Oracle Fusion Applications Properties

Configuration Group	Property	Description	Sample Value
	Fusion Runtime Schema Password File	Location of password file which contains schema password to login to the database	/tmp/fusion_runtime_password.txt
Fusion Applications Front end Host URL Configuration	FusionAppsFrontendHostUrl	<p>The external OHS URL for all the Oracle Fusion applications human workflow task flows. The format of FusionAppsFrontendHostUrl attribute is as follows:</p> <p><i>FANamespace1=https://host1:port1::FANamespace2=https://host2:port2::*=https://host1:port1</i></p> <p>where * is the wild card namespace. If the wildcard is specified, the frontend host URL of the wild card will be applied to all Oracle Fusion Applications servers, including the SOA server.</p> <p>If <i>FANamespace</i> is <i>soa.server.url</i>, then the provided URL will be used for override the SOA server URL.</p> <p>For example, if SupplierPortal has namespace SUPPLIER, the format for adding FrontendHostUrl for Supplier Portal FusionAppsFrontendHostUrl will be:</p> <p><i>SUPPLIER=https://supportportalFrontendHostName:port::*=https://OtherFAFrontendHostHame:port1</i></p>	*=https://agc.example.com:10608
	SupplierPortalFrontendHostUrl	The external OHS URL for all the Oracle Fusion applications human workflow task flows (ProcurementDomain)	https://acd.example.com:10610

17.6.6.4 Editing Oracle BI EE Properties

This section describes the move plan properties for Oracle BI Enterprise Edition, Oracle BI EE Data Warehouse Administration Console, Oracle Essbase, EPM Registry, and Oracle BI Action Framework.

Table 17–16 shows the move plan properties that you can change for Oracle BI Enterprise Edition.

Table 17–16 Move Plan Properties for Oracle BI EE

Properties	Description	Sample Value
Properties in the XMLP-SERVER-CONFIG configGroup:	Oracle BI Publisher configuration	
SAW_SERVER	The name of the host that is running the Oracle BI Presentation Services to which you must connect	example_host
SAW_PORT	The port number for connecting to Oracle BI Presentation Services	10217
SAW_PASSWORD	The absolute path of a secure file that contains the password for Oracle BI Presentation Services	/scratch/oracle/bip_pass.txt
SAW_USERNAME	The user name for Oracle BI Presentation Services	user1
Properties in the connection configProperty:	connection is a sub-property of dataSource. Specifies that the data source is a Connection type. Each data source can be either a Connection or File type.	
url	The URL of the connection	jdbc:oracle:thin:@host:port:sid
driver	The driver to use for the connection	oracle.jdbc.OracleDriver
username	The user name for the connection	user1
password	The absolute path of a secure file that contains the password for the connection	/scratch/oracle/ds_conn_pass.txt
Properties in the file configProperty:	file is a sub-property of dataSource. Specifies that the data source is a File type.	
path	The file system path that points to the relevant data source file	/scratch/APPTOP/instance/domains/example.com/BIDomain/config/bipublisher/repository/DemoFiles
Properties in the XDO-CLIENT_CONFIG configGroup:	Oracle BI Publisher client	
XMLPClientDirPath	The absolute path to the Oracle BI Publisher client directory	/scratch/APPTOP/instance/domains/example.com/CommonDomain/config
* (any name)	The Oracle BI Publisher configured endpoint connecting URL. The move plan may have more than one endpoint	http://example.com:10621/xmlpserver
Properties in the XMLP-SCHEDULER-JMS-CONFIG configGroup:	The Scheduler configuration	
JMS_WEBLOGIC_JNDI_URL	The Oracle WebLogic Server JNDI URL for the Oracle BI EE Managed Server	cluster:t3://bi_cluster
JMS_Shared_Temp_Directory	The JMS shared temporary directory used in an Oracle BI EE cluster environment	/scratch/APPTOP/instance/BIShared/BIPublisher/biptemp
Properties in the XMLP-PROVIDER-CONFIG configGroup, in the provider configProperty:	Oracle BI Publisher provider There is a separate configProperty for each BI Publisher provider configured.	

Table 17–16 (Cont.) Move Plan Properties for Oracle BI EE

Properties	Description	Sample Value
uri	The URI for the Oracle BI Publisher provider	http://example.com:10603/financialCommon
nonSSOUri	The non-SSO URI for the Oracle BI Publisher provider	http://example.com:7404/financialCommon
Properties in the XDO-SERVER_CONFIG configGroup:		
XMLPServerConfigDirPath	The absolute path to the Oracle BI Publisher server configuration directory	/scratch/APPTOP/instance/domains/example.com/BIDomain/config/bipublisher
	The following property is in the RepositoryDirPath configProperty, the directory path of the Oracle BI Publisher server repository	
path	The directory of the Oracle BI Publisher server configuration repository. (It can be located outside of the BIDomain that is specified as the server xdo.server.config.dir system property.) The default value is \${xdo.server.config.dir}/repository.	/scratch/APPTOP/instance/BIShared/BIPublisher/repository
	The following property is in the SawSourcePath configProperty. The Oracle BI Publisher server connection resource details for the Oracle BI EE Presentation Server.	
source	The path for the Oracle BI Publisher repository directory	/scratch/APPTOP/instance/domains/example.com/BIDomain/config/bipublisher/repository
Properties in the OracleInstances configGroup:		
instanceHome	The path of the Oracle instance in which Oracle BI EE is deployed	/scratch/APPTOP/instance/BIInstance
host	The host where the Oracle BI EE instance is configured	example.com
Properties in the BIInstanceDeployment configProperty:		
listenAddress	The listen address for the host. It can be set to a virtual IP address or a subset on a multi-homed computer. You can specify an asterisk to specify multiple network addresses for the host	example.com *
portRangeStart	The start of the range of ports used by the Oracle BI EE system components	10206
portRangeEnd	The end of the range of ports used by the Oracle BI EE system components	10214

Table 17–16 (Cont.) Move Plan Properties for Oracle BI EE

Properties	Description	Sample Value
Properties in the BIInstance configGroup, in the EmailOptions configProperty:	Oracle BI EE instance configuration	
smtpServerName	The host name of the SMTP server	example.com
port	The port number of the SMTP server	25
fromDisplayName	The sender's name that is used as the display name by the Oracle BI EE system when it sends email	Oracle Business Intelligence
emailAddressOfSender	The email address used by the Oracle BI EE system when it sends email	defaultuser@defaultmailserver.com
Property in the BIInstance configGroup, in the MarketingOptions configProperty:	Oracle BI EE instance configuration	
url	The base URL used by the Oracle BI EE system when the emails have embedded URLs	http://example.com:7012/_dav/cs/idcplg
Property in the BIInstance configGroup, in the PresentationServerOptions configProperty:	Presentation Server configuration.	
webCatalogLocation	The absolute path of the location of the Oracle BI Presentation Catalog	/scratch/APPTOP/instance/BIShared/OracleBIPresentationServicesComponent/coreapplication_obips1/catalog/OracleBIApps
Property in the BIInstance configGroup, in the Scheduler configProperty:	Scheduler configuration	
dataSource	The connection details for the Oracle BI Scheduler data source	(DESCRIPTION=(ADDRESS_LIST=(ADDRESS=(PROTOCOL=TCP)(HOST=example.com)(PORT=1565)))(CONNECT_DATA=(SERVICE_NAME=d8b41fc1))
Properties in the BIInstance configGroup, in the ServerOptions configProperty:	Server options configuration	
repositorySharedLocation	The shared location for the Oracle BI repository	/scratch/APPTOP/instance/BIShared/OracleBIServerComponent/coreapplication_obis1/repository
Properties in the BIInstance configGroup, in the PerformanceOptions configProperty:	Performance options configuration	
globalCacheStoragePath	The global location of the Oracle BI EE server cache	/scratch/APPTOP/instance/BIShared/OracleBIServerComponent/coreapplication_obis1/cache
Properties in the DEPLOY_USER_CREDENTIALS configGroup:	Oracle RTD Inline Services (BI_RTD_SPE_ILS_DEPLOY_CONFIG)	

Table 17–16 (Cont.) Move Plan Properties for Oracle BI EE

Properties	Description	Sample Value
username	The user name used to deploy the RTD SPE inline service	weblogic
password	The absolute path of a secure file that contains the password for the connection	/scratch/oracle/rtd_pass.txt
Properties in the CONNECTIONPOOLS configGroup:	RPD configuration	
user	Connection pool user name (the database schema name). The name may be a variable, in the format VALUE_OF (varname) which would then appear in the VARIABLES configuration group.	VALUEOF (ORACLE_INITBLOCK_USER)
datasource	Oracle BI repository connection pool data source name or definition. The name may be a variable, in the format VALUE_OF (varname) which would then appear in the VARIABLES configuration group.	VALUEOF (ORACLE_INITBLOCK_DSN)
appServerName	If this is an ADF connection, the Business Component URL	http://example.com:10603/fscmAnalytics/obieebroker
password	The absolute path of a secure file that contains the password for the connection to the Oracle BI repository data source	/scratch/oracle/rpd_ds_conn_pass.txt
Properties in the VARIABLES configGroup:	Definition of variables	
name	The name of the variable that is used in the Oracle BI repository connection pool definitions. There can be multiple name/value pairs.	ORACLE_INITBLOCK_USER
value	The value of the variable that is used in the Oracle BI repository connection pool definitions. There are multiple name/value pairs.	'ORA_INIT_USER'

Table 17–17 describes the move plan properties that you can change for Oracle Business Intelligence Data Warehouse Administration Console (DAC).

Table 17–17 Move Plan Properties for Oracle Business Intelligence Data Warehouse Administration Console (DAC)

Properties	Description	Example
Properties in the DAC-SERVER-CONFIGURATION configGroup:	DAC configuration	
jdbc.url	The URL to connect to the DAC repository	jdbc:oracle:thin:@example.com:1521/example.com
jdbc.driver	The name of the JDBC driver	oracle.jdbc.driver.OracleDriver
Username	The user name used to connect to the DAC repository	IMPORT_DAC

Table 17–17 (Cont.) Move Plan Properties for Oracle Business Intelligence Data Warehouse Administration Console (DAC)

Properties	Description	Example
Password File	The absolute path of a secure file containing the password for the user to connect to the DAC repository. You must provide a password file, even if you are not changing the configuration.	/scratch/biplatform/cloning/password_DAC.txt
Properties in the EMAIL-CONFIGURATION configGroup:	Email configuration.	
email_host	The host name of the email server	example
email_protocol	The protocol for the email server	smtp
email_address	The email address of the user	test@test.te
needs_authentication	The flag indicating whether the corporate email server requires authentication. Valid values are true or false.	true
needs_ssl	The flag indicating whether an SSL connection is required. Valid values are true or false.	false
email_host_port	The port where the email server listens	5555
email_user	User name for the email account The email user needs to be provided even though email is not configured.	test
email_password	The absolute path of a secure file containing the password for the user of the email server. (Only required if needs_authentication is set to true.). The secure file containing the password for the user should contain the password in clear text and not an obfuscated password. The email password needs to be provided even though email is not configured.	/scratch/biplatform/cloning/password_email.txt
Properties in the DATAWAREHOUSE-CONFIGURATION configGroup:	Data Warehouse configuration	
jdbc.url	The URL to connect to the Data Warehouse	jdbc:oracle:thin:@example.com:1521/example.com
jdbc.driver	The name of the JDBC driver	oracle.jdbc.driver.OracleDriver
Username	The user name used to connect to the Data Warehouse	IMPORT_DW
Password File	The absolute path of a secure file containing the password for the user to connect to the Data Warehouse. You must provide a password file, even if you are not changing the configuration.	/scratch/biplatform/cloning/password_DW.txt
Properties in the INFORMATICA-CONFIGURATION configGroup:	Informatica configuration	
Informatica server home	The Informatica server home	/scratch/infahome/

Table 17–17 (Cont.) Move Plan Properties for Oracle Business Intelligence Data Warehouse Administration Console (DAC)

Properties	Description	Example
Domains infa file location	The domain's infa file location	/scratch/infahome/domains.info
InformaticaParameterFileLocation	The directory where the Informatica parameter files are stored (or DEFAULT)	DEFAULT
Properties in the DATASOURCES-CONNECTION-DETAILS configGroup:	Data source connection information	
type	The physical data source type. Possible values are: Source, Warehouse, Informatica Repository, DAC Repository, or Other.	Source
Connection Type	The type of database connection. Possible values are: BI Server, Oracle (OCI8), Oracle (Thin), DB2, DB2-390, MSSQL, Teradata, and Flat File.	Oracle (Thin)
Connection String	The data source connection string. If you are using: <ul style="list-style-type: none"> ■ Oracle (OCI8): Use the tnsnames.ora file entry. ■ Oracle (Thin): Use the instance name. SQL Server: Use the database name. ■ DB2-UDB/DB2-390: Use the connect string as defined in the DB2 configuration. ■ Teradata: Use the database name. 	orcl.example.com
Table Owner	The name of the table owner	DB_USER
Host	The host name of the server where the database resides	example.com
Port	The port where the database receives requests	1521
JDBC Driver (Optional)	The JDBC driver for the data source connection. The value in this field must conform to the database specifications.	oracle.jdbc.driver.OracleDriver
URL (Optional)	The JDBC URL for the data source connection. The value in this field must conform to the database specifications.	jdbc:oracle:thin:@example.com:1521/orcl.example.com
Password File	The absolute path of a secure file containing the password for the user to connect to data source. You must provide a password file, even if you are not changing the configuration.	/scratch/biplatform/cloning/password_ds.txt
Connection Pool Name (BIPool)	The connection pool name	FSCM_OLTP."Connection Pool"
Database Type (BIPool)	Database type of the transactional data source	Oracle

Table 17–17 (Cont.) Move Plan Properties for Oracle Business Intelligence Data Warehouse Administration Console (DAC)

Properties	Description	Example
Properties in the EXTERNAL-EXECUTORS configGroup:	External executors configuration	
Execution type	The execution type for the tasks that will be executed by the external executor	ODI 11g Embedded Agent
name	The name of the property that must be configured to integrate DAC with other Extract, Transform, and Load (ETL) tools. There are multiple properties for the external executors. Name is the name of the property. Value is the value that defines the property.	<name>ODIUser</name> <value>TestUser</value>

[Table 17–18](#) describes the move plan properties that you can change for Oracle Essbase.

Table 17–18 Move Plan Properties for Oracle Essbase

Properties	Description	Example
Properties in the EssbaseAgentConfig configGroup:	Oracle Essbase configuration.	
EssbaseAdminUserName	The administration user name for Oracle Essbase	weblogic
EssbaseAdminPassword	The absolute path of a secure file containing the password for the Oracle Essbase administration password	/scratch/oracle/essbase_passwd.txt
Properties in the ASOAppsTableSpaceCustomizations configGroup:	Aggregate Storage (ASO) application configuration. Each aggregate storage (ASO) application has a name and Table Space property, followed by the properties in the configGroup.	
file_location	The absolute path of the application file	/scratch/oracle/aso
Properties in the BSOAppsDiskVolumeCustomizations configGroup:	Block Storage (BSO) application configuration. Each Essbase block storage (BSO) application has a name and database name property, followed by the properties in the configGroup.	
volume	The location of the disk volume	/scratch/biplatform
file_type	The file type of the disk volume, such as index, data, or index_data	index_data

[Table 17–19](#) describes the move plan properties that you can change for the EPM Registry.

Table 17–19 Move Plan Properties for the EPM Registry

Properties	Description	Example
Properties in the reg.properties configGroup:	EPM Registry	
jdbc.url	The URL to connect to the EPM Registry database	jdbc:oracle:thin:@example.com:1570/db20258
jdbc.driver	The name of the JDBC driver	oracle.jdbc.OracleDriver
username	The user name used to connect to the EPM database	FUSION_BIPLATFORM
password	The absolute path of a secure file containing the password for the user to connect to the EPM database. You must provide a password file, even if you are not changing the configuration.	/scratch/biplatform/password_epm_jdbc.txt
Properties in the EPM_COMPONENTS configGroup, in the DATABASE_CONN configProperty:	EPM components configuration	
host	The database server host name	example.com
dbUserName	The database user name	FUSION_BIPLATFORM
dbJdbcUrl	The JDBC URL to connect to the database	jdbc:oracle:thin:@example.com:1570/db20258
dbName	The database name. For Oracle Database, use the service name or SID	db20258
dbPort	The database port number	1570
dbPassword	The absolute path of a secure file containing the password for the user to connect to the database. You must provide a password file, even if you are not changing the configuration.	/scratch/biplatform/password_epm_db.txt
Properties in the EPM_COMPONENTS configGroup, in the Default configProperty:	EPM components configuration.	
host	The host name of the front-end web server or load balancer	example2.com
port	The port of the front-end web server or load balancer. This is a subentry to the Default property.	10621
isSSL	The flag indicating whether the front end is in SSL mode. Valid values are true and false.	false
SSLPort	The SSL port of the front-end web server or load balancer	10218
Properties in the WORKSPACE_APP configProperty:	Workspace configuration.	
host	The host name of the server hosting the Workspace web application	example.com
port	The port where the Workspace web application is running	10217

Table 17–19 (Cont.) Move Plan Properties for the EPM Registry

Properties	Description	Example
SSLPort	The SSL port (if configured for SSL) where the Workspace web application is running	10218
Properties in the EPM_COMPONENTS configGroup, in the WEB_SERVER configProperty:	Web server configuration.	
host	The host name of the web server that the web application is configured to use	example.com
port	The port where the web application is running	10217
isSSL	The flag indicating whether the front end is in SSL mode. Valid values are true and false.	false
Properties in the EPM_COMPONENTS configGroup, in the CALC_WEBAPP configProperty:	EPM components configuration	
host	The host name of the server hosting the Oracle Calculation Manager web application	example.com
port	The port where the Oracle Calculation Manager web application is running	10217
SSL_Port	The SSL port (if configured for SSL) where the Oracle Calculation Manager web application is running	10218
name	The name of the Oracle Essbase cluster. There may be more than one cluster	EssbaseCluster-1
Properties in the essbaseservern configProperty:	Oracle Essbase server	
host	The host name of the Oracle Essbase server	example.com
arborPath	The ARBORPATH of the Oracle Essbase server	/scratch/rmunugal/shared_essbase
ess_AppLocation	The location of the Oracle Essbase application location	/scratch/biplatform/instances/instance1/Essbase/essbaseserver1
agent_PortNumber	The agent port number of the Oracle Essbase server	9511
agent_StartPort	The start of the range of ports used by agent for Oracle Essbase server	9000
agent_StopPort	The end of the range of ports used by agent for Oracle Essbase server	9499
Properties in the BIEE_WEBAPP configProperty:	Oracle BI EE Web application configuration.	
host	The host name of the server hosting the Oracle BI EE web application.	example.com
port	The port where the Oracle BI EE web application is running	10217
SSL_Port	The SSL port (if configured for SSL) where the Oracle BI EE web application is running	10218

Table 17–19 (Cont.) Move Plan Properties for the EPM Registry

Properties	Description	Example
Properties in the PROVIDER_SERVICES_WEB_APP configProperty:	The Oracle Essbase APS Web application configuration	
host	The host name of the server hosting the Oracle Essbase APS web application	example.com
port	The port where the Oracle Essbase APS web application is running	10217
SSL_Port	The SSL port (if configured for SSL) where the Oracle Essbase APS web application is running	10218
Properties in the PFINANCIAL_REPORTING_WEB_APP configProperty:	The Financial Reporting Web application configuration.	
host	The host name of the server hosting the Financial Reporting web application.	example.com
port	The port where the Financial Reporting web application is running	10217
SSL_Port	The SSL port (if configured for SSL) where the Financial Reporting web application is running	10218

Table 17–20 describes the move plan properties that you can change for Oracle BI Action Framework

Table 17–20 Oracle BI Action Framework

Properties	Description	Sample Value
Property in the location-alias configGroup:	Action Framework configuration.	
alias_name	The URL corresponding to the action name Note that there may be more than one name/value pair	http://adc10001.us.oracle.com:9704/analytics

17.6.6.5 Editing Oracle Data Integrator Properties

You must modify the move plan to specify the Oracle Data Integrator properties for the target environment before you apply the copied configurations from the source environment to the target environment.

Table 17–21 describes the Oracle Data Integrator properties. Edit the values of each property.

Table 17–21 Oracle Data Integrator Properties

Configuration Group	Property	Description	Sample Value
Master Repository	Master Repository Id	ODI Master Repository ID. It should be different than source repository ID	502
	SUPERVISOR password file	Password for the SUPERVISOR user in ODI repository when target ODI repository was created.	/scratch/tank/src_bin/passwd_for_SUPERVISOR.txt
Physical Data Servers	Schema name	Directory location for the data source of "file" technology	/temp
	Work schema	Directory location for the data source of "file" technology	/temp
	Driver Class	Drive class for "Oracle" technology	oracle.jdbc.RedDriver
	Url	JDBC URL for connecting to the data server	jdbc:red:thin:@abc12300:1521:lion
	User	User name for the physical data servers connection	fox_odi
	Password File	Password file for the physical data servers connection	/scratch/tank/src_bin/passwd_fox_odi.txt
Agents	Host name	Agent host name	localhost
	Host port	Agent host port	12311

17.6.6.6 Editing Oracle Enterprise Crawl and Search Framework Properties

The Oracle Enterprise Crawl and Search Framework (ECSF) seed data (data in the ECSF tables) is moved from the source environment to the target environment using the configuration scripts. Running the `copyConfig` script creates a configuration archive that contains the snapshot of the ECSF seed data. The `extractMovePlan` script extracts the environment specific ECSF configuration properties from the source environment into the move plan. You must modify the move plan to specify the ECSF parameter values for the target environment before you execute the `pasteConfig` script.

ECSF plug-in execution varies from domain to domain. In the `CommonDomain` domain, it extracts all the data from ECSF tables from source database and imports the data into the target database. In the `CRMDomain`, `HCMDomain`, and `FinancialDomain` domains, if the corresponding search service component is active, it updates certain records specified in the move plan. The plug-in does not do anything in other domains.

[Table 17–22](#) describes the ECSF properties. Edit the values of the ECSF properties in the move plan.

Table 17–22 ECSF Properties

Property	Description	Sample Value
SES_ADMIN_SERVICE	URL to Oracle SES administrative web service endpoint.	http://example.com:12345/search/api/admin/AdminService

Table 17-22 (Cont.) ECSF Properties

Property	Description	Sample Value
SES_QUERY_SERVICE	URL to Oracle SES Query web service end point.	http://example.com:12345/search/query/OracleSearch
ECSF_DATA_SERVICE	URL to ECSF servlets. This URL is used to construct config feeds for Oracle SES.	http://example.com:67890/crmSearch/searchfeedservlet
ECSF_SECURITY_SERVICE	URL to ECSF servlets. This URL is used to construct config feeds for Oracle SES. ECSF runtime URL and ECSF Security Service URL can be set to two different servers, which is also what Oracle SES supports.	http://example.com:67890/crmSearch/searchfeedservlet
ECSF_REDIRECT_SERVICE	URL to ECSF feed servlets. This URL is used to redirect drilldowns or actionable results.	http://example.com:67890/crmSearch/searchfeedservlet

Note: Any properties that are internally derived do not appear in the move plan.

17.6.6.7 Editing Oracle Enterprise Scheduler Properties

Modify the move plan by editing the `environment.properties` specific entries within the `movableComponent` of type `ESS-EXT` according to the target environment.

[Table 17-23](#) describes the Oracle Enterprise Scheduler properties. Edit the following properties, if available in move plan. If one or more properties are not available, they will be auto-populated in the `environment.properties` file of the target environment. Most of the properties with values relative to target environment's Middleware Home are derived automatically by the plug-in.

Note: You do not need to change the value of the ID attribute in `configProperty` that corresponds to location of the `environment.properties` file.

Table 17-23 ESS Properties

Property	Description	Sample Value
ESSBASEPATH	The top level directory of the product	/net/mount1/appbase/fusionapps/bi/clients/epm/Essbase/EssbaseRTC
LD_LIBRARY_PATH	Colon separated set of library directories	/net/prod/oracle/lib:/usr/openwin/lib:/usr/dt/lib
ORACLE_CSF_WALLET_LOC	The CSF wallet location in the Oracle Fusion Applications Middleware home	file:/net/mount1/appbase/fusionapps/wlserver_10.3/server/lib

Table 17–23 (Cont.) ESS Properties

Property	Description	Sample Value
ORACLE_JPS_CONFIG_LOC	The location of the <code>jps-config.xml</code> file in the product family domain directory. You must provide the value for only one product family domain, and it can be from any of the domains.	<code>/net/mount1/appbase/instance/domains/abc05.us.example.com/FinancialDomain/config/fmwconfig/jps-config.xml</code>
PATH	The full path of the spawned job. In Windows environments, <code>PATH</code> must include all directories that are normally part of <code>LD_LIBRARY_PATH</code> .	<code>/usr/local/bin:/bin:/usr/bin:/net/prod/oracle/bin</code>
PAY_TOP	The top level directory of the product	<code>/net/mount1/appbase/fusionapps/applications/thirdparty/hcm/vertex</code>
TWO_TASK	The TNS name identifying the database to which spawned jobs should connect. In Windows environments, the environment variable is <code>LOCAL</code> .	<code>proddb</code>
VERTEX_TOP	The top level directory of the product	<code>/net/mount1/appbase/fusionapps/applications/thirdparty/hcm/vertex</code>

17.6.6.8 Editing Oracle Global Order Promising Properties

You must modify the move plan to specify the Oracle Global Order Promising properties for the target environment before you apply the copied configurations from the source environment to the target environment.

[Table 17–24](#) describes the Oracle Global Order Promising property you must edit.

Table 17–24 Editing Oracle Global Order Promising Properties

Property	Description	Sample
<code>GOPServerPort</code>	Global Order Promising Server listen port	<code>39000</code>

17.6.6.9 Editing Oracle SES Properties

You must modify entries in the move plan to change the writable attributes of configuration properties (`configProperty`) such as server names, user names, passwords, and so on for the target environment before you apply the copied configurations from the source environment to the target environment.

[Table 17–25](#) describes the Oracle SES properties. Edit the values of these properties in the move plan.

Note: Oracle SES is dependent on the `SEARCH_TOP` property of the `ESS-EXT` component. You must modify the value of the `SEARCH_TOP` property. For information, see [Section 17.6.6](#).

Table 17–25 Oracle SES Properties

Configuration Group	Property	Description	Sample Value
identityPlugin	HTTP endpoint for authentication	The HTTP endpoint for authentication. This is the URL for the search application deployment endpoint for the Oracle Fusion customer.	<code>http://abcdef05.us.example.com:10615/crmSearch/searchfeedservlet/SecurityService</code>
userDefinedSource	Configuration URL	The file or HTTP URL of the configuration file. For Oracle Fusion sources, the value is the URL for the search application deployment endpoint. For WebCenter sources, the value is the WebCenter RSS crawl servlet endpoint.	<p>For Oracle Fusion sources:</p> <code>http://abcdef05.us.example.com:10615/crmSearch/searchfeedservlet/oracle.apps.contracts.dashboard.textSearch.publicModel.view.ContractHeaderSearchVO/ConfigFeed?engineInstId=100010032752886</code> <p>For WebCenter sources:</p> <code>http://abcdef05.us.example.com:10613/rsscrawl?excludedServiceIds=oracle.webcenter.peopleconnections.profile</code> <p>For WebCenter Content sources:</p> <code>http://abcdef05.us.example.com:10613/cs/idcplg?IdcService=SES_CRAWLER_DOWNLOAD_CONFIG&source=default</code>
	Oracle SSO Login URL	(Optional) The Oracle Single Sign-On login URL that protects all Oracle Single Sign-On applications. This parameter is relevant when the authentication type is ORASSO.	<code>https://login.example.com/myssso/signon.jsp</code>
	Oracle SSO Action URL	(Optional) The Oracle Single Sign-On action URL that authenticates Oracle Single Sign-On user credentials. This is the URL to which the Oracle Single Sign-On login form is submitted. This parameter is relevant when the authentication type is ORASSO.	<code>https://login.example.com/sso/auth</code>
	Scratch Directory	(Optional) Local directory where status files can be temporarily written	<code>/net/adc12345/dir123</code>
	Database Connection String	The JDBC connect string for the database	<code>jdbc:wolf:thin:@xyz2120830.us.example.com:1521/ntlywcdb</code>

Table 17–25 (Cont.) Oracle SES Properties

Configuration Group	Property	Description	Sample Value
	Password	The password to log in to the database	
	URL Prefix	The prefix to the content of the URL column to form display URL. For Oracle WebCenter sources, use the Oracle WebCenter instance URL.	http://abcdef05.us.example.com:10613/webcenter
authorizationPlugin	HTTP endpoint for authorization	The HTTP endpoint for Oracle Fusion authorization. For Oracle Fusion sources, the value is the URL for the search application deployment endpoint.	For Oracle Fusion sources: http://abcdef05.us.example.com:10615/crmSearch/searchfeedservlet/SecurityService?engineInstanceId=100010032752886 For WebCenter Content sources: http://abcdef05.us.example.com:10613/cs/idcplg
	Display URL prefix	(Optional) The HTTP host to prefix the access URL to form the display URL. For Oracle Fusion sources, the value is the URL for the search application deployment endpoint.	For Oracle Fusion sources: http://abcdef05.us.example.com:10615/crmSearch/searchfeedservlet/Redirect For WebCenter Content sources: http://abcdef05.us.example.com:10613/cs
	Authorization Endpoint	The URL servicing the lookup of authorization information. For WebCenter sources, the value is the WebCenter auth servlet endpoint.	http://abcdef05.us.example.com:10613/sesUserAuth
	Authorization Database Connection String	The JDBC connect string for the database	jdbc:wolf:thin:@xyz2120830.us.example.com:1521/ntlywcdB
	Password	The password to log in to the database.	
docServiceInstance	WebCenter URL Prefix	The prefix to the content of the URL column to form the display URL. For WebCenter sources, use the WebCenter instance URL.	http://abcdef05.us.example.com:10613/webcenter

17.6.6.10 Editing Oracle Topology Manager Properties

Configuration information for Oracle Topology Manager is moved from a source environment to a target environment using the configuration scripts. However, you must modify the move plan to specify the Oracle Topology Manager properties for the target environment before you apply the copied configurations from the source environment to the target environment.

Table 17–26 describes the Oracle Topology Manager properties. Edit the values of each property.

Table 17–26 Oracle Topology Manager Properties

Configuration Group	Property	Description	Sample Value
deployed-application	external-server-protocol	Protocol used by the external server for this deployed application	http
	external-endpoint-host	Host name of the external server for this deployed application	fusionapps-systemfin-us.example.com
	external-endpoint-port	Port number of the external server for this deployed application	80
	default-url	Default URL for this deployed application	http://fusionapps-systemfin-us.example.com:80
deployed-app-without-domain	external-server-protocol	Protocol used by the external server for this deployed application	http
	external-endpoint-host	Host name of the external server for this deployed application	fusionapps-systemfin-us.example.com
	external-endpoint-port	Port number of the external server for this deployed application	80
	default-url	Default URL for this deployed application	http://fusionapps-systemfin-us.example.com:80
deployed-domain	admin-server-name	Name of the administrator server for this domain	AdminServer
	admin-server-protocol	Protocol used by the administrator server for this domain	http
	admin-host	Administrator server host name for this domain	serverhost.us.example.com
	admin-port	Administrator server port number for this domain	7101
	internal-server-protocol	Protocol used by the internal server for this domain	http
	internal-endpoint-host	Host name of the internal server for this domain	hostname.us.example.com
	internal-endpoint-port	Port number of the internal server for this domain	7102
	external-server-protocol	Protocol used by the external server for this domain	https
	external-endpoint-host	Host name of the external server for this domain	fusionapps-systemfin-us.example.com
	external-endpoint-port	Port number of the external server for this domain	80

Table 17–26 (Cont.) Oracle Topology Manager Properties

Configuration Group	Property	Description	Sample Value
	node-manager-protocol	Protocol used by the node manager for this domain	https
	node-manager-port	Port number of the node manager for this domain	5556
	jmx-port	JMX port number for this domain	9001
	em-home-url	Full path of the location where Oracle Enterprise Manager is hosted for this domain	http://emhome.us.example.com:7101/em
deployed-environment	name	Name of the environment	oracle

17.6.6.11 Editing Oracle WebCenter Content: Imaging Properties

You must modify the move plan to specify the Oracle WebCenter Content: Imaging properties for the target environment before you apply the copied configurations from the source environment to the target environment.

Table 17–27 describes the Imaging properties. Edit the values of each property.

Table 17–27 Oracle WebCenter Content: Imaging Properties

Configuration Group	Property	Description	Sample Value
	AdminUser	Administrative user ID used during the pasteConfig operation to seed system security. If left blank, the domain administrator user provided on the pasteConfig command line is used. This property should be used for situations where the Oracle I/PM administrative user must be a user other than the domain administrator. For example, during a pasteConfig operation, the Oracle I/PM database being used for the target has rows in the SYSTEM_SECURITY table. These rows grant system security privileges to the user FUSION_APPS_PROV_PATCH_APPID and no other users or groups. Therefore, for this case, you would set this value to FUSION_APPS_PROV_PATCH_APPID.	FUSION_APPS_PROV_PATCH_APPID
MBean Settings	InputAgentInputDirectories	A comma separate list of directories where input sources look for work	IPM/InputAgent/Input
	InputSampleDirectory	Directory that holds the sample data for the input UI	IPM/InputAgent/Input/Samples

Table 17-27 (Cont.) Oracle WebCenter Content: Imaging Properties

Configuration Group	Property	Description	Sample Value
	RenderGDFontPath	Location of the TrueType (TTF) font files used by the OIT rendering package	/usr/share/X11/fonts/TT F
UCM Connection	repository.machine	Location of the repository. The value must be localhost if the connection is configured for "Use Local Content Server."	localhost
	repository.port	WebCenter Content server port used when the local content server is used. If not using local content server connection, remove the configuration property.	4444
	repository.useSSL	Specifies (true or false) whether the connection to WebCenter Content systems use SSL.	false
WORKFLOW Connection	bpel.front.address	HTTP Front End Address used in the IPM SOA: Connection Settings UI	http://abc1234567:8001
	bpel.csf.key	Credential alias used in the IPM SOA: Connection Settings UI	basic.credentials

17.6.7 Applying the Configurations to the Target Environment

Apply the copied configurations from the source environment to the target environment by running the `pasteConfig` script for each source entity (for example, Oracle WebLogic Server domain, OHS instance, or Node Manager) separately, using the same modified move plan. You must run the `pasteConfig` script on the target environment.

Run the `pasteConfig` script for the source entities in the following order:

1. CommonDomain
2. Oracle Fusion Applications domains, such as `CRMDomain` and `FinancialDomain`
3. All system component instances
4. All Node Managers

For Oracle WebLogic Server domains, run the `pasteConfig` script on the Administration Server host. For system components like OHS, run the script on the host where that component runs. For Node Manager, run the script on each host.

This section covers the following topics:

- [Running pasteConfig for Oracle WebLogic Server Domains](#)
- [Running pasteConfig for System Components](#)
- [Running pasteConfig for Node Manager](#)
- [Running pasteConfig for Oracle BI Enterprise Edition](#)
- [Running pasteConfig for Oracle Global Order Promising Server Instance](#)

Note: If you extracted multiple move plans, you must use the move plan corresponding to each domain or system component instance.

17.6.7.1 Running pasteConfig for Oracle WebLogic Server Domains

Running `pasteConfig` applies the copy of the source Oracle WebLogic Server Java EE component to the target environment by unpacking and recreating the domain configuration, MDS, and component-specific configuration into the specified Middleware home.

Note: Make sure the Oracle WebLogic Server administrator user name is the same on both source and target environments. The password will be the same across the environments if the authenticator is on an embedded LDAP, but the password can be different if the authenticator is on an external LDAP. The hostname and listen port of any server can be different on the target environment, and you can modify those entries accordingly in the move plan.

You can only invoke `pasteConfig` once for each target entity (that is, Oracle WebLogic Server domain, OHS instance, or Node Manager). If you need to rerun `pasteConfig` on the same target entity, you must clean up the target environment before invoking `pasteConfig` again. For information about cleaning up target entities, see Section 24.1.9.

Use the following syntax to run `pasteConfig` for Oracle WebLogic Server domains:

```
pasteConfig -javaHome path_of_jdk
            -archiveLoc archive_location
            -targetDomainLoc trgt_domain_path
            -targetMWHomeLoc trgt_Middleware_Home_path
            -domainAdminPassword path_of_file_with_domain_admin_server_user_
password
            -movePlanLoc move_plan_path
            -appDir WLS_application_directory]
            [-logDirLoc log_dir_path]
            [-silent {true | false}]
```

The following example shows how to apply the copy of the domain to the `/fusionapps` Middleware home:

```
pasteConfig -javaHome USER_HOME/jrockit_160_20_D1.0.1-1705/
            -archiveLoc /FIN_T2P/FIN_FinancialDomain1.jar
            -targetDomainLoc /net/mount2/instance/domains/server_
name/FinancialDomain
            -targetMWHomeLoc /net/mount1/appbase/fusionapps
            -domainAdminPassword /home/oracle/password.txt
            -movePlanLoc /FIN_T2P/moveplan.xml
            -appDir /net/mount2/instance/applications/FinancialDomain
            -logDirLoc /FIN_T2P/logs
            -silent true
```

`/net/mount1/appbase` represents the top level applications base directory for binaries, while `/net/mount2` represents the top level applications configuration directory for configuration files.

Table 17–28 describes the options for the `pasteConfig` script for Oracle WebLogic Server Java EE components.

Table 17–28 Options for the `pasteConfig` Script for Oracle WebLogic Server Java EE Components

Option	Shortcut	Description	Mandatory or Optional?
<code>-javaHome</code>	None	The absolute path of the Java Developer's Kit.	Mandatory
<code>-d64</code>	None	If the operating system is 64-bit UNIX, pass the <code>-d64</code> option to the scripts in the command line. To set the runtime property, you can specify the <code>-d64</code> option in the <code>T2P_JAVA_OPTIONS</code> environment variable. For example: <pre>setenv T2P_JAVA_OPTIONS "-d64 -Djava.io.tmpdir=/home/t2p/temp"</pre>	Optional
<code>-archiveLoc</code>	<code>-al</code>	The absolute path of the archive location. Use this option to specify the location of the archive file created by the <code>copyConfig</code> script.	Mandatory
<code>-targetDomainLoc</code>	<code>-tdl</code>	The absolute path of the target domain. The domain location must not already exist for the specified Middleware home. The domain directory may be located outside of the directory structure of the Middleware home.	Mandatory
<code>-targetMWHomeLoc</code>	<code>-tmw</code>	The absolute path of the target Middleware home in which the copy of the domain will be applied.	Mandatory
<code>-domainAdminPassword</code>	<code>-dap</code>	The absolute path of the file containing the domain Administration Server user password.	Mandatory
<code>-movePlanLoc</code>	<code>-mpl</code>	The absolute path of the modified move plan in the target environment.	Mandatory
<code>-appDir</code>	<code>-ad</code>	The absolute path of the Oracle WebLogic Server application directory on the target.	Optional
<code>-logDirLoc</code>	<code>-ldl</code>	The location of an existing directory. A new log file is created in the directory.	Optional
<code>-silent</code>	None	Specifies whether the operation operates silently. That is, it does not prompt for confirmation. The default is that the operation prompts for confirmation. To specify that it not prompt for confirmation, specify this option with the value of <code>true</code> .	Optional

17.6.7.2 Running `pasteConfig` for System Components

Running `pasteConfig` applies the copy of a system component, such as OHS, to the target environment by pasting the configuration files of the source component into the specified Oracle instance.

Note: You must use an Oracle home that contains the OHS binaries when you paste the OHS configuration files into the specified Oracle instance.

Use the following syntax to run `pasteConfig` for system components:

```
pasteConfig -javaHome path_of_jdk  
            -archiveLoc archive_location
```

```

-movePlanLoc move_plan_path
-targetComponentName trgt_component_name
-targetInstanceHomeLoc trgt_Instance_path
[-targetInstanceName trgt_Instance_name]
[-targetOracleHomeLoc trgt_ORACLE_HOME_path]
[-silent {true | false}]
[-logDirLoc log_dir_path]
[-domainHostName domain_host_name]
[-domainPortNum domain_port_number]
[-domainAdminUserName domain_admin_username]
[-domainAdminPassword domain_admin_password_file]

```

The following example shows how to apply the copy of the OHS instance to the Oracle instance `CommonDomain_webtier` and how to name the moved OHS instance `ohs1`:

```

pasteConfig -javaHome USER_HOME/jrockit_160_17_R28.0.0-679/
            -archiveLoc /FIN_T2P/FIN_CommonDomain_webtier4.jar
            -movePlanLoc /FIN_T2P/moveplan.xml
            -targetComponentName ohs1
            -targetInstanceHomeLoc /net/mount2/instance/CommonDomain_webtier
            -targetInstanceName CommonDomain_webtier
            -targetOracleHomeLoc /net/mount1/appbase/webtier_mwhome/webtier
            -silent true
            -logDirLoc /FIN_T2P/logs

```

`/net/mount1/appbase` represents the top level applications base directory for binaries, while `/net/mount2` represents the top level applications configuration directory for configuration files.

Table 17–29 describes the options for the `pasteConfig` script for system components.

Table 17–29 Options for the `pasteConfig` Script for System Components

Option	Shortcut	Description	Mandatory or Optional?
<code>-javaHome</code>	None	The absolute path of the Java Developer's Kit.	Mandatory
<code>-d64</code>	None	If the operating system is 64-bit UNIX, pass the <code>-d64</code> option to the scripts in the command line. To set the runtime property, you can specify the <code>-d64</code> option in the <code>T2P_JAVA_OPTIONS</code> environment variable. For example: <pre>setenv T2P_JAVA_OPTIONS "-d64 -Djava.io.tmpdir=/home/t2p/temp"</pre>	Optional
<code>-archiveLoc</code>	<code>-al</code>	The absolute path of the archive location. Use this option to specify the location of the archive file created by the <code>copyConfig</code> script.	Mandatory
<code>-movePlanLoc</code>	<code>-mpl</code>	The absolute path of the modified move plan in the target environment.	Mandatory
<code>-targetComponentName</code>	<code>-tcn</code>	The name of the target component to be moved. The name must be unique in the instance.	Mandatory
<code>-targetInstanceName</code>	<code>-tin</code>	The name of the target Oracle instance. The name must be unique in the domain.	Mandatory for OHS

Table 17–29 (Cont.) Options for the pasteConfig Script for System Components

Option	Shortcut	Description	Mandatory or Optional?
-targetInstanceHomeLoc	-tih	The absolute path of the target Oracle instance. If the Oracle instance does not exist at that location, the command creates the instance.	Mandatory
-targetOracleHomeLoc	-toh	The absolute path of the target Oracle home. The target Oracle home must exist and it must contain the binaries for the component you are moving.	Optional, if the targetInstanceHomeLoc exists. In this case, the operation retrieves the value from the configuration.
-logDirLoc	-ldl	The location of an existing directory. A new log file is created in the directory.	Optional
-silent	None	Specifies whether the operation operates silently. That is, it does not prompt for confirmation. The default is that the operation prompts for confirmation. To specify that it not prompt for confirmation, specify this option with the value of <code>true</code> .	Optional
Domain-Detail Options		These parameters are optional. If you choose to use the domain parameters, you must use all four parameters.	
-domainHostName	-dhn	The name of the host on which the domain is configured. Use this option if you want to register the component with the domain.	Optional, if you do not want to register the component with the domain.
-domainPortNum	-dpn	The port number of the domain. Use this option if you want to register the component with the domain. The domain port number is listed in the following file as the <code>adminPort</code> . <i>ORACLE_</i> <i>INSTANCE/config/OPMN/opmn/instance.properties</i> For example: <code>adminPort=7001</code>	Optional, if you do not want to register the component with the domain.
-domainAdminUserName	-dau	The name of the administrative user for the domain. For example, <code>faadmin</code> .	Optional, if you do not want to register the component with the domain.
-domainAdminPassword	-dap	The password file for the administrative user for the domain. For example, <code>/FIN_T2P/domain_admin_server_pw.txt</code> . Use this option if you want to register the component with the domain.	Optional, if you do not want to register the component with the domain.

17.6.7.3 Running pasteConfig for Node Manager

Running `pasteConfig` applies the copy of Node Manager to the target environment by pasting the configuration files of the source Node Manager into the specified Node Manager home.

Note: All the domains that are to be managed by the Node Manager should be moved before applying the copy of Node Manager to the target environment, and the Administration Server must be in running state.

After running `pasteBinary`, if the Node Manager directory exists in `WL_HOME/common/nodemanager` then remove it before running `pasteConfig`.

Use the following syntax to run `pasteConfig` for Node Manager on each machine within the topology in the target environment:

```
pasteConfig -javaHome path_of_jdk
            -archiveLoc archive_location
            -targetNMHomeLoc trgt_Node_Manager_Home_path
            -targetMWHomeLoc trgt_Middleware_Home_path
            -movePlanLoc move_plan_path
            [-logDirLoc log_dir_path]
            [-silent {true | false}]
```

The following example shows how to apply the copy of Node Manager to the Node Manager home located in `CommonDomain_webtier`:

```
pasteConfig -javaHome USER_HOME/jrockit_160_17_R28.0.0-679/
            -archiveLoc /tmp/a.jar
            -targetNMHomeLoc /net/mount1/appbase/fusionapps/wlserver_
10.3/common/nodemanager/xyz456.us.example.com
            -targetMWHomeLoc /net/mount1/appbase/fusionapps/
            -movePlanLoc /FIN_T2P/moveplan.xml
            -silent true
```

[Table 17–30](#) describes the options for the `pasteConfig` script for Node Manager.

Table 17–30 Options for the `pasteConfig` Script for Node Manager

Option	Shortcut	Description	Mandatory or Optional?
<code>-javaHome</code>	None	The absolute path of the Java Developer's Kit.	Mandatory
<code>-d64</code>	None	If the operating system is 64-bit UNIX, pass the <code>-d64</code> option to the scripts in the command line. To set the runtime property, you can specify the <code>-d64</code> option in the <code>T2P_JAVA_OPTIONS</code> environment variable. For example: <pre>setenv T2P_JAVA_OPTIONS "-d64 -Djava.io.tmpdir=/home/t2p/temp"</pre>	Optional
<code>-archiveLoc</code>	<code>-al</code>	The absolute path of the archive location. Use this option to specify the location of the archive file created by the <code>copyConfig</code> script.	Mandatory
<code>-targetNMHomeLoc</code>	<code>-tnh</code>	The absolute path of the target Node Manager.	Mandatory
<code>-targetMWHomeLoc</code>	<code>-tmw</code>	The absolute path of the target Middleware home in which the copy of Node Manager will be applied.	Mandatory

Table 17–30 (Cont.) Options for the pasteConfig Script for Node Manager

Option	Shortcut	Description	Mandatory or Optional?
-movePlanLoc	-mpl	The absolute path of the modified move plan in the target environment.	Mandatory
-logDirLoc	-ldl	The location of an existing directory. A new log file is created in the directory.	Optional
-silent	None	Specifies whether the operation operates silently. That is, it does not prompt for confirmation. The default is that the operation prompts for confirmation. To specify that it not prompt for confirmation, specify this option with the value of true.	Optional

17.6.7.4 Running pasteConfig for Oracle BI Enterprise Edition

You paste the configuration of the following Oracle BI Enterprise Edition components:

- Oracle BI Server
- Oracle BI Presentation Services
- Oracle BI Cluster Controller
- Oracle BI Scheduler
- JavaHost
- Oracle Essbase server, if it is installed in your environment

Running `pasteConfig` applies the copy of Oracle BI Enterprise Edition to the target environment by pasting the configuration files of the source Oracle BI Enterprise Edition into the specified Oracle BI Enterprise Edition home.

Use the following syntax to run `pasteConfig` for Oracle BI Enterprise Edition on each machine within the topology in the target environment:

```
pasteConfig -javaHome path_of_jdk
            -archiveLoc archive_location
            -targetInstanceHomeLoc trgt_Node_Manager_Home_path
            -targetInstanceName trgt_Instance_name
            -targetOracleHomeLoc trgt_Middleware_Home_path
            -domainHostName domain_host_name
            -domainPortNum domain_port_number
            -domainAdminUserName domain_admin_user_name
            -domainAdminPassword domain_admin_password_file
            [-silent {true | false}]
            [-logDirLoc log_dir_path]
```

The following example shows how to apply the copy of Oracle BI Enterprise Edition to the Oracle BI Enterprise Edition home located in `/net/mount1/appbase/instance/BIInstance` on UNIX:

```
pasteConfig -javaHome USER_HOME/jrockit_160_17_R28.0.0-679/
            -archiveLoc /FIN_T2P/BIInstance.jar
            -targetInstanceHomeLoc /net/mount1/appbase/instance/BIInstance
            -targetInstanceName BIInstance
            -targetOracleHomeLoc /net/mount1/appbase/fusionapps/bi
            -domainHostName server.us.example.com
            -domainPortNum 10201
            -domainAdminUserName faadmin
            -domainAdminPassword /tmp/welcome1.txt -silent true
```

Table 17–31 describes the options for the `pasteConfig` script for Oracle BI Enterprise Edition.

Table 17–31 Options for the `pasteConfig` Script for Oracle BI Enterprise Edition

Option	Shortcut	Description	Mandatory or Optional?
<code>-javaHome</code>	None	The absolute path of the Java Developer's Kit.	Mandatory
<code>-archiveLoc</code>	<code>-al</code>	The absolute path of the archive location. Use this option to specify the location of the archive file created by the <code>copyConfig</code> script.	Mandatory
<code>-targetInstanceHomeLoc</code>	<code>-tih</code>	The absolute path of the target Oracle instance. If the Oracle instance does not exist at that location, the command creates the instance	Mandatory
<code>-targetInstanceName</code>	<code>-tin</code>	The name of the target Oracle instance. The name must be unique in the domain.	Mandatory
<code>-targetOracleHomeLoc</code>	<code>-toh</code>	The absolute path of the target Oracle home. The target Oracle home must exist and it must contain the binaries for the component you are moving.	Mandatory
<code>-logDirLoc</code>	<code>-ldl</code>	The location of an existing directory. A new log file is created in the directory.	Optional
<code>-silent</code>	None	Specifies whether the operation operates silently. That is, it does not prompt for confirmation. The default is that the operation prompts for confirmation. To specify that it not prompt for confirmation, specify this option with the value of <code>true</code> .	Optional
Domain-Detail Options		These parameters are optional. If you choose to use the domain parameters, you must use all four parameters.	
<code>-domainHostName</code>	<code>-dhn</code>	The name of the host on which the domain is configured.	Mandatory
<code>-domainPortNum</code>	<code>-dpn</code>	The port number of the source domain.	Mandatory
<code>-domainAdminUserName</code>	<code>-dau</code>	The name of the administrative user for the source domain.	Mandatory
<code>-domainAdminPassword</code>	<code>-dap</code>	The password file location for the administrative user for the domain.	Mandatory

17.6.7.5 Running `pasteConfig` for Oracle Global Order Promising Server Instance

Running `pasteConfig` applies the copy of the Oracle Global Order Promising Server configuration files on the source environment to the target environment.

Use the following syntax to run `pasteConfig` for Oracle Global Order Promising Server on each computer within the topology in the target environment:

```
pasteConfig -javaHome path_of_jdk
            -archiveLoc archive_location
            -movePlanLoc move_plan_path
            -targetInstanceHomeLoc target_instance_path
            -movePlanLoc move_plan_path
            -targetComponentName GlobalOrderPromisingServer1
            -targetInstanceHomeLoc trgt_Instance_path
```

```
-targetInstanceName gop_1
-targetOracleHomeLoc trgt_ORACLE_HOME_path
-domainHostName domain_host_name
-domainPortNum domain_port_number
-domainAdminUserName domain_admin_user_name
-domainAdminPassword domain_admin_password_file
[-logDirLoc existing_log_directory_location]
[-silent {true | false}]
```

The following example shows a sample pasteconfig operation configuration for an Oracle Global Order Promising Server instance:

```
pasteConfig -javaHome USER_HOME/jrockit_160_17_R28.0.0-679/
-archiveloc USER_HOME/gop_1.jar
-movePlanLoc USER_HOME/t2pPlans/moveplan.xml
-targetComponentName GlobalOrderPromisingServer1
-targetInstanceHomeLoc /net/mount1/appbase/instance/gop_1
-targetInstanceName gop_1
-targetOracleHomeLoc /net/mount1/appbase/fusionapps/gop
-logDirLoc USER_HOME/clone_logs
-silent false
-domainHostName server.us.example.com
-domainPortNo 7051
-domainAdminUserName weblogic
-domainAdminPassword USER_HOME/p1.txt
```

Table 17–32 describes the options for the pasteconfig script for Oracle Global Order Promising Server.

Table 17–32 Options for the pasteConfig Script for Oracle Global Order Promising Server

Option	Shortcut	Description	Mandatory or Optional?
-javaHome	None	The absolute path of the Java Developer's Kit.	Mandatory
-archiveloc	-al	The absolute path of the archive location. Use this option to specify the location of the archive file created by the copyConfig script.	Mandatory
-movePlanLoc	-mpl	The absolute path of the modified move plan in the target environment.	Mandatory
-targetComponentName	-tcn	The name of the existing component which will be archived, GlobalOrderPromisingServer1.	Mandatory
-targetInstanceHomeLoc	-tin	The absolute path of the Oracle Global Order Promising Server instance for the source component.	Mandatory
-logDirLoc	-ldl	The location of an existing directory. A new log file is created in the directory.	Optional
-silent	None	Specifies whether the operation operates silently. That is, it does not prompt for confirmation. The default is that the operation prompts for confirmation. To specify that it not prompt for confirmation, specify this option with the value of true.	Optional
Domain-Detail Options			
-domainHostName	-dhn	The name of the host on which the domain is configured.	Mandatory

Table 17–32 (Cont.) Options for the pasteConfig Script for Oracle Global Order Promising Server

Option	Shortcut	Description	Mandatory or Optional?
-domainPortNum	-dpn	The port number of the source domain.	Mandatory
-domainAdminUserName	-dau	The name of the administrative user for the domain. For example, faadmin.	Mandatory
-domainAdminPassword	-dap	The password file location for the administrative user for the domain.	Mandatory

17.6.8 Completing the Component-Specific Configuration Move

Moving the Oracle WebLogic Server domains from a source environment to a target environment also moves some Oracle Fusion Middleware components, but additional tasks are required for completing the move.

- Oracle BI EE. See [Section 17.6.8.1](#).
- Oracle BI Applications. See [Section 17.6.8.2](#).
- Oracle Governance, Risk and Compliance Controls. See [Section 17.6.8.3](#), "Moving the Oracle Governance, Risk and Controls Client Configuration".
- Oracle SES. See [Section 17.6.8.4](#).
- Oracle WebLogic Communication Services. See [Section 17.6.8.5](#).

17.6.8.1 Moving Oracle BI Enterprise Edition Components

To complete the movement of Oracle BI EE, perform the following tasks:

- [Task 1, "Update the App ID passwords"](#)
- [Task 2, "Enable New Agents"](#)

Task 1 Update the App ID passwords

If the passwords for following App IDs are not properly provided in move plan, you must the update target environment with the new passwords:

- FUSION_APPS_PROV_PATCH_APPID
- FUSION_APPS_PRC_BI_APPID
- FUSION_APPS_BI_APPID

To update the APP ID passwords:

1. Obtain the APP ID passwords through the credential store. Execute the following Execute following listCred commands to retrieve passwords for FUSION_APPS_PROV_PATCH_APPID, FUSION_APPS_PRC_BI_APPID, and FUSION_APPS_BI_APPID .

- a. At an operating system command prompt, navigate to the location of the WLST script:

```
(UNIX) FA_MW_HOME/oracle_common/common/bin
(Windows) FA_MW_HOME\oracle_common\common\bin
```

FA_MW_HOME is named fusionapps and located under the *APPLICATIONS_BASE*.

- b. Run the wlst script:

```
(UNIX) wlst.sh
(Windows) wlst.cmd
```

- c. Connect to Oracle WebLogic Server with the following WLST command:

```
connect([options])
```

See the "connect" section in the *Oracle Fusion Middleware WebLogic Scripting Tool Command Reference*.

- d. Use the following WLST commands from the fusionapps Middleware directory:

```
listCred(map="oracle.patching", key="FUSION_APPS_PATCH_WLS_ADMIN-KEY")
listCred(map="oracle.apps.security", key="FUSION_APPS_PRC_BI_APPID-KEY")
listCred(map="oracle.apps.security", key="FUSION_APPS_BI_APPID-KEY")
```

See the "listCred" section in the *Oracle Fusion Middleware WebLogic Scripting Tool Command Reference*.

2. Change JDBC Data for the source configuration for FUSION_APPS_PROV_PATCH_APPID and FUSION_APPS_PRC_BI_APPID .

In Oracle Fusion Applications, there are two data sources which are configured using APP IDs:

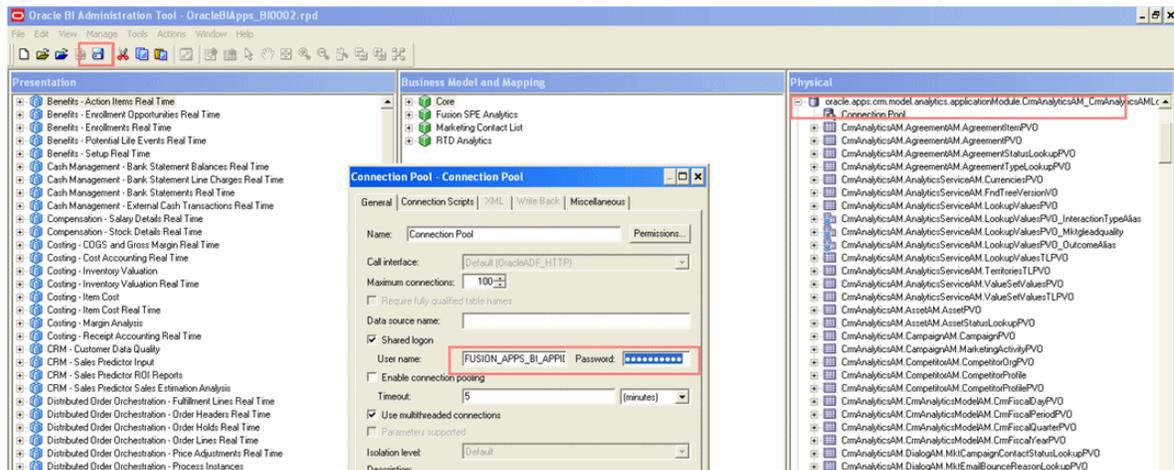
- BIDomain: BIANalytics data-source
User: FUSION_APPS_PROV_PATCH_APPID
- ProcurementDomain: BIANalyticsServer data-source
User: FUSION_APPS_PRC_BI_APPID

To change JDBC data for the source configuration through the Oracle WebLogic Server Administration Console for the BIDomain and ProcurementDomain domains:

- a. Locate the **Change Center** in the upper left of the Administration Console screen
 - b. Click **Lock & Edit** to lock the configuration edit hierarchy for the domain.
 - c. From the **Domain Structure** pane, choose **Services > Data Sources**.
The Summary of JDBC Data Sources page displays.
 - d. In the **Data Sources** table, click the **BIAnalytics** data source for the BIDomain domain or **BIAnalyticsServer** data source for the ProcurementDoamin domain.
 - e. Click the **Connection Pool** sub-tab.
 - f. Locate the **Password** and **Confirm Password** fields, and modify the passwords.
 - g. Click **Save**.
 - h. Restart all the data source's target Managed Servers.
3. Change repository publishing directory connection pools through the Oracle BI Administration Tool.
 - a. Choose **Start > Programs > Oracle Business Intelligence > BI Administration**.
 - b. In the Administration Tool, choose **File > Open > Offline**.
 - c. Navigate to the repository you want to open, and then select **Open**.
 - d. In the Open Offline dialog box, type a valid user ID and password, and then click **OK**.

This opens the repository in offline mode. This mode enables you to view and modify a repository while it is not loaded into the Oracle BI Server. If you attempt to open a repository in offline mode while it is loaded into the Oracle BI Server, the repository opens in read-only mode. Only one Administration Tool session at a time can edit a repository in offline mode.

- e. In **Physical** layer, select an Oracle Fusion application, for example, `oracle.apps.crm*`, `oracle.apps.fscrm*`, `oracle.apps.hcm*`.
- f. Click the connection pool icon in the toolbar.



- g. In the Connection Pool dialog, update the new password in the Connection Pool dialog, and then click **OK**.
- h. Perform sub-steps e through g for other Oracle Fusion applications.

For more information about the Oracle BI administration Tool, see the section "Before You Begin" in the *Oracle Fusion Middleware Metadata Repository Builder's Guide for Oracle Business Intelligence Enterprise Edition (Oracle Fusion Applications Edition)*.

- 4. Upload the newly update repository publishing directory to the BIDomain domain. See the section "Configuring Repositories" in the *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

Task 2 Enable New Agents

If required, enable new agents. If new agents were created in the source environment, click each agent in the Oracle BI Presentation Services Catalog Manager (in the target environment) to enable it. For more information about enabling new agents, see the "Working with the Properties of Catalog Objects" section of the *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

17.6.8.2 Moving Oracle BI Applications Components

To complete the movement of Oracle BI Applications, see the section "Moving Oracle BI Applications Using the Full-Movement Scenario," in the *Oracle Fusion Middleware Reference Guide for Oracle Business Intelligence Applications*.

17.6.8.3 Moving the Oracle Governance, Risk and Controls Client Configuration

If you have custom provisioned Oracle Governance, Risk and Compliance Controls, use an Oracle Fusion Functional Setup Manager task in the target environment to move the Oracle Governance, Risk and Compliance Controls client configuration to the target Oracle Identity Manager server.

To move the GRC configuration:

1. Sign into Oracle Fusion Applications using a user account provisioned with an enterprise role that inherits the GRC Setup Management Duty role. For example, in the security reference implementation, the IT Security Manager and Application Implementation Consultant job roles inherit this duty role. Contact your security administrator to ensure your user account is provisioned appropriately.
2. In Oracle Fusion Functional Setup Manager, navigate to the **Manage Application Access Controls** task and select a GRC setup configuration.

The **Manage Application Access Controls** task is included in the **Define Governance, Risk and Performance Controls** task list, which is visible under the **Common Applications...** task list only if the offering was configured to include the Governance, Risk and Compliance feature prior to the creation of the implementation project.

You can also manually add the task to the implementation project.

3. Select the desired configuration detail (the one appropriate for the target environment). If none is appropriate, create a new configuration detail for the target environment. For information, see the *Oracle Fusion Applications Information Technology Management, Implement Applications Guide*.
4. Perform the **Publish to OIM Server** action using the configuration detail. This requires the connection information for the Oracle WebLogic Server instance that is hosting the Oracle Identity Manager server.

17.6.8.4 Moving Oracle SES

Source types of portal, email, and mailing list are not moved from the source Oracle SES instance to the target Oracle SES instance. If you purchased an additional Oracle SES license and want to enable search non-Fusion sources, you must create these data sources in the target Oracle SES instance manually using the Oracle SES Administration GUI in the target Oracle SES instance. For information, see the *Oracle Secure Enterprise Search Administrator's Guide*.

17.6.8.5 Moving Oracle WebLogic Communication Services

Configure the following Oracle WebLogic Communication Services properties on the target environment:

- `pAssertedIdentityUri`
- `controllerUri`
- `preloadedRouteUri`

Configure these attributes with the `ThirdPartyCallConfiguration` MBean. See [Section 8.2](#) for configuration instructions.

17.6.9 Applying Domain Configurations to Other Machines

The domain directory is local to each machine, and `pasteConfig` is performed only on the Administration Server. Subsequently, if the domain has multiple machines, you

must re-create the local domain by using the Oracle WebLogic Server pack and unpack commands. For more information, see *Oracle Fusion Middleware Creating Templates and Domains Using the Pack and Unpack Commands*.

Note: If the Administration Server is using the shared storage, and the other Managed Servers on the Administration Server machine are using the local storage, then you must run the pack and unpack commands once even on the Admin Server machine to create the domain directory on the local storage.

17.7 Moving Binary and Configuration Files to a Non-Distributed Target Environment

Some environments are non-distributed, whereby the Administration Servers of all the domains and all the system component instances run on the same host. For these environments, you can automate moving binary and configuration files with a different set of these steps than those provided in [Section 17.5](#) and [Section 17.6](#). Please note these scripts are not intended for production, distributed environments, whereby the topology is located on multiple computers. The scripts for a non-distributed environment are intended for moving binary files and configuration files from one QA environment to another QA environment.

These scripts run in nonparallel mode rather. For example, while the CopyConfig command moves the configuration archives in one command, it runs the operation in nonparallel mode. Therefore, the operation can take some time to complete all the moves.

To move the binary files and configuration files from the source environment to the target environment:

1. Set the following variables on both the source and the target:
 - HOME: Location of the user's home directory
 - APPTOP: Location of the applications base, the top-level directory where the Oracle Fusion Applications binaries are located
 - JAVA_HOME: Location of the JDK
 - T2P_DIR: Location of the working directory for the movement scripts. You do not need to use this variable if you use the `-workdir` command is specified when running the commands.
2. For either UNIX or Windows, at an operating system command prompt, navigate to the following directories on both the source and the target:

```
(UNIX) FA_ORACLE_HOME/applications/lcm/t2p/bin
(Windows) FA_ORACLE_HOME\applications\lcm\t2p\bin
```

`FA_ORACLE_HOME` is located under the fusionapps Oracle Fusion Middleware home directory.

3. Use the following command to copy the binary files from the source environment to the target environment:

```
java -DT2PAdmin=user -DT2PAdminPassword=password -classpath
T2PAutomation.jar:cloningclient.jar:obfuscatepassword.jar:ojmisc.jar
oracle.t2p.automation.T2PAutomationMain -archive archiveLocation
-task CopyBinary -workdir working_dir
```

The options for the CopyBinary command follow:

Option	Description
-DT2PAdmin	The Oracle Fusion Middleware administration user.
-DT2PAdminPassword	The password for the Oracle Fusion Middleware administration user.
-classpath	The dependent library files containing classes required to run the command. Use the following value for the CopyBinary command: T2PAutomation.jar:cloningclient.jar:obfuscatepassword.jar:ojmisc.jar oracle.t2p.automation.T2PAutomationMain
-archive	<p>The absolute path of the archive location. Use this option to specify the location for the T2Pcontrol.dat file to be created with the CopyBinary script.</p> <p>The archive location must not already exist, but its parent directory must exist and have write permission. Ensure that the archive location is not within the Middleware home structure.</p> <p>The T2Pcontrol.dat file has the following structure:</p> <pre>task=copyConfig, domain=BI Domain, jarfile=configBIDomain.jar, order=2 task=copyConfig, subtask=NodeManager, jarfile=configNodeManager.jar, order=4 task=copyConfig, subtask=OHS, jarfile=configOHS.jar, order=3 task=copyConfig, subtask=BIInstance, jarfile=configBIInstance.jar, order=3 task=copyConfig, domain=CommonDomain, jarfile=configCommonDomain.jar, order=1 task=copyConfig, domain=FinancialDomain, jarfile=configFinancialDomain.jar, order=2 task=copyConfig, domain=ProjectsDomain, jarfile=configProjectsDomain.jar, order=2 task=copyConfig, domain=ProcurementDomain, jarfile=configProcurementDomain.jar, order=2 task=copyConfig, domain=SCMDomain, jarfile=configSCMDomain.jar, order=2 task=copyConfig, domain=HCMDomain, jarfile=configHCMDomain.jar, order=2 task=copyConfig, domain=CRMDomain, jarfile=configCRMDomain.jar, order=2 task=copyBinary, jarfile=cloneBinary.jar, order=1</pre> <p>Data from this file will be used in PasteBinary phase.</p>
-task	The command to execute. Enter CopyBinary.
-workdir	The directory used for log files and any other temporary files generated by the automation framework. The log files are located at <i>working_dir</i> \logs (UNIX) or <i>working_dir</i> \logs folder. This command is not needed if the T2P_DIR environment variable is set.

This command moves the following from the source environment to the target environment:

- Binaries from the fusionapps Middleware home directory:
 - (UNIX) APPLICATIONS_BASE/fusionapps
 - (Windows) APPLICATIONS_BASEE\fusionapps

- Binaries from the `webtier_mwhome` Middleware home directory:

(UNIX) `APPLICATIONS_BASE/webtier_mwhome`
 (Windows) `APPLICATIONS_BASEE\webtier_mwhome`

- Files from the Oracle Database client home:

(UNIX) `APPLICATIONS_BASE/dbclient`
 (Windows) `APPLICATIONS_BASEE\dbclient`

- Binay from the Webgate Oracle home:

(UNIX) `APPLICATIONS_BASE/webgate`
 (Windows) `APPLICATIONS_BASEE\webgate`

where `APPLICATIONS_BASE (/net/mount1/appbase)` represents the top-level applications base directory for binaries. For information about the Oracle Fusion Applications directory structure, see [Section 1.2.3](#).

4. Use the following command to unpack all binary files in the target environment:

```
java -classpath
T2PAutomation.jar:cloningclient.jar:obfuscatepassword.jar:ojmisc.jar
oracle.t2p.automation.T2PAutomationMain
-archive archiveLocation -control T2PControl.dat_location
-task PasteBinary -workdir working_dir
```

The options for the `CopyBinary` command follow:

Options	Description
<code>-classpath</code>	The dependent library files containing classes required to run the command. Use the following value for the <code>CopyBinary</code> command: <code>T2PAutomation.jar:cloningclient.jar:obfuscatepassword.jar:ojmisc.jar oracle.t2p.automation.T2PAutomationMain</code>
<code>-archive</code>	The absolute path of the archive location. Use this option to specify the location of the <code>T2Pcontrol.dat</code> file created by the <code>CopyBinary</code> command. The location must exist. Ensure that the archive location is not within the Middleware home structure.
<code>-task</code>	The command to execute. Enter <code>PasteBinary</code> .
<code>-workdir</code>	The directory used for log files and any other temporary files generated by the automation framework. The log files are located at <code>working_dir\logs</code> (UNIX) or <code>working_dir\logs</code> folder. This command is not needed if the <code>T2P_DIR</code> environment variable is set.

5. Extract the move plan by running the `extractMovePlan` script on the list of configuration archive locations, separated by comma. See [Section 17.6.5](#).
6. Modify the move plan to specify properties for the target environment. See [Section 17.6.6](#).
7. Create the configuration archive from the source environment by running the `CopyConfig` command in the source environment:

```
java -classpath
T2PAutomation.jar:cloningclient.jar:obfuscatepassword.jar:ojmisc.jar
oracle.t2p.automation.T2PAutomationMain -archive archiveLocation
-task CopyConfig -workdir working_dir
```

The options for the CopyConfig command follow:

Option	Description
-classpath	The dependent library files containing classes required to run the command. Use the following value for the CopyConfig command: T2PAutomation.jar:cloningclient.jar:obfuscatepassword.jar:ojmisc.jar oracle.t2p.automation.T2PAutomationMain
-archive	The absolute path of where all the JARs will be generated. Use this option to specify the location of the T2Pcontrol.dat file to be created with the CopyConfig script. The archive location must not already exist, but its parent directory must exist and have write permission. Ensure that the archive location is not within the Middleware home structure. Data from this file will be used in PasteConfig phase.
-task	The command to execute. Enter CopyConfig.
-workdir	The directory used for log files and any other temporary files generated by the automation framework. The log files are located at <i>working_dir</i> \logs (UNIX) or <i>working_dir</i> \logs folder. This command is not needed if the T2P_DIR environment variable is set.

8. Apply the copied configurations from the source environment to the target environment by running the PasteConfig command:

```
java -classpath
T2PAutomation.jar:cloningclient.jar:obfuscatepassword.jar:ojmisc.jar
oracle.t2p.automation.T2PAutomationMain
-archive archiveLocation
-control T2PControl.dat location -task PasteConfig
-plan moveplan -bihost bihost -workdir working_dir
```

The options for the PasteConfig command follow:

Options	Description
-classpath	The dependent library files containing classes required to run the command. Use the following value for the CopyBinary command: T2PAutomation.jar:cloningclient.jar:obfuscatepassword.jar:ojmisc.jar oracle.t2p.automation.T2PAutomationMain
-archive	The absolute path of the archive location. Use this option to specify the location of the T2Pcontrol.dat file created by the CopyConfig command. The directory must contain all files produced by the copyConfig command and referenced in the T2PControl.dat file. Ensure that the archive location is not within the Middleware home structure.
-task	The command to execute. Enter PasteConfig.
-plan	The absolute path of the modified move plans in the target environment.
-bihost	The host name of the target.

Options	Description
-workdir	The directory used for log files and any other temporary files generated by the automation framework. The log files are located at <i>working_dir</i> \logs (UNIX) or <i>working_dir</i> \logs folder. This command is not needed if the T2P_DIR environment variable is set.

9. Validate the domain, Managed Servers, Oracle HTTP Server, and Node Manager server status by running the Validate command:

```
java -classpath
T2PAutomation.jar:cloningclient.jar:$APPTOP/fusionapps/wlserver_
10.3/server/lib/weblogic.jar:obfuscatepassword.jar:ojmisc.jar
oracle.t2p.automation.T2PAutomationMain -task Validate -props property_file
-hostfqdn host -workdir working_dir
```

The options for the Validata command follow:

Option	Description
-classpath	The dependent library files containing classes required to run the command. Use the following value for the PostMovement command: T2PAutomation.jar:cloningclient.jar:\$APPTOP/fusionapps/wlserver_10.3/server/lib/weblogic.jar:obfuscatepassword.jar:ojmisc.jar oracle.t2p.automation.T2PAutomationMain
-task	The command to execute. Enter Validate.

Option	Description
-props	<p>The absolute path to the properties file name. This file provides all the required user names, passwords and has the following structure:</p> <pre>#Properties file for T2P post movement setup parameters #Script location, needs to be the absolute path to scripts folder, inclusive. scripts-folder=/scratch/aime/work/t2p/scripts #Node manager information NMUser=bootstrap_admin NMPassword=welcome1 NMPort=5556 #Process id of node manager process, required on unix/linux NMPid=<pid> #Service name of nodemanager, required on windows NMSvcName=<Node Manager Service> #With the following true, reconcile GUID operation is run on this host. For other hosts set this to false. run-reconcile-guid=true #Path to domain from which to use the jps-config-jse.xml reconcile-guid-domain-path=/scratch/aime/work/APPTOP/insta nce/domains/adcdat10.us.oracle.com/CommonDomain #Database info to run reconcile GUID dburl=jdbc:oracle:thin:@adcdaf12.us.oracle.com:5593:DB5593 dbuser=fusion dbpassword=<dbpassword> #If following 2 props are set, then it is assumed use name and passwords are the same for all domains. AdminUser=FAAdmin AdminPassword=fusionfa1 #Domain admin user credentials, if the above is not set CommonDomainAdminUser=FAAdmin CommonDomainAdminPassword=fusionfa1 #Put all other domain user name and password below, if needed.</pre>
-hostfqdn	The current host name where this operation runs.
-workdir	The directory used for log files and any other temporary files generated by the automation framework. The log files are located at <i>working_dir</i> \logs (UNIX) or <i>working_dir</i> \logs folder. This command is not needed if the T2P_DIR environment variable is set.

10. Complete the component-specific configuration move by performing additional movement tasks on the Oracle Fusion Middleware components. See [Section 17.6.8](#).
11. If there are multiple machines for the domain, recreate the local domain directory using the Oracle WebLogic Server pack and unpack mechanisms. See [Section 17.6.9](#).

17.8 Moving Functional Setup Data

Functional Setup Manager enables users to implement and configure all Oracle Fusion Applications as well as to move the setup data across instances. An Application Implementation Consultant user can move the setup data using Functional Setup Manager export and import functionality. This functionality can be accessed from a command line.

To move setup data from the source environment to the target environment using command line scripts:

- [Task 1, "Prerequisite"](#)
- [Task 2, "Export Data from the Source"](#)
- [Task 3, "Import Data on the Target"](#)

Task 1 Prerequisite

To prepare the environment prior to running the scripts:

1. Identify the implementation projects used to enter and define the setup data in the source environment.

Before you move the Functional Setup Data to the target environment, you must identify the implementation projects used to enter and define the setup data in the source environment. Ask your Application Implementation Consultant user to gather the codes of the implementation projects using the following procedure for each project:

You can gather the implementation codes using Oracle Fusion Functional Setup Manager or SQL.

To gather the implementation codes using Oracle Fusion Functional Setup Manager:

- a. Sign in to Oracle Fusion Applications with a user account that is provisioned with the necessary role, such as the predefined Application Implementation Consultant role. Contact your security administrator for details.
- b. From the **Administration** menu in the work area of Oracle Fusion Applications, choose **Setup and Maintenance**.
- c. In the **Tasks** pane, under the **Implementations** group, click **Manage Implementation Projects**.
- d. Click an implementation project name in the **Search Results**.
The Implementation Project page displays.
- e. In the **Basic Information** section, copy the value for the **Code** field.
- f. Perform Steps 4 and 5 for each project.

To obtain a list of the implementation codes using SQL, use the following query:

```
SQL> SELECT impl_project_name, short_name from asm_impl_projects_v1;
```

2. Ensure all the Managed Servers are started in the environment. See [Section 3.3.2.1, Task 4, "Start the Managed Servers"](#).

Task 2 Export Data from the Source

To export the data from the source:

1. Set the USER_MEM_ARG variable as follows:

```
export USER_MEM_ARGS="-Xms256m -Xmx1024m -XX:CompileThreshold=8000
-XX:PermSize=128m -XX:MaxPermSize=1024m"
```

- For either UNIX or Windows, at an operating system command prompt, navigate to the location of the `fsmConfigurationPackageExport` executable:

```
(UNIX) ATGPF_ORACLE_HOME/fsm/scripts
(Windows) ATGPF_ORACLE_HOME\fsm\scripts
```

`ATGPF_ORACLE_HOME` is located under the fusionapps Oracle Fusion Middleware home directory (`FA_MW_HOME`).

- Enter `fsmConfigurationPackageExport.sh` (for UNIX) or `fsmConfigurationPackageExport.cmd` (for Windows), using the following syntax:

```
fsmConfigurationPackageExport
  commonDomainInternalUrl=http://host.domain:port
  jdbcUrl=jdbc:oracle:thin:@DB_host_name.domain/sid
  fusionRuntimeSchema=fusion_runtime
  fusionRuntimeSchemaPasswordFile=USER_HOME/fusionRuntimeSchemaPassword.txt
  fsmUsername=user
  fsmPasswordFile=USER_HOME/fusionRuntimeSchemaPassword.txt
  implementationProjectCode=IMPLEMENTATION_PROJECT_CODE_TO_EXPORT
  configurationPackageFile=USER_HOME/cp.zip
  logFile=USER_HOME/log_file
  jpsConfig=COMMON_DOMAIN_HOME/config/fmwconfig/jps-config-jse.xml
  mwHomeLoc=MW_HOME
  atgpfHomeLoc=ATGPF_ORACLE_HOME
  [pollingInterval=interval]
  [pollingLimit=limit]
```

The following example shows how to export the configuration package from the source instance for an implementation project named `FINANCIALS_IMPLEMENTATION`:

```
fsmConfigurationPackageExport.sh
commonDomainInternalUrl=http://fusionssystemtest-w-internal-fs.example.com:7011
jdbcUrl=jdbc:oracle:thin:@db.example.com:1528:proddb
fusionRuntimeSchema=fusion_runtime
fusionRuntimeSchemaPasswordFile=USER_HOME/fusionRuntimeSchemaPassword.txt
fsmUsername=FUSION
fsmPasswordFile=USER_HOME/fusionRuntimeSchemaPassword.txt
implementationProjectCode=FINANCIALS_IMPLEMENTATION
configurationPackageFile=USER_HOME/cp.zip
logFile=USER_HOME/cpExport.log
jpsConfig=COMMON_DOMAIN_HOME/config/fmwconfig/jps-config-jse.xml
mwHomeLoc=/net/mount1/appbase/fusionapps
atgpfHomeLoc=/net/mount1/appbase/atgpf
```

The following table describes the options for the `fsmConfigurationPackageExport` command.

Option	Description	Optional or Mandatory?
-commonDomainInternalUrl	<p>The CommonDomain domain internal URL on the source registered in Oracle Topology Manager in the format of <i>Internal_Server_Protocol://Internal_Server_Host:Internal_Server_Port</i>.</p> <p>To obtain the URL, use Oracle Fusion Functional Setup Manager:</p> <ol style="list-style-type: none"> 1. Sign in to Oracle Fusion Applications with a user account that is provisioned with the necessary role, such as the predefined Application Implementation Consultant role. Contact your security administrator for details. 2. From the Administration menu in the work area of Oracle Fusion Applications, choose Setup and Maintenance. 3. In the Tasks pane, under the Topology Registration group, click Register Domains. 4. In the Search Results section, click the CommonDomain entry. The Edit Domain: CommonDomain dialog displays. 5. Use the values from the Internal Server Protocol, Internal Server Host, and Internal Server Port fields to obtain the elements of the URL. 	Mandatory
-jdbcUrl	The JDBC URL in the format of <i>jdbc:oracle:thin:@//DB_host_name.domain/sid</i> to the connect to the source database.	Mandatory
-fusionRuntimeSchema	The Oracle database schema name (user name) to login to the source database. This value should always be <i>fusion_runtime</i> .	Mandatory
-fusionRuntimeSchemaPasswordFile	The absolute path to the password file, which contains the schema password to login to the source database.	Mandatory
-fsmUsername	User name to login to Oracle Fusion Functional Setup Manager that has Application Implementation Consultant job role.	Mandatory
-fsmPasswordFile	The absolute path to the password file, which contains password to login to Functional Setup Manager.	Mandatory
-implementationProjectCode	The implementation project code that needs to be exported from the source instance.	Mandatory
-configurationPackageFile	The absolute path to the configuration package ZIP file to save exported data.	Mandatory
-logFile	The absolute path to the source log file.	Mandatory
-jpsConfig	The absolute path to the <i>jps-config.xml</i> file.	Mandatory
-mwHomeLoc	The absolute location of the source Oracle Fusion Applications Middleware Home (<i>FA_MW_HOME</i>). <i>FA_MW_HOME</i> is located under the appbase Applications base directory (<i>APPLICATIONS_BASE</i>).	Mandatory
-atgpfHomeLoc	The absolute location of the source Applications Core Oracle Home (<i>ATGPF_ORACLE_HOME</i>). <i>ATGPF_ORACLE_HOME</i> is located under the <i>fusionapps</i> Oracle Fusion Middleware home directory (<i>FA_MW_HOME</i>).	Mandatory

Option	Description	Optional or Mandatory?
-clientPolicy	The Web Service Client Policy to be used to invoke the Oracle Enterprise Scheduler Web service to submit the export process. The default is oracle/wss10_saml_token_client_policy.	Optional
-pollingInterval	The polling interval to be used to periodically check if export process is complete. The default is 5 seconds.	Optional
-pollingLimit	The polling limit to be used until reached to check if the export process is complete. The default is 3600 seconds.	Optional

Task 3 Import Data on the Target

The `fsmConfigurationPackageImport` script moves the setup data for an implementation to a target instance.

To import the data on the target:

1. For either UNIX or Windows, at an operating system command prompt on the target, navigate to the location of the `fsmConfigurationPackageImport` executable:

```
(UNIX) ATGPF_ORACLE_HOME/fsm/scripts
(Windows) ATGPF_ORACLE_HOME\fsm\scripts
```

`ATGPF_ORACLE_HOME` is located under the fusionapps Oracle Fusion Middleware home directory (`FA_MW_HOME`).

2. Enter `fsmConfigurationPackageImport.sh` (for UNIX) or `fsmConfigurationPackageImport.cmd` (for Windows), using the following syntax:

```
fsmConfigurationPackageImport
commonDomainInternalUrl=http://host.domain:port
jdbcUrl=jdbc:oracle:thin:@DB_host_name.domain/sid
fusionRuntimeSchema=fusion_runtime
fusionRuntimeSchemaPasswordFile=USER_HOME/fusionRuntimeSchemaPassword.txt
fsmUsername=FUSION
fsmPasswordFile=USER_HOME/fusionRuntimeSchemaPassword.txt
configurationPackageFile=USER_HOME/cp.zip
logfile=USER_HOME/cpImport.log
jpsConfig=COMMON_DOMAIN_HOME/config/fmwconfig/jps-config-jse.xml
mwHomeLoc=MW_HOME
atgpfHomeLoc=ATGPF_ORACLE_HOME
```

The following example shows how to import the configuration package from the source instance:

```
fsmConfigurationPackageImport.sh
commonDomainInternalUrl=http://fusionsystemtest-w-internal-fs.example.com:7011
jdbcUrl=jdbc:oracle:thin:@db.example.com:1528:proddb
fusionRuntimeSchema=fusion_runtime
fusionRuntimeSchemaPasswordFile=USER_HOME/fusionRuntimeSchemaPassword.txt
fsmUsername=FUSION
fsmPasswordFile=USER_HOME/fusionRuntimeSchemaPassword.txt
configurationPackageFile=USER_HOME/cp.zip
logfile=USER_HOME/cpImport.log
jpsConfig=COMMON_DOMAIN_HOME/config/fmwconfig/jps-config-jse.xml
mwHomeLoc=/net/mount1/appbase/fusionapps
atgpfHomeLoc=/net/mount1/appbase/atgpf
```

The following table describes the options for the `fsmConfigurationPackageImport` command.

Option	Description	Optional or Mandatory?
<code>-commonDomainInternalUrl</code>	<p>The CommonDomain domain internal URL on the target registered in Oracle Topology Manager in the format of <code>Internal_Server_Protocol://Internal_Server_Host:Internal_Server_Port</code>.</p> <p>To obtain the URL, use Oracle Fusion Functional Setup Manager:</p> <ol style="list-style-type: none"> 1. Sign in to Oracle Fusion Applications with a user account that is provisioned with the necessary role, such as the predefined Application Implementation Consultant role. Contact your security administrator for details. 2. From the Administration menu in the work area of Oracle Fusion Applications, choose Setup and Maintenance. 3. In the Tasks pane, under the Topology Registration group, click Register Domains. 4. In the Search Results section, click the CommonDomain entry. The Edit Domain: CommonDomain dialog displays. 5. Use the values from the Internal Server Protocol, Internal Server Host, and Internal Server Port fields to obtain the elements of the URL. 	Mandatory
<code>-jdbcUrl</code>	The JDBC URL in the format of <code>jdbc:oracle:thin:@//DB_host_name.domain/sid</code> to connect to the database.	Mandatory
<code>-fusionRuntimeSchema</code>	The Oracle database schema name (user name) to login to the target database. This value should always be <code>fusion_runtime</code> .	Mandatory
<code>-fusionRuntimeSchemaPasswordFile</code>	The absolute path to the password file, which contains the schema password to login to the target database.	Mandatory
<code>-fsmUsername</code>	User name to login to Oracle Fusion Functional Setup Manager that has Application Implementation Consultant job role.	Mandatory
<code>-fsmPasswordFile</code>	The absolute path to the password file, which contains password to login to Functional Setup Manager.	Mandatory
<code>-fsmPasswordFile</code>	The absolute path to the password file, which contains password to login to Functional Setup Manager.	Mandatory
<code>-configurationPackageFile</code>	The absolute path to the configuration package ZIP file to save exported data.	Mandatory
<code>-logFile</code>	The absolute path to the target log file.	Mandatory
<code>-jpsConfig</code>	The absolute path to the <code>jps-config.xml</code> file.	Mandatory
<code>-mwHomeLoc</code>	The absolute location of the target Oracle Fusion Applications Middleware Home (<code>FA_MW_HOME</code>). <code>FA_MW_HOME</code> is located under the <code>appbase</code> Applications base directory (<code>APPLICATIONS_BASE</code>).	Mandatory
<code>-atgpfHomeLoc</code>	The absolute location of the target Applications Core Oracle Home (<code>ATGPF_ORACLE_HOME</code>). <code>ATGPF_ORACLE_HOME</code> is located under the <code>fusionapps</code> Oracle Fusion Middleware home directory (<code>FA_MW_HOME</code>).	Mandatory

Option	Description	Optional or Mandatory?
-clientPolicy	The Web Service Client Policy to be used to invoke the Oracle Enterprise Scheduler Web service to submit the import process. The default is oracle/wss10_saml_token_client_policy.	Optional
-pollingInterval	The polling interval to be used to periodically check if import process is complete. The default is 5 seconds.	Optional
-pollingLimit	The polling limit to be used until reached to check if the import process is complete. The default is 3600 seconds.	Optional

17.9 Completing Post-Movement Tasks

This section covers the following topics related to post-movement tasks:

- [Completing Post-Movement Tasks for a Target Distributed Environment](#)
- [Tips for Completing the Oracle RAC Configuration in the Target Environment](#)

17.9.1 Completing Post-Movement Tasks for a Target Distributed Environment

After you move the Oracle Fusion Applications components across environments, you must complete the following tasks:

This section covers the following topics:

- [Completing Post-Movement Tasks for a Target Distributed Environment](#)
- [Completing Post-Movement Tasks on a Target Non-Distributed Environment](#)

For production, distributed environments, whereby the topology is located on multiple computers, use [Section 17.9.1.1](#). Some environments are non-distributed, whereby the topology in its entirety is on the same machine or a shared file system. For these environments, you can automate the post-movement tasks, as described in [Section 17.9.1.2](#).

17.9.1.1 Completing Post-Movement Tasks for a Target Distributed Environment

In a distributed environment, after you move the Oracle Fusion Applications components across environments, you must complete the following tasks:

- Reconcile GUIDs security policies.
Since GUIDs are not preserved during movement to a target environment, you must reconcile the GUIDs in the security policies. See [Section 4.4.3](#).
- Apply any necessary topology changes, such as scaling up or scaling out the topology. For information, see [Section 16.5](#).
- Complete the Oracle RAC configuration in the target environment. See [Section 17.9.2](#).
- Start the Managed Servers as needed. See [Section 3.3.2.1, Task 4, "Start the Managed Servers"](#).
- Start a full index on all index schedules. For information, see [Section 7.4.2.7](#).

17.9.1.2 Completing Post-Movement Tasks on a Target Non-Distributed Environment

In a non-distributed environment, after you move the Oracle Fusion Applications components across environments, you must complete the following tasks:

- Reconcile GUIDs security policies.
- Stop the Administration Server and Node Manager.
- Restart the Administration Server and Node Manager.

To complete the post-movement tasks:

1. For either UNIX or Windows, at an operating system command prompt, navigate to the following locations on the target:

```
(UNIX) FA_ORACLE_HOME/applications/lcm/t2p/bin
(Windows) FA_ORACLE_HOME\applications\lcm\t2p\bin
```

FA_ORACLE_HOME is located under the fusionapps Oracle Fusion Middleware home directory.

2. Use the following command to perform post-movement tasks on the target environment:

```
java -classpath
T2PAutomation.jar:cloningclient.jar:obfuscatepassword.jar:ojmisc.jar
oracle.t2p.automation.T2PAutomationMain
-task PostMovement -hostfqdn host -props properties_file -workdir working_dir
```

[Table 17–33](#) describes the options for the `PostMovement` command.

Table 17–33 Options for the `PostMovement` Command

Options	Description
<code>-classpath</code>	The dependent library files containing classes required to run the command. Use the following value for the <code>PostMovement</code> command: <code>T2PAutomation.jar:cloningclient.jar:obfuscatepassword.jar:ojmisc.jar oracle.t2p.automation.T2PAutomationMain</code>
<code>-task</code>	The command to execute. Enter <code>PostMovement</code> .
<code>-hostfqdn</code>	The current host name where this operation runs.

Table 17-33 (Cont.) Options for the PostMovement Command

Options	Description
-props	<p>The absolute path to the properties file name. This file provides all the required user names, passwords and has the following structure:</p> <pre>#Properties file for T2P post movement setup parameters #Script location, needs to be the absolute path to scripts folder, inclusive. scripts-folder=/scratch/aime/work/t2p/scripts #Node manager information NMUser=bootstrap_admin NMPassword=welcome1 NMPort=5556 #Process id of node manager process, required on unix/linux NMPid=<pid> #Service name of nodemanager, required on windows NMSvcName=<Node Manager Service> #With the following true, reconcile GUID operation is run on this host. For other hosts set this to false. run-reconcile-guid=true #Path to domain from which to use the jps-config-jse.xml reconcile-guid-domain-path=/scratch/aime/work/APPTOP/insta nce/domains/adcdat10.us.oracle.com/CommonDomain #Database info to run reconcile GUID dburl=jdbc:oracle:thin:@adcdaf12.us.oracle.com:5593:DB5593 dbuser=fusion dbpassword=<dbpassword> #If following 2 props are set, then it is assumed use name and passwords are the same for all domains. AdminUser=FAAdmin AdminPassword=fusionfa1 #Domain admin user credentials, if the above is not set CommonDomainAdminUser=FAAdmin CommonDomainAdminPassword=fusionfa1 #Put all other domain user name and password below, if needed.</pre>
-workdir	<p>The directory used for log files and any other temporary files generated by the automation framework. The log files are located at <i>working_dir</i>\logs (UNIX) or <i>working_dir</i>\logs folder. This command is not needed if the T2P_DIR environment variable is set.</p>

This task also deletes unnecessary files from the following directories:

- (UNIX) *DOMAIN_HOME*/compositeConfigPlans
- (Windows) *DOMAIN_HOME*\compositeConfigPlans

17.9.2 Tips for Completing the Oracle RAC Configuration in the Target Environment

After completing the movement process, you must complete the Oracle RAC configuration in the target environment for any of the following scenarios that apply:

- For non-RAC to RAC movement: After moving Oracle Fusion Applications components from a source environment with no Oracle RAC to a target environment configured with Oracle RAC, you must configure the target environment with Oracle RAC. For information, see [Section 16.4.2](#).
- For RAC n instances to non-RAC movement: After completing the movement process, remove the Oracle RAC configuration from the target environment.
- For RAC m instances to RAC n instances movement, where the number of RAC instances in the source environment is less than or equal to the number of RAC instances in the target environment ($m \leq n$): After completing the movement process, add any additional $n-m$ generic data sources to the target environment. For information, see [Section 16.5.3.2](#).

For example, if the source environment contains three Oracle RAC instances, and the target environment contains four Oracle RAC instances, then you will have three generic data sources that are named `mds-soa-rac1`, `mds-soa-rac2`, and `mds-soa-rac3`. You must add one additional generic data source to the target environment.

- For RAC m instances to RAC n instances movement, where the number of Oracle RAC instances in the source environment is greater than the number of Oracle RAC instances in the target environment ($m > n$): After completing the movement process, remove the extra $m - n$ generic data sources from the last Oracle RAC instance in the target environment.

For example, if the source environment contains four Oracle RAC instances, and the target environment contains three Oracle RAC instances, then you will have four generic data sources that are named `mds-soa-rac1`, `mds-soa-rac2`, `mds-soa-rac3`, and `mds-soa-rac4`. The extra generic data source (`mds-soa-rac4`) points to the third Oracle RAC instance in the target environment (the third Oracle RAC instance will contain both `mds-soa-rac3` and `mds-soa-rac4`). You must remove `mds-soa-rac4` from the last Oracle RAC instance in the target environment.

17.10 A Case Study: Moving Oracle Fusion Supply Chain Management Components to a Target Distributed Environment

This section provides a case study of moving Oracle Fusion Supply Chain Management components from a source environment to a target distributed environment, whereby the Administration Servers of the Oracle WebLogic Server domains and all the system component instances run on the multiple hosts.

In this case study, the Oracle Fusion Supply Chain Management topology, shown in [Table 17-34](#), represents a sample topology in the source environment. It does not include a comprehensive list of applications.

Table 17–34 Sample Oracle Fusion Supply Chain Management Topology

System Component/Tier/Domain	Host in the Domain	Applications, Managed Servers, and System Components Running on the Host
Web Tier	WebHost-1	OHS
Oracle Identity Management Tier	IdmHost-1	Oracle Identity Management components: <ul style="list-style-type: none"> ■ Oracle Internet Directory ■ Oracle Virtual Directory ■ Oracle Access Manager ■ Oracle Identity Manager ■ Oracle Identity Federation
Mid Tier (BIDomain)	MidHost-1	Oracle BI EE applications: <ul style="list-style-type: none"> ■ coreapplication_obips1 (bi_server1/bi_server2) ■ coreapplication_obis1 (bi_server1/bi_server2) ■ coreapplication_obisch1 (bi_server1/bi_server2) ■ coreapplication_obijh1 (bi_server1/bi_server2) ■ coreapplication_obicc1 (bi_server1/bi_server2) Oracle BI EE components: <ul style="list-style-type: none"> ■ Oracle BI Server ■ Oracle BI Presentation Services ■ Oracle BI Cluster Controller ■ Oracle BI Scheduler ■ JavaHost ■ Oracle Essbase server
Mid Tier (CommonDomain)	MidHost-2	ApplicationsCoreSetupApp (HelpPortalServer_1/HelpPortalServer_2) GrcSetupApp (HelpPortalServer_1/HelpPortalServer_2) HelpPortalApp (HelpPortalServer_1/HelpPortalServer_2) HomePageApp (HomePageServer_1/HomePageServer_2) SetupApp (SalesServer_1/SalesServer_2) SetupEssApp (ess_server1/ess_server2) TopologyManagerServices (FunctionalSetupServer_1/FunctionalSetupServer_1)
Mid Tier (CRMDomain)	MidHost-3	ContractManagementApp (ContractManagmentServer_1/ContractManagmentServer_2) ContractManagementEssApp (ess_server1/ess_server2) CrmCommonApp (CRMCommonServer_1/CRMCommonServer_2) CrmEssApp (ess_server1/ess_server2)
Mid Tier (FinancialDomain)	MidHost-4	FinancialCommonApp (FinancialCommonServer_1/FinancialCommonServer_2) FscmAnalyticsApps (FinancialAnalyticsServer_1/FinancialAnalyticsServer_2) LedgerApp (LedgerServer_1/LedgerServer_2) PayablesApp (PayablesServer_1/PayablesServer_2) ReceivablesApp (ReceivablesServer_1/ReceivablesServer_2)
Mid Tier (HCMDomain)	MidHost-5	HcmAnalyticsApp (HCMAnalyticsServer_1/HCMAnalyticsServer_2) HcmCoreApp (CoreProcessesServer_1/CoreProcessesServer_1) HcmCoreExternalApp (CoreProcesses_Server_1/CoreProcesses_Server_2) HcmCoreSetupApp (CoreSetupServer_1/CoreSetupServer_2) HcmEssApp (ess_server1/ess_server2)
Mid Tier (ProcurementDomain)	MidHost-6	ProcurementApp (ProcurmentServer_1/ProcurementServer_2) ProcurementEssApp (ess_server1/ess_server2) SupplierPortalApp (SupplierPortal_Server_1/SupplierPortal_Server_2)

Table 17–34 (Cont.) Sample Oracle Fusion Supply Chain Management Topology

System Component/Tier/Domain	Host in the Domain	Applications, Managed Servers, and System Components Running on the Host
Mid Tier (ProjectsDomain)	MidHost-7	ProjectsFinancialsApps (ProjectsFinancialsServer_1/ProjectsFinancialsServer_2) ProjectsFinancialsEssApp (ess_server1/ess_server2)
Mid Tier (SCMDomain)	MidHost-8	AdvancedPlanningApp (AdvancedPlanningServer_1/AdvancedPlanningServer_2) CostManagementApp (CostManagementServer_1/CostManagementServer_2) GlobalizationApp (GlobalizationServer_1/GlobalizationServer_2) LogisticsApp (LogisticsServer_1/LogisticsServer_2) LogisticsCommonApp (SCMCommonServer_1/SCMCommonServer_2) OrderOrchestrationApp (OrderOrchestrationServer_1/OrderOrchestrationServer_2) ProductManagementCommonApp (ProductManagementServer_1/ProductManagementServer_2) ScmEssApp (ess_server1/ess_server2) ProductManagementCommonApp (ProductManagementServer_1/ProductManagementServer_2)
Oracle Database Tier	DbHost-1	

This case study describes the following steps for moving the Oracle Fusion Supply Chain Management components from a source environment to a target environment:

- Task 1, "Prepare the Source Environment"
- Task 2, "Prepare the Target Environment"
- Task 3, "Move Oracle Identity Management Domain Component Artifacts"
- Task 4, "Move the Binary Files"
- Task 6, "Create the Configuration Archives"
- Task 7, "Extract the Move Plan"
- Task 8, "Modify the Move Plan"
- Task 9, "Apply the Configurations to the Target Environment"
- Task 10, "Completing the Component-Specific Configuration Move"
- Task 11, "Move the Functional Setup Data"
- Task 12, "Complete the Post-Movement Tasks"

Task 1 Prepare the Source Environment

Before performing the full-movement tasks, make sure that the necessary preparations are in place in the source environment. For information, see [Section 17.3.1](#).

Task 2 Prepare the Target Environment

Before performing the full-movement tasks, complete the necessary preparations in the target environment. For information, see [Section 17.3.2](#).

Task 3 Move Oracle Identity Management Domain Component Artifacts

Move the Oracle Fusion Applications related artifacts for Oracle Identity Management for Oracle Identity Management from the source Oracle Identity Management

domains to the target Oracle Identity Management domains. For information, see [Section 17.4](#).

Task 4 Move the Binary Files

You must move the binary files of each Middleware home (`fusionapps` and `webtier_mwhome`) from the source environment to the target environment.

To move the binary files:

1. Using the `copyBinary` script, create a separate copy of each source Middleware home by copying the installed binaries and patches into an archive file. For more information, see [Section 17.5.1](#).
2. Using the `pasteBinary` script, apply the archive containing the binary files of the source Middleware home to the target environment. For more information, see [Section 17.5.2](#).
3. Move the Oracle Database client installation. For information, see [Section 17.5.3](#).

Task 5 Move the Webgate Installation

After moving the OHS instance home, you must move the Webgate installation. For information, see [Section 17.6.3](#).

Task 6 Create the Configuration Archives

Creating a configuration archive copies the source Oracle WebLogic Server domain home, OHS instance home, and Node Manager configuration files into an archive file. You must create a separate configuration archive for each Oracle WebLogic Server domain and each OHS instance in the source environment.

To create the configuration archives, run the `copyConfig` script for the OHS instances, node managers, and Oracle BI EE components, and each of the following Oracle WebLogic Server domains:

- `CRMDomain`
- `CommonDomain`
- `BIDomain`
- `HCMDomain`
- `FinancialDomain`
- `ProcurementDomain`
- `ProjectsDomain`
- `SCMDomain`

For information about running `copyConfig` for Oracle WebLogic Server domains, see [Section 17.6.4.1](#).

For information about running `copyConfig` for OHS instances, see [Section 17.6.4.2](#).

You must also run `copyConfig` for Node Manager on each machine within the topology of the source environment. For information, see [Section 17.6.4.3](#).

Task 7 Extract the Move Plan

Extracting the move plan consolidates into a single XML file the configuration settings of the source environment from the configuration archives, as well as from other configuration files specific to components such as SOA composite plans, adapter plans, and so on.

To extract the move plan, run the `extractMovePlan` script on the list of configuration archive locations, separated by comma.

In this case study, a total of thirteen configuration archives were created; each of the Oracle WebLogic Server domains (six archives for six domains), each of Node Manager (six archived for six node managers) and OHS instances. You must extract the configuration information from all thirteen configuration archives into a single move plan:

In this case study, a total of eighteen configuration archives were created; each of the Oracle WebLogic Server domains (eight archives for eight domains), each of Node Manager (eight archived for eight node managers), Oracle BI EE, and OHS instances. You must extract the configuration information from all eighteen configuration archives into a single move plan:

```
extractMovePlan.sh -javaHome USER_HOME/jrockit_160_20_D1.0.1-1705
                  -archiveLoc /FSCM_T2P/FSCM_WebTier.jar,/FSCM_T2P/FSCM_
CRMDomain.jar,/FSCM_T2P/FSCM_CommonDomain.jar,/FSCM_T2P/FSCM_BIDomain.jar,/FSCM_
T2P/HCMDomain.jar,/FSCM_T2P/FSCM_FinancialDomain.jar,/FSCM_T2P/FSCM_
ProcurementDomain.jar,/FSCM_T2P/FSCM_ProjectsDomain.jar,/FSCM_T2P/FSCM_
SCMDomain.jar,/FSCM_T2P/CRMDomain_NodeManager.jar,/FSCM_T2P/CommonDomain_
NodeManager.jar,/FSCM_T2P/BIDomain_NodeManager.jar,/FSCM_T2P/HCMDomain_
NodeManager.jar,/FSCM_T2P/FinancialDomain_NodeManager.jar,/FSCM_
T2P/ProcurementDomain_NodeManager.jar,/FSCM_T2P/ProjectsDomain_
NodeManager.jar,/FSCM_T2P/SCMDomain_NodeManager.jar,/FSCM_T2P/FSCM_
BIComponents.jar
                  -planDirLoc /FSCM_T2P/move_plan
                  -logDirLoc /FSCM_T2P/logs
                  -optimizationHints fusionApps,sameSchemaNameSinglePassword
```

The eighteen configuration archive locations, separated by comma, are specified to create a single move plan called `moveplan.xml` or multiple archive move plans in `/FSCM_T2P/move_plan`.

For more information about extracting move plans, see [Section 17.6.5](#).

Task 8 Modify the Move Plan

You must modify the move plan to specify properties (such as data source definitions, host names, port numbers, and end point addresses) for the target environment before you apply the copied configurations from the source environment to the target environment. For more information, see [Section 17.6.6](#).

Task 9 Apply the Configurations to the Target Environment

Apply the copied configurations from the source environment to the target environment by running the `pasteConfig` script for each of the following Oracle WebLogic Server domains, OHS instances, and Oracle BI EE components separately, using the same move plan (`moveplan.xml`):

- CRMDomain
- CommonDomain
- BIDomain
- HCMDomain
- FinancialDomain
- ProcurementDomain
- ProjectsDomain

- SCMDomain

You must also run `pasteConfig` for Node Manager on each machine within the topology in the target environment. For more information, see [Section 17.6.7.3](#).

Task 10 Completing the Component-Specific Configuration Move

You must move the component-specific configurations of the Oracle Fusion Middleware components.

For more information, see [Section 17.6.8](#).

Task 11 Move the Functional Setup Data

Moving the functional setup data applies the setup configuration data of the source environment to the target environment. Move the functional setup data for all the offerings. For information, see [Section 17.8](#).

Task 12 Complete the Post-Movement Tasks

After you move the Oracle Fusion Applications components across environments, you must complete various post-movement tasks. For information, see [Section 17.9](#).

Part V

Appendixes

This part contains the following appendixes:

- [Appendix A, "High Availability for Oracle Fusion Middleware Extensions for Applications"](#)
- [Appendix B, "High Availability for Oracle Enterprise Scheduler"](#)
- [Appendix C, "Seeded Searchable Objects, Search Categories, Index Schedules, and Facets"](#)

High Availability for Oracle Fusion Middleware Extensions for Applications

This appendix describes high availability considerations for Oracle Fusion Middleware Extensions for Applications (Applications Core). For more information about Oracle Fusion Applications high availability, see [Chapter 16](#). For more information on Oracle Fusion Middleware Extensions for Applications, see "Oracle Fusion Middleware Extensions for Applications" in *Oracle Fusion Applications Concepts Guide*.

High availability refers to the ability of users to access a system without loss of service. Oracle Fusion Middleware has an extensive set of high availability features, which protect its components and applications from unplanned down time and minimize planned downtime. For more information about Oracle Fusion Middleware high availability, see *Oracle Fusion Middleware High Availability Guide*.

This appendix includes the topic, [Section A.1, "How Oracle Fusion Middleware Extensions Components Use Fusion Middleware Components for High Availability and Failover."](#)

A.1 How Oracle Fusion Middleware Extensions Components Use Fusion Middleware Components for High Availability and Failover

With the approach of only building with use of other existing or lower level components, Applications Core obtains complete high availability without the need to implement any high availability features itself. However, it is then important to know what the base features each component is built upon, to understand how it behaves for high availability.

A.1.1 MDS

Applications Core relies almost entirely on MDS for storage of metadata used by Customization, Flexfields, and Menus. It is very important that MDS is configured for high availability. MDS for these operations should be based on a database rather than file based MDS architecture.

The MDS database-based repository can be configured for high availability Oracle database access. With this configuration, failure detection, recovery, and retry by MDS, as well as by the WebLogic infrastructure, result in the application's read-only MDS operations being protected from Oracle RAC database planned and unplanned downtimes. For more information about configuring multi data sources for MDS repositories, see "Configuring Multi Data Sources for MDS Repositories" in the *Oracle Fusion Middleware High Availability Guide*.

A.1.2 ADF

Most runtime state information of an application is held in ADF. Applications Core does not create its own components, so it can achieve failover support by following the ADF rules for high availability development. When you are designing an application to run in a clustered environment, you must:

- Ensure that all managed beans with a life span longer than one request are serialized. When the Fusion web application runs in a clustered environment, a portion of the application's state is serialized and copied to another server or a data store at the end of each request so that the state is available to other servers in the cluster.
- Ensure that Oracle ADF is aware of changes to managed beans stored in ADF scopes (view scope and page flow scope) and enable the tracking of changes to ADF memory scopes. When a value within a managed bean in either view scope or page flow scope is modified, the application needs to notify Oracle ADF so that it can ensure the bean's new value is replicated.

For more information about configuring high availability for ADF, see the "Configuring High Availability for Oracle ADF and WebCenter Applications" in the *Oracle Fusion Middleware High Availability Guide*.

A.1.3 WebLogic Server Failover for Session Content

Applications Core Session Management does cache the `AppSession` using the `HttpSession`, but it is marked as non-serializable so that it does not get persisted in the event of failover/replication. The `AppSession` will be re-established from the session cookie from the client that is set when the session is first created.

High Availability for Oracle Enterprise Scheduler

This chapter describes how to configure and manage a highly available Oracle Enterprise Scheduler environment.

This appendix includes the following topics:

- [Introduction to High Availability for Oracle Enterprise Scheduler](#)
- [Oracle Enterprise Scheduler Concepts](#)
- [Configuring High Availability for Oracle Enterprise Scheduler](#)
- [Managing an Oracle Enterprise Scheduler Cluster](#)

B.1 Introduction to High Availability for Oracle Enterprise Scheduler

A highly available cluster of Oracle Enterprise Scheduler servers is recommended for optimal job performance. This is especially useful for running asynchronous jobs remotely, which may require returning a status message upon completion.

For example, suppose an asynchronous ADF Business Components job runs remotely. Oracle Enterprise Scheduler expects the job to send a status upon completion using a web services callback. If Oracle Enterprise Scheduler runs on only one node, if that node is down, the callback message does not arrive and the status of the job is unknown. The job would then require manual intervention to mark its status as complete.

A two node cluster, however, allows all callbacks to process and arrive at their destination even if one server is down. A clustered Oracle Enterprise Scheduler environment allows callbacks to be delivered as required, and jobs to complete with the correct status automatically assigned by the system.

The main steps required for configuring a highly available Oracle Enterprise Scheduler environment are as follows:

1. Use the Oracle Fusion Applications Install and Configuration Wizard to set up a domain and configure a cluster.
2. Add nodes to the cluster as required in order to enhance scalability, allowing more processing power for jobs.

When a cluster node is added, the new node's processor configuration might have to be adjusted to assign appropriate work assignments.

For more information, see the Oracle WebLogic Server documentation.

3. Configure the load balancer. For more information, see the Oracle HTTP Server documentation.

Note: For information about troubleshooting an Oracle Enterprise Scheduler cluster, see the following chapters:

- "Troubleshooting Oracle Enterprise Scheduler" chapter in the *Oracle Fusion Applications Administrator's Troubleshooting Guide*.
 - "Troubleshooting Oracle Enterprise Scheduler" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Enterprise Scheduling Service*.
-

B.2 Oracle Enterprise Scheduler Concepts

In order to configure an Oracle Enterprise Scheduler environment, it helps to understand concepts such as the architecture of Oracle Enterprise Scheduler, its components and life cycle.

For more information about Oracle Enterprise Scheduler architecture, components, life cycle and life cycle tools, see the section "Oracle Enterprise Scheduler Concepts" in the chapter "High Availability for Oracle Enterprise Scheduler" in *Oracle Fusion Middleware Administrator's Guide for Oracle Enterprise Scheduling Service*.

B.3 Configuring High Availability for Oracle Enterprise Scheduler

In order to enable a highly available environment, it is recommended to run Oracle Enterprise Scheduler in a cluster of at least two nodes.

This section includes the following topics:

- [Oracle Enterprise Scheduler Configuration and Deployment Artifacts](#)
- [Oracle Enterprise Scheduler Logging](#)
- [Backup and Recovery](#)

For more information about cluster architecture, failover requirements, scalability, and load balancing for Oracle Enterprise Scheduler, see the section "Configuring High Availability for Oracle Enterprise Scheduler" in the chapter "High Availability for Oracle Enterprise Scheduler" in *Oracle Fusion Middleware Administrator's Guide for Oracle Enterprise Scheduling Service*.

B.3.1 Oracle Enterprise Scheduler Configuration and Deployment Artifacts

Configuration files are as follows:

- `ess.xml`: This file is part of the Oracle Enterprise Scheduler EAR file deployed to the Oracle Enterprise Scheduler cluster.
- `connections.xml`: This file is part of the Oracle Enterprise Scheduler EAR file deployed to the Oracle Enterprise Scheduler cluster. This file is used only with Oracle Enterprise Scheduler deployments running in an Oracle Fusion Applications environment.
- **MDS repository**: The Oracle Metadata repository stores Oracle Enterprise Scheduler job metadata. Oracle Enterprise Scheduler supports both database and file-based MDS repositories.

Deployment artifacts are as follows:

- J2EE application for core runtime and hosting applications.
 - Job metadata within Oracle MDS for core runtime and jobs loaded at startup.
- The Oracle WebLogic Server deployment is non-staged.

B.3.2 Oracle Enterprise Scheduler Logging

Use standard Oracle WebLogic Server logging for an Oracle Enterprise Scheduler cluster. Use logs in Oracle WebCenter Content to examine Oracle Enterprise Scheduler behavior. Oracle Enterprise Scheduler logging is configured by default in Oracle WebLogic Server.

The default location for log files for Oracle Enterprise Scheduler spawned jobs on UNIX servers is `/tmp/ess/requestFileDirectory`. Oracle Enterprise Scheduler operational log files can be found under `<DOMAIN_HOME>/servers/<SERVER_HOME>/logs/<server name>-diagnostic.log`, and `<MW_HOME>\user_projects\domains\<DOMAIN_HOME>\servers\<SERVER_HOME>\logs\<server name>-diagnostic.log` on Windows.

B.3.3 Backup and Recovery

Following are the backup and recovery guidelines for various components:

- Components stored on the file system: Product binaries, deployed application EAR files and standard Oracle WebLogic Server files in the domain root.
- Changes to the file system: The file system artifacts change when new EAR files are deployed or when the product is patched.
- Data stored in the database: The database stores all metadata and runtime data.
- Changes to database artifacts: Metadata changes when metadata is created and deployed from Oracle JDeveloper or Fusion Applications Control. Run time data changes when jobs are submitted, undergo state changes, and so on.

There is no consistency requirement between the artifacts stored on the file system and those in database. The file system stores EAR files and temporarily stores scheduled job output and log files. Job output and log files are saved to Oracle WebCenter Content upon job completion. Only one database is used, and two phase commit is not used in an Oracle Fusion Applications environment.

B.4 Managing an Oracle Enterprise Scheduler Cluster

Managing an Oracle Enterprise Scheduler cluster involves starting the cluster, propagating configuration changes throughout the cluster, deploying applications and handling unexpected behavior.

For more information about managing an Oracle Enterprise Scheduler cluster, see the section "Managing an Oracle Enterprise Scheduler Cluster" in the chapter "High Availability for Oracle Enterprise Scheduler" in *Oracle Fusion Middleware Administrator's Guide for Oracle Enterprise Scheduling Service*.

Note: Regarding failures in Oracle Java Transaction API migration and Oracle Java Message Service, Oracle JTA migration is unnecessary. When a node fails, there is no need to failover the node on another machine. Oracle Fusion Applications uses Oracle Enterprise Scheduler in a way does not require JTA recovery.

Oracle Enterprise Scheduler does not use Oracle Java Message Service (JMS) such that JMS recovery is not needed.

Seeded Searchable Objects, Search Categories, Index Schedules, and Facets

This appendix provides information about the searchable objects, search categories, index schedules, and facets that are available in Oracle Fusion Applications. These searchable objects, search categories, index schedules, and facets are available for Oracle Fusion Applications Search.

This appendix includes the following topics:

- [Searchable Objects by Product Family](#)
- [Search Categories by Product Family](#)
- [Index Schedules by Product Family](#)
- [Facets by Product Family](#)

The Searchable View Objects (SVOs) are stored in the JAR files in the following fusionapps Middleware subdirectories:

```
(UNIX) FA_MW_HOME/applications/product_family/deploy/Earproduct_
familySearch.ear/APP-INF/lib
(Windows) FA_MW_HOME\applications\product_family\deploy\Earproduct_
familySearch.ear\APP-INF\lib
```

C.1 Searchable Objects by Product Family

Searchable objects are sets of data that make view objects available for text search. They are used in an abstract way for exposing business data to search engines.

[Table C-1](#) provides a list of searchable objects available for Oracle Fusion Applications Search. The list is organized by product family and indicates the scope (global, focused, or both) of the search.

Note: While Oracle Fusion Financials searchable objects are seeded, they are not available for search. Do not use Oracle Fusion Financials searchable objects to enable search.

Table C-1 Searchable Objects by Product Family

Product Family	Searchable Object Name	Scope	Notes
Oracle Fusion Applications Help	TopicSearchPVO	Global	Do not modify this searchable object.
Oracle Fusion Customer Relationship Management	Billing Accounts - oracle.apps.customerCenter.infrastructure360.publicViewEcsf.CustomerAccountPVO	Both	
	Campaigns - oracle.apps.marketing.coreMarketing.campaigns.publicViewEcsf.CampaignsPVO	Global	
	Competitors - oracle.apps.sales.baseSales.competitors.publicViewEcsf.CompetitorPVO	Global	
	Contacts - oracle.apps.customerCenter.infrastructure360.publicViewEcsf.CustomerContactPVO	Both	
	Contract Documents - oracle.apps.contracts.dashBoard.textSearch.publicModel.view.ContractDocumentSearchVO	Focused	
	Customer Assessments - oracle.apps.customerCenter.infrastructure360.publicViewEcsf.CustomerAssessmentPVO	Both	
	Customer Contracts - oracle.apps.contracts.dashBoard.textSearch.publicModel.view.Customer360ContractPVO	Focused	
	Customer Interactions - oracle.apps.customerCenter.infrastructure360.publicViewEcsf.CustomerInteractionPVO	Both	
	Customer Leads - oracle.apps.marketing.leadMgmt.leads.publicViewEcsf.MklCustomer360LeadPVO	Focused	
	Customer Notes - oracle.apps.customerCenter.infrastructure360.publicViewEcsf.CustomerNotePVO	Both	
	Customer Opportunities - oracle.apps.sales.opptyMgmt.opportunities.publicViewEcsf.Customer360OpportunityPVO	Focused	
	Customer Profiles - oracle.apps.customerCenter.infrastructure360.publicViewEcsf.CustomerPVO	Both	
	Customer References - oracle.apps.sales.baseSales.references.publicViewEcsf.Customer360ReferencePVO	Focused	
Enterprise Contracts - oracle.apps.contracts.dashBoard.textSearch.publicModel.view.ContractHeaderSearchVO	Both		

Table C-1 (Cont.) Searchable Objects by Product Family

Product Family	Searchable Object Name	Scope	Notes
	Leads - oracle.apps.marketing.leadMgmt.leads.publicViewEcsf.MklLeadPVO	Both	
	Opportunities - oracle.apps.sales.opptyMgmt.opportunities.publicViewEcsf.OpportunityPVO	Both	
	Partner Profiles - oracle.apps.partnerFinder.publicViewEcsf.view.PartnerProfilePVO	Focused	
	Partners - oracle.apps.partnerMgmt.partnerCenter.publicViewEcsf.view.PartnerPVO	Global	
	Purchasing Contract Documents - oracle.apps.contracts.dashBoard.textSearch.publicModel.view.PoDocumentSearchVO	Focused	
	Purchasing Contracts - oracle.apps.contracts.dashBoard.textSearch.publicModel.view.PoHeaderSearchVO	Focused	
	References - oracle.apps.sales.baseSales.references.publicViewEcsf.ReferencePVO	Global	
	Sourcing Contract Documents - oracle.apps.contracts.dashBoard.textSearch.publicModel.view.AuctionDocumentSearchVO	Focused	
	Sourcing Contracts - oracle.apps.contracts.dashBoard.textSearch.publicModel.view.AuctionHeaderSearchVO	Focused	
Oracle Fusion Financials	Payables Credit Memo	Global	This searchable object is not available for search.
	Payables Payment	Global	This searchable object is not available for search.
	Payables Prepayment	Global	This searchable object is not available for search.
	Payables Refund	Global	This searchable object is not available for search.
	Payables Standard Invoice	Global	This searchable object is not available for search.
	Receivables Adjustment	Global	This searchable object is not available for search.
	Receivables ChargeBack	Global	This searchable object is not available for search.
	Receivables Credit Memo	Global	This searchable object is not available for search.
	Receivables Debit Memo	Global	This searchable object is not available for search.
	Receivables Invoice	Global	This searchable object is not available for search.

Table C-1 (Cont.) Searchable Objects by Product Family

Product Family	Searchable Object Name	Scope	Notes
	Receivables Miscellaneous Receipt	Global	This searchable object is not available for search.
	Receivables Standard Receipt	Global	This searchable object is not available for search.
	Collections Delinquency	Global	This searchable object is not available for search.
	Expense	Global	This searchable object is not available for search.
	Expense Report	Global	This searchable object is not available for search.
	Fixed Asset	Global	This searchable object is not available for search.
	Intercompany Transaction	Global	This searchable object is not available for search.
	Journal	Global	This searchable object is not available for search.
	Miscellaneous Receipt	Global	This searchable object is not available for search.
Oracle Fusion Human Capital Management	Worker	Global	Like Aria.
	Worker	Global	Secure.
Oracle Fusion Procurement	Purchase Requisitions - oracle.apps.prc.por.publicViewEcsf.RequisitionHeaderPVO	Global	
	Purchase Orders - oracle.apps.prc.po.publicViewEcsf.PurchaseOrderHeaderPVO	Global	
	Purchase Agreements - oracle.apps.prc.po.publicViewEcsf.PurchaseAgreementHeaderPVO	Global	
	Negotiations - oracle.apps.prc.pon.publicViewEcsf.NegotiationHeaderPVO	Global	
	Suppliers - oracle.apps.prc.poz.publicViewEcsf.SupplierPVO	Global	
Oracle Fusion Project	ProjectSearchPVO	Global	
	ProjectBudgetSearchPVO	Global	
	ProjectAssetPVO	Global	
	ProjectContractInvoicePVO	Global	
	ProjectEventPVO	Global	
	ProjectExpenditureItemPVO	Global	
	ProjectForecastSearchPVO	Global	
	ProjectTaskPVO	Global	

Table C-1 (Cont.) Searchable Objects by Product Family

Product Family	Searchable Object Name	Scope	Notes
Oracle Fusion Supply Chain Management	Orchestration Orders - oracle.apps.scm.doo.workbench.publicViewEcsf.view.OrchestrationOrderPVO	Global	
	Items - oracle.apps.scm.productModel.publicViewEcsf.ItemPVO	Global	
	Cost Adjustments - oracle.apps.scm.costing.distProcessors.reviewDistributions.publicViewEcsf.CostAdjustmentsPVO	Global	
	Material Transactions - oracle.apps.scm.costing.distProcessors.reviewDistributions.publicViewEcsf.MaterialTransactionsPVO	Global	
	Item Costs - oracle.apps.scm.costing.itemCosts.publicViewEcsf.ItemCostsPVO	Global	
	COGS Transactions - oracle.apps.scm.costing.distProcessors.reviewDistributions.publicViewEcsf.CogsTransactionPVO	Global	
	Receiving Transactions - oracle.apps.scm.receiptAccounting.receiptAccounting.reviewAccounting.publicViewEcsf.RecptAcctgPVO	Global	
	Accrual Adjustments - oracle.apps.scm.receiptAccounting.accrualClearing.adjustAccruals.publicViewEcsf.AdjustAccrualsPVO	Global	
	Receipts - oracle.apps.scm.receiving.publicViewEcsf.ReceiptPVO	Global	
	Inbound Shipments - oracle.apps.scm.receiving.publicViewEcsf.InboundShipmentPVO	Global	
Outbound Shipments - oracle.apps.scm.shipping.shipConfirm.deliveries.publicViewEcsf.ShipmentPVO	Global		

C.2 Search Categories by Product Family

Search categories are logical collections of searchable objects that facilitate group search on related items.

[Table C-2](#) provides a list of search categories available for Oracle Fusion Applications Search. The list is organized by product family and includes the business objects for each search category.

Note: While Oracle Fusion Financials search categories are seeded, they are not available for search. Do not use Oracle Fusion Financials search categories to enable search.

Table C-2 Search Categories by Product Family

Product Family	Search Category Name	Business Object	Scope		
Oracle Fusion Applications Help	Help	TopicSVO	GLOBAL		
Oracle Fusion Customer Relationship Management	All Contracts	Enterprise Contracts	LOCAL		
		Purchasing Contracts			
		Sourcing Contracts			
		Campaigns	Campaigns	GLOBAL	
		Competitors	Competitors	GLOBAL	
		Contacts	Contacts	GLOBAL	
		Contract Documents	Sourcing Contract Documents	LOCAL	
		Purchasing Contract Documents			
		Contract Documents			
	Customer	Customer	Customer Interactions	LOCAL	
			Contacts		
			Billing Accounts		
			Customer Assessments		
			Customer Contracts		
			Customer Profiles		
			Customer References		
			Customer Opportunities		
			Customer Notes		
			Customer Leads		
			Customers	Customers	Customer Profiles
Contacts					
Customer Assessments					
Customer Notes					
Billing Accounts					
Customer Interactions					
ECM Contracts	Enterprise Contracts	GLOBAL			
Lead	Leads	LOCAL			
Leads	Leads	GLOBAL			
Opportunities	Opportunities	GLOBAL			
Opportunity	Opportunities	LOCAL			
Partners	Partners	GLOBAL			
Public Partners	Public Partners	LOCAL			
References	References	GLOBAL			

Table C-2 (Cont.) Search Categories by Product Family

Product Family	Search Category Name	Business Object	Scope
Oracle Fusion Financials	Payables	Payables Standard Invoice	GLOBAL
		Payables Payment	
		Payables Credit Memo	
		Payables Prepayment	
		Payables Refund	
	Expenses	Expense Report	GLOBAL
		Expense	GLOBAL
	Fixed Assets	Fixed Asset	GLOBAL
	General Accounting	Journal	GLOBAL
		InterCompany Transaction	GLOBAL
	Account Receivables	Receivables Invoice	GLOBAL
		Receivables Miscellaneous Receipt	GLOBAL
		Receivables Receipt	GLOBAL
		Receivables Credit Memo	GLOBAL
		Receivables Debit Memo	GLOBAL
		Receivables Adjustments	GLOBAL
		Receivables ChargeBack	GLOBAL
		Collections	Collections Delinquency
	Oracle Fusion Human Capital Management	Worker	Worker
Global Worker			GLOBAL
Oracle Fusion Partner Relationship Management	Partners	PartnerProfileSVO	
Oracle Fusion Procurement	Procurement	Purchase Requisitions	GLOBAL
		Purchase Agreements	GLOBAL
		Purchase Orders	GLOBAL
		Negotiations	GLOBAL
		Suppliers	GLOBAL
Oracle Fusion Project	Projects	Project	GLOBAL
		Project Task	GLOBAL
		Project Asset	GLOBAL
		Project Contract Invoice	GLOBAL
		Project Budget Version	GLOBAL
		Project Forecast Version	GLOBAL
		Project Event	GLOBAL
		Project Expenditure Item	GLOBAL
Oracle Fusion Supply Chain Management	Cost Accounting	Cost Adjustments	GLOBAL

Table C-2 (Cont.) Search Categories by Product Family

Product Family	Search Category Name	Business Object	Scope
		Material Transactions	GLOBAL
		COGS Transactions	GLOBAL
		Item Costs	GLOBAL
	Receipt Accounting	Receiving Transactions	GLOBAL
		Accrual Adjustments	GLOBAL
	Orchestration Order	Orchestration Orders	GLOBAL
	Receiving	Receipts	GLOBAL
		Inbound Shipments	GLOBAL
	Shipping	Outbound Shipments	GLOBAL
	Items	Items	GLOBAL

C.3 Index Schedules by Product Family

Index schedules define the frequency at which the Oracle Secure Enterprise Search search engine crawls the searchable objects in Oracle Fusion Applications and indexes them as documents.

The following tables provide a list of index schedules available for Oracle Fusion Applications Search. The tables are organized by product family.

- [Table C-3, "Index Schedules for Oracle Fusion Customer Relationship Management"](#)
- [Table C-4, "Index Schedules for Oracle Fusion Applications Help"](#)
- [Table C-5, "Index Schedules for Oracle Fusion Human Capital Management"](#)
- [Table C-6, "Index Schedules for Personal Information Manager \(PIM\)"](#)
- [Table C-7, "Index Schedules for Oracle Fusion Procurement"](#)
- [Table C-8, "Index Schedules for Oracle Fusion Project"](#)
- [Table C-9, "Index Schedule for Oracle Fusion Supply Chain Management"](#)

Note: While Oracle Fusion Financials index schedules are seeded, they are not available for search. Do not use Oracle Fusion Financials index schedules to enable search.

Table C-3 Index Schedules for Oracle Fusion Customer Relationship Management

Schedule Name	Searchable Object	Business Object	Repeat Interval
Campaigns	Campaigns	oracle.apps.marketing.coreMarketing.campaigns.publicViewEcsf.CampaignsPVO	Manual
Competitors	Competitors	oracle.apps.sales.baseSales.competitors.publicViewEcsf.CompetitorPVO	Manual
Contacts	Contacts	oracle.apps.customerCenter.infrastructure360.publicViewEcsf.CustomerContactPVO	Manual
Contract Documents	Contract Documents	oracle.apps.contracts.dashboard.textSearch.publicModel.view.ContractDocumentSearchVO	Manual
Customer Assessments	Customer Assessments	oracle.apps.customerCenter.infrastructure360.publicViewEcsf.CustomerAssessmentPVO	Manual
Billing Accounts	Billing Accounts	oracle.apps.customerCenter.infrastructure360.publicViewEcsf.CustomerAccountPVO	Manual
Customer Contracts	Customer Contracts	oracle.apps.contracts.dashboard.textSearch.publicModel.view.Customer360ContractPVO	Manual
Customer Interactions	Customer Interactions	oracle.apps.customerCenter.infrastructure360.publicViewEcsf.CustomerInteractionPVO	Manual
Customer Leads	Customer Leads	oracle.apps.marketing.leadMgmt.leads.publicViewEcsf.MklCustomer360LeadPVO	Manual
Customer Notes	Customer Notes	oracle.apps.customerCenter.infrastructure360.publicViewEcsf.CustomerNotePVO	Manual
Customer Opportunities	Customer Opportunities	oracle.apps.sales.opptyMgmt.opportunities.publicViewEcsf.Customer360OpportunityPVO	Manual
Customer References	Customer References	oracle.apps.sales.baseSales.references.publicViewEcsf.Customer360ReferencePVO	Manual
Customer Profiles	Customer Profiles	oracle.apps.customerCenter.infrastructure360.publicViewEcsf.CustomerPVO	Manual
Enterprise Contracts	Enterprise Contracts	oracle.apps.contracts.dashboard.textSearch.publicModel.view.ContractHeaderSearchVO	Manual

Table C-3 (Cont.) Index Schedules for Oracle Fusion Customer Relationship Management

Schedule Name	Searchable Object	Business Object	Repeat Interval
Leads	Leads	oracle.apps.marketing.leadMgmt.leads.publicViewEcsf.MkLeadPVO	Manual
Opportunities	Opportunities	oracle.apps.sales.opptyMgmt.opportunities.publicViewEcsf.OpportunityPVO	Manual
Partner Profiles	Partner Profiles	oracle.apps.partnerFinder.publicViewEcsf.view.PartnerProfilePVO	Manual
Partners	Partners	oracle.apps.partnerMgmt.partnerCenter.publicViewEcsf.view.PartnerPVO	Manual
Purchasing Contracts	Purchasing Contracts	oracle.apps.contracts.dashboard.textSearch.publicModel.view.PoHeaderSearchVO	Manual
References	References	oracle.apps.sales.baseSales.references.publicViewEcsf.ReferencePVO	Manual
Sourcing Contracts	Sourcing Contracts	oracle.apps.contracts.dashboard.textSearch.publicModel.view.AuctionHeaderSearchVO	Manual

Table C-4 Index Schedules for Oracle Fusion Applications Help

Schedule Name	Searchable Object	Business Object	Repeat Interval
Help Schedule	TopicSearch	oracle.apps.atk.helpPortal.publicView.TopicSearchVO	Manual

Table C-5 Index Schedules for Oracle Fusion Human Capital Management

Schedule Name	Searchable Object	Business Object	Repeat Interval
All Workers Schedule	GlobalWorker	GlobalWorkerPVO	Manual
My Reports Schedule	Worker	WorkderPVO	Manual

Table C-6 Index Schedules for Personal Information Manager (PIM)

Product	Schedule Name	Search Category	Business Object	Repeat Interval
Product Management	Items Schedule	Items	oracle.apps.scm.productModel.publicViewEcsf.ItemPVO	Manual

Table C-7 Index Schedules for Oracle Fusion Procurement

Schedule Name	Searchable Object	Business Object	Repeat Interval
Negotiations Schedule	Negotiations	NegotiationsPVO	Manual
Purchase Agreement Schedule	Purchase Agreements	PurchaseAgreementPVO	Manual
Purchase Order Schedule	Purchase Orders	PurchaseOrderPVO	Manual
Purchase Requisition Schedule	Purchase Requisitions	PurchaseRequisitionPVO	Manual
Suppliers Schedule	Suppliers	SuppliersPVO	Manual

Table C-8 Index Schedules for Oracle Fusion Project

Schedule Name	Searchable Object	Business Object	Repeat Interval	Repeat Interval Unit	Schedule Launch Interval	Schedule Launch Time
Projects	ProjectSearch	ProjectSearchPVO	1	Daily		12:00 AM
Project Assets	ProjectAsset	ProjectAssetPVO	1	Weekly	Saturday	12:00 PM
Project Budgets	ProjectBudgetSearch	ProjectBudgetSearchPVO	1	Daily		12:00 AM
Project Contract Invoices	ProjectContractInvoice	ProjectContractInvoicePVO	2	Weekly	Saturday	12:00 PM
Project Events	ProjectEvent	ProjectEventPVO	2	Weekly	Saturday	12:00 PM
Project Expenditure Items	ProjectExpenditureItem	ProjectExpenditureItemPVO	1	Daily		12:00 AM
Project Forecasts	ProjectForecastSearch	ProjectForecastSearchPVO	1	Daily		12:00 AM
Project Tasks	ProjectTask	ProjectTaskPVO	1	Daily		12:00 AM

Table C-9 Index Schedule for Oracle Fusion Supply Chain Management

Schedule Name	Product	Search Category	Business Object	Repeat Interval
Inbound Schedule	Receiving	Receiving	oracle.apps.scm.receiving.publicViewEcsf.InboundShipmentPVO	Manual
Receipt Schedule	Receiving	Receiving	oracle.apps.scm.receiving.publicViewEcsf.ReceiptPVO	Manual
Shipping Schedule	Shipping	Shipping	oracle.apps.scm.shipping.shipConfirm.deliveries.publicViewEcsf.ShipmentPVO	Manual
Orchestration Orders Schedule	DOO	Orchestration Orders	oracle.apps.scm.doo.workbench.publicViewEcsf.view.OrchestrationOrderPVO	Manual

Table C–9 (Cont.) Index Schedule for Oracle Fusion Supply Chain Management

Schedule Name	Product	Search Category	Business Object	Repeat Interval
Cost Adjustments Schedule	Cost Management	Cost Accounting	oracle.apps.scm.costing.distProcessors.reviewDistributions.publicViewEcsf.CostAdjustmentsPVO	Manual
Cogs Transactions Schedule	Cost Management	Cost Accounting	oracle.apps.scm.costing.distProcessors.reviewDistributions.publicViewEcsf.CogsTransactionPVO	Manual
Material Transactions Schedule	Cost Management	Cost Accounting	oracle.apps.scm.costing.distProcessors.reviewDistributions.publicViewEcsf.MaterialTransactionsPVO	Manual
Item Costs Schedule	Cost Management	Cost Accounting	oracle.apps.scm.costing.itemCosts.publicViewEcsf.ItemCostsPVO	Manual
Accrual Adjustments Schedule	Cost Management	Receipt Accounting	oracle.apps.scm.receiptAccounting.accrualClearing.adjustAccruals.publicViewEcsf.AdjustAccrualsPVO	Manual
Receiving Transactions Schedule	Cost Management	Receipt Accounting	oracle.apps.scm.receiptAccounting.reviewAccounting.publicViewEcsf.RecptAcctgPVO	Manual

C.4 Facets by Product Family

Facets are predefined aspects, or dimensions, of a searchable object. Facet-based navigation of the search results, or Faceted Navigation, allows end users to narrow down their search by navigating through structured search results.

[Table C–10](#) provides a list of facets available for Oracle Fusion Applications Search. The list is organized by product family and includes the facets for each search category.

Note: While Oracle Fusion Financials facets are predefined, they are not available for search. Do not use Oracle Fusion Financials facets to enable search.

Table C-10 Facets by Product Family

Product Family	Search Category	Facet	
Oracle Fusion Customer Relationship Management		oracle.apps.customerCenter.infrastructure360.publicViewEcsf.CustomerInteractionPVO	
		oracle.apps.customerCenter.infrastructure360.publicViewEcsf.CustomerNotePVO	
		oracle.apps.sales.opptyMgmt.opportunities.publicViewEcsf.OpportunityPVO	
Oracle Fusion Supply Chain Management	Shipping	oracle.apps.scm.shipping.shipConfirm.deliveries.publicViewEcsf.ShipmentPVO	
	Receiving	oracle.apps.scm.receiving.publicViewEcsf.ReceiptPVO	
		oracle.apps.scm.receiving.publicViewEcsf.InboundShipmentPVO	
	Projects		oracle.apps.projects.foundation.publicView.ecsf.ProjectAssetPVO
			oracle.apps.projects.foundation.publicView.ecsf.ProjectBudgetSearchPVO
			oracle.apps.projects.foundation.publicView.ecsf.ProjectContractInvoicePVO
			oracle.apps.projects.foundation.publicView.ecsf.ProjectEventPVO
			oracle.apps.projects.foundation.publicView.ecsf.ProjectExpenditureItemPVO
			oracle.apps.projects.foundation.publicView.ecsf.ProjectForecastSearchPVO

