Oracle® Fusion Applications

Administrator's Guide

11*g* Release 1 (11.1.1.5)

E14496-02

September 2011

ORACLE®

Oracle Fusion Applications Administrator's Guide, 11*g* Release 1 (11.1.1.5)

E14496-02

Contents

Part I   Understanding Oracle Fusion Applications

1   Introduction to Oracle Fusion Applications for Systems Administrators

Part II   Basic Administration

2   Getting Started with Administering Oracle Fusion Applications

5   Managing Oracle Enterprise Scheduler Service and Jobs

6   Managing Report Delivery Servers

7 Managing Search with Oracle Enterprise Crawl and Search Framework

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

Part III   Monitoring

9   Monitoring Oracle Fusion Applications

10   Monitoring the Oracle Fusion Applications Middle Tier

11   Diagnosing the Oracle Fusion Applications Middle Tier Performance

12   Monitoring and Tuning Oracle Database for Oracle Fusion Applications

13 Managing Oracle Fusion Applications Log Files and Diagnostic Tests

Part IV Advanced Administration

14 Backing Up and Recovering Oracle Fusion Applications

15   Configuring High Availability and Scaling Out Oracle Fusion Applications

16   Moving Components for Oracle Fusion Applications Across Environments

Part V   Troubleshooting

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

18 Troubleshooting Oracle Business Intelligence

19 Troubleshooting Oracle Enterprise Crawl and Search Framework

20 Troubleshooting Oracle Enterprise Scheduler

21 Troubleshooting Oracle SOA Suite

22   Troubleshooting Oracle Identity Management

23   Troubleshooting Oracle WebCenter Portlets

24   Troubleshooting Performance, Scalability, Reliability

Part VI   Appendixes

A   High Availability for Oracle Fusion Middleware Extensions for Applications

B   High Availability for Oracle Enterprise Scheduler

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

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 Middleware Administrator's 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.
italic	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.

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"

1

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
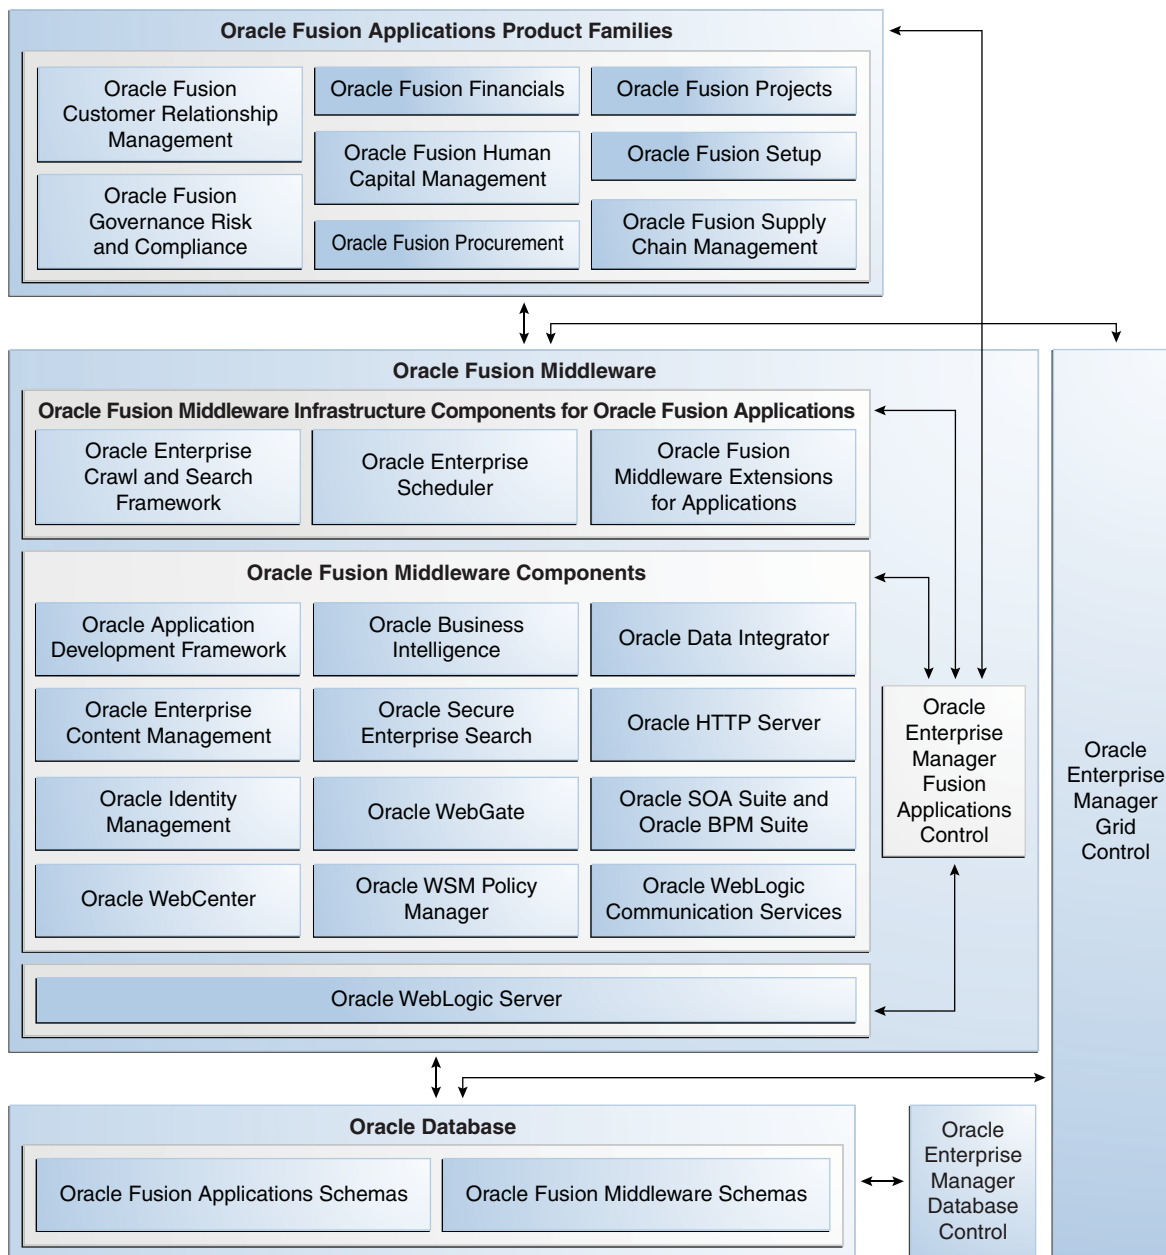
The section contains the following topics

- 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 Grid Control
- Oracle Enterprise Manager Database Control

1.1.1 Oracle Fusion Applications Product Families

During installation, you select a **product offering** (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
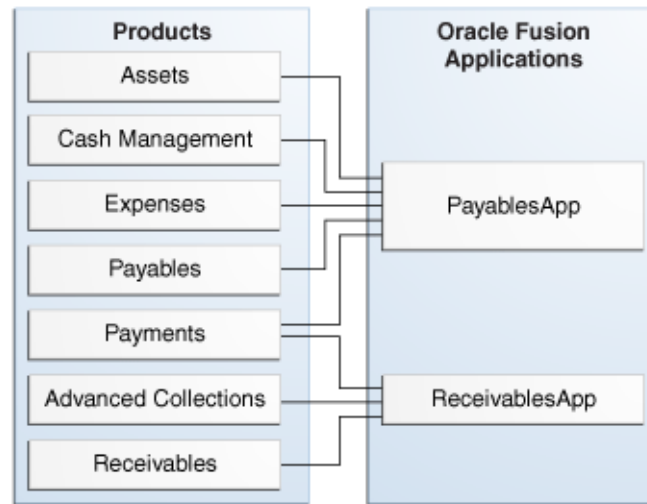
Oracle Fusion Applications include the following product families described in Table 1–1.

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	Supports the other product families. In addition to Oracle Fusion Functional Setup Manager for setting up functional data, this product family includes applications to assist application users: - 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 *Oracle Fusion Applications Installation Guide* and the *Oracle Fusion Applications Common Implementation Guide*.

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)	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.
	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.
	See the "Oracle Fusion Middleware Extensions for Applications" section in the *Oracle Fusion Applications Concepts Guide*.
Oracle Enterprise Scheduler	Enables you to manage and schedule jobs for Oracle Fusion Applications.
	See Chapter 5, "Managing Oracle Enterprise Scheduler Service and Jobs."
Oracle Enterprise Crawl and Search Framework (ECSF)	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.
	See Chapter 7, "Managing Search with Oracle Enterprise Crawl and Search Framework."

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)	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.
	See the *Oracle Fusion Middleware Getting Started with Oracle Data Integrator*.

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

Oracle Fusion Middleware Components	Description
Oracle Business Intelligence	Oracle Business Intelligence provides a complete, integrated solution of analytics and reporting for Oracle Fusion Applications.
	See the following documentation on Oracle Business Intelligence components in Oracle Fusion Applications:
	■ Oracle BI Enterprise Edition: *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*
	■ Oracle Business Intelligence Publisher: Chapter 6
	■ Oracle Essbase: *Oracle Essbase Database Administrator's Guide* and *Oracle Essbase Technical Reference*
	■ Oracle Enterprise Performance Management: http://www.oracle.com/technetwork/middleware/performance-management/documentation/index.html
	■ Oracle Real-Time Decisions: *Oracle Fusion Middleware Platform Developer's Guide for Oracle Real-Time Decisions*
	■ Oracle Transactional Business Intelligence:
	After the Oracle Fusion applications are running, you can configure the Oracle Business Intelligence Applications for analyzing historical data. See the *Oracle Fusion Middleware Configuration Guide for Oracle Business Intelligence Applications*.
Oracle Data Integrator	Extracts, transforms, and loads data for the product families.
	See the *Oracle Fusion Middleware Developer's Guide for Oracle Data Integrator*.
Oracle Enterprise Content Management	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.
	For more information about documentation for Oracle Enterprise Content Management components in Oracle Fusion Applications, see:
	■ Oracle Universal Content Management: *Oracle Fusion Middleware System Administrator's Guide for Oracle Content Server*
	■ Oracle Imaging and Process Management: *Oracle Fusion Middleware Administrator's Guide for Oracle Imaging and Process Management*
Oracle HTTP Server	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.
	See the *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server*.

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

Oracle Fusion Middleware Components	Description
Oracle Identity Management	Provides a shared infrastructure for all applications, enabling developers to incorporate identity management into applications.
	For more information about documentation for Oracle Identity Management components in Oracle Fusion Applications, see the following guides:
	■ Oracle Platform Security Services (OPSS: *Oracle Fusion Middleware Application Security Guide* and *Oracle Fusion Middleware Security Overview*
	■ Oracle Access Manager: *Oracle Fusion Middleware Administrator's Guide for Oracle Access Manager with Oracle Security Token Service*
	■ Oracle Authorization Policy Manager: *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)*
	■ Oracle Identity Manager: *Oracle Fusion Middleware Administrator's Guide for Oracle Identity Manager*
	■ Oracle Internet Directory: *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*
	■ Oracle Virtual Directory: *Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory*
Oracle SOA Suite	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.
	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.
	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.
	See the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite.*

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

Oracle Fusion Middleware Components	Description
Oracle Secure Enterprise Search (Oracle SES)	Provides a search engine for Oracle Fusion Applications Search.
	For more information on Oracle SES, see the following guides:
	■ *Oracle Secure Enterprise Search Administrator's Guide*
	■ *Oracle Secure Enterprise Search Administration API Guide*
	■ *Oracle Secure Enterprise Search Java API Reference*
	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."
Oracle WebCenter	Enables you to create social applications, enterprise portals, collaborative communities, and composite applications, built on a standards-based, service-oriented architecture. Oracle WebCenter 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 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.
	See the *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter*.
Oracle WebGate	Acts as a communicator plug-in that accepts users requests through Oracle HTTP Server and communicates with Oracle Access Manager.
	See the *Oracle Fusion Middleware Administrator's Guide for Oracle Access Manager with Oracle Security Token Service*.
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.
	For more information for enabling these services, see Chapter 8, "Managing Oracle WebLogic Communication Services for Click-to-Dial Functionality."
Oracle WebLogic Server	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).
	See the *Oracle Fusion Middleware Introduction to Oracle WebLogic Server*.

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 Grid Control.

1.1.6 Oracle Enterprise Manager Grid Control

Oracle Enterprise Manager Grid Control (Grid Control) enables you to monitor the Oracle database and Oracle Fusion Middleware components across the entire Oracle Fusion Applications environment. 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 Installation 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 Installation 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 (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 Oracle Fusion Applications Repository Creation Utility, 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
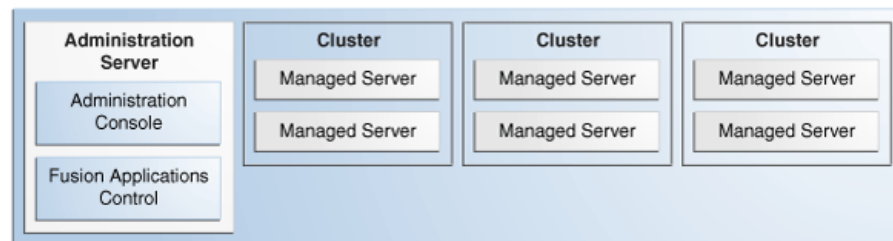
As an example of greater detail for a product family, Figure 1–4 shows a portion of the domain for the Oracle Fusion Financials product family with an Administration Server and eight Managed Servers in a cluster. The applications are each deployed to a Managed Server within a cluster. In addition to the applications for Oracle Fusion Financials, the domain also contains an Oracle Enterprise Scheduler Service 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

Cluster
(FinancialCommonCluster)
Managed Server
(FinancialCommonServer_1)
FinancialCommonApp
application

Cluster
(FinancialAnalyticsCluster)
Managed Server
(FinancialSearchServer_1)
FinancialAnalyticsServer_1
application

Cluster
(FinancialSearchCluster)
Managed Server
(FinancialSearchServer_1)
FinancialSearchServer_1
application

Administration Server
Administration Console
Fusion Applications Control

Cluster
(GeneralLedgerCluster)
Managed Server
(GeneralLedgerServer_1)
GeneralLedgerApp
application

Cluster
(PayablesCluster)
Managed Server
(PayablesServer_1)
PayablesApp
application

Cluster
(ReceivablesCluster)
Managed Server
(ReceivablesServer_1)
ReceivablesApp
application

Cluster
(FinancialsESSCluster)
Managed Server
(FinancialsESSServer_1)
ESSAPP
application

Cluster
(FinancialsSOACluster)
Managed Server
(FinancialsSOAServer_1)
soa_infra
application

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
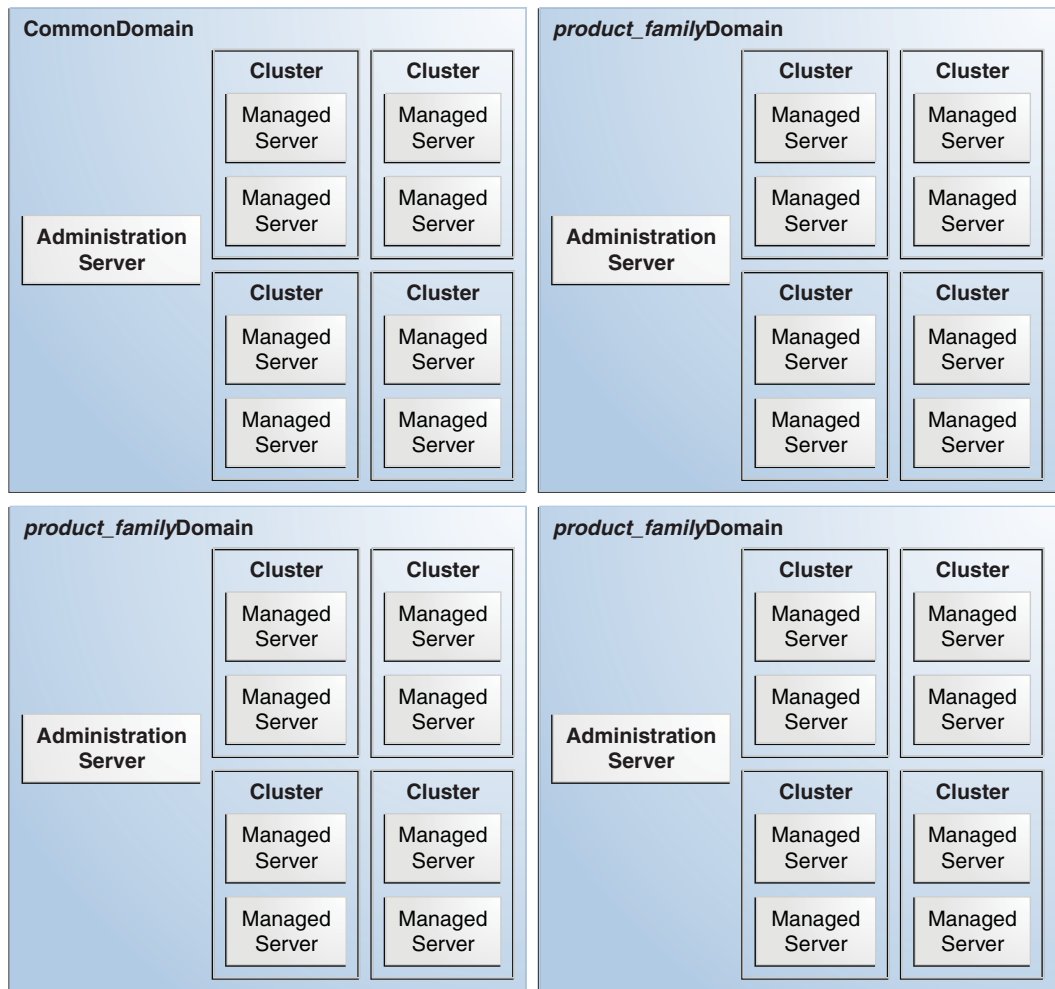
1.2.3 Provisioned Oracle Fusion Applications Home Directories

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 `/net/mount1/appbase` for the purpose of showing an example directory structure.	`APPLICATIONS_BASE`
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 `/net/mount2` for the purpose of showing an example directory structure. Provisioning automatically creates a subdirectory named `instance` as a container for the configuration files.	`APPLICATIONS_CONFIG`

Middleware home	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).
	
	An installation of Oracle Fusion Applications creates two Middleware homes:
	
	■ `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
	
	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 `FA_MW_HOME`, 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 `WT_MW_HOME`. In cases where differentiation is not necessary, this guide refers to the Middleware home and uses the `MW_HOME` variable.
		`MW_HOME`
		`FA_MW_HOME`
		`WT_MW_HOME`

Oracle Business Intelligence configuration home	This directory contains the configuration files for the Oracle Business Intelligence installation.	`BI_CONFIG_HOME`
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).	`ORACLE_COMMON_HOME`

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.	`DBCLIENT_ORACLE_HOME`
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.	`WL_HOME`

Oracle home	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 `soa`, and the application artifacts for Oracle Fusion Applications are stored in an Oracle home directory named `applications`.
	
	An Oracle home resides within the directory structure of the Middleware home. There can be multiple Oracle homes within each Middleware home.
	
	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 `soa` Oracle home directory and its variable is `SOA_ORACLE_HOME`, while the Oracle home for Oracle Fusion Applications is referred to as the Oracle Fusion Applications Oracle home or the `applications` Oracle home and its variable is `FA_ORACLE_HOME`.
		`ATGPF_ORACLE_HOME`
		`FA_ORACLE_HOME`
		`JROCKIT_ORACLE_HOME`
		`SOA_ORACLE_HOME`
		`WC_ORACLE_HOME`
		`WG_ORACLE_HOME`
		`WT_ORACLE_HOME`

Oracle Web Tier configuration home	This directory contains the configuration files for the Oracle Web Tier installation.	`WT_CONFIG_HOME`
Oracle WebLogic Server domain home	This directory contains the configuration files for the Administration Server and the Managed Servers for the product family.	`DOMAIN_HOME`

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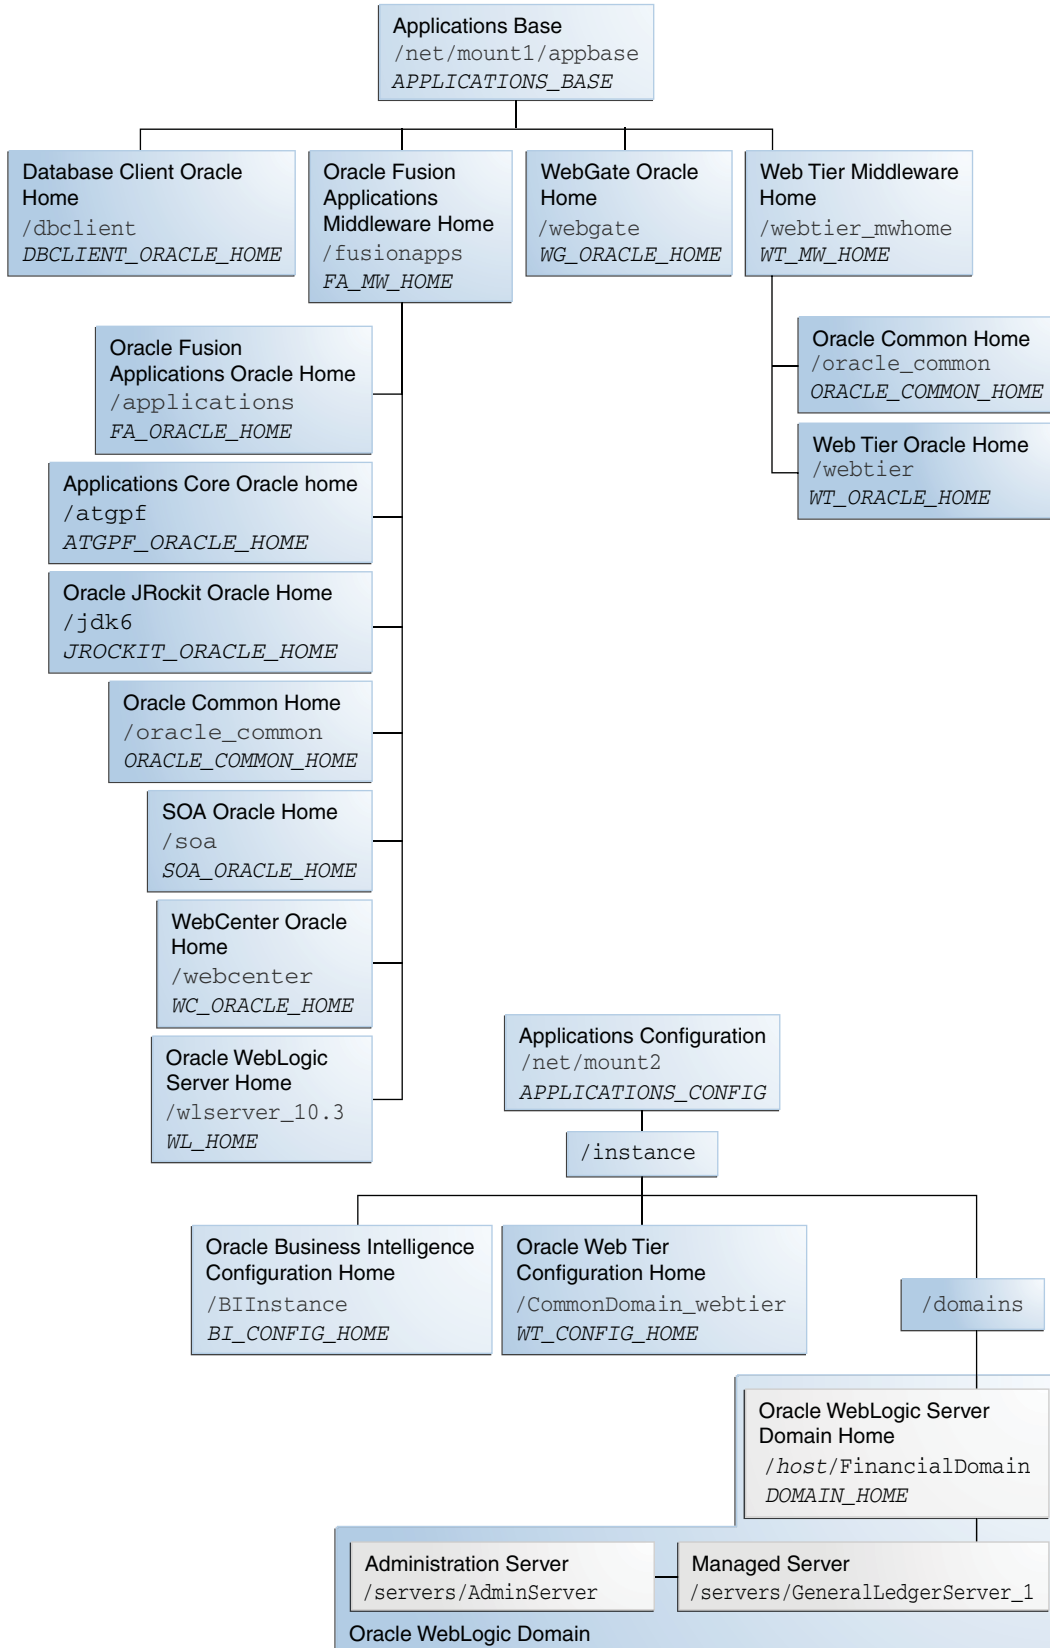
For more information about the detailed directory structure, see the *Oracle Fusion Applications Installation Guide*.

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 14.

10. Configure load balancing between different components or applications. See Chapter 15.

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 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

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"

2

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 Oracle Enterprise Manager Grid 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 Grid 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 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 *Oracle Fusion Middleware Administrator's Guide*.
Oracle Enterprise Manager Grid Control (Grid Control)	Enables you to monitor historical data for Oracle Fusion Middleware components across all Oracle WebLogic Servers for the Oracle Fusion Applications environment. See Section 2.2 to understand the key difference between Grid Control and Fusion Applications Control, and see Section 2.7 to get started with using Grid 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 *Oracle Fusion Middleware Administrator's Guide* 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	The Oracle Fusion Applications environment includes the following tools for managing application security:
	■ Oracle WebLogic Server Administration Console enables you to manage policies of enterprise applications. See the *Oracle Fusion Middleware Securing Resources Using Roles and Policies for Oracle WebLogic Server*.
	■ Oracle Authorization Policy Manager enables you to manage application policies. See the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)*.
	■ 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 *Oracle Fusion Middleware Administrator's Guide for Oracle Access Manager with Oracle Security Token Service*.
	■ 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 *Oracle Fusion Middleware Administrator's Guide for Oracle Directory Integration Platform*.
	■ Oracle Web Service Policy Manager enables you manage and secure web services for SOA composite applications. See the *Oracle Fusion Middleware Security and Administrator's Guide for Web Services*.
	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 *Oracle Fusion Applications Security Guide*.

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 *Oracle Database 2 Day DBA*.
Grid Control	Enables you to monitor historical data for the Oracle Database. See Section 2.7.

2.2 Key Differences Between Fusion Applications Control and Oracle Enterprise Manager Grid Control

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

Grid Control enables you to monitor the Oracle database and Oracle Fusion Middleware components across the entire Oracle Fusion Applications environment. To use Grid Control, you need to separately install Oracle Enterprise Manager 11*g* Grid Control with the Application Management Pack for Oracle Fusion Applications. For more information about this control, see Section 2.7. For more information about installing Grid Control, see the *Oracle Enterprise Manager Grid Control Basic Installation Guide*.

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:**
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

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

Task	Tool to Use
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:

`DOMAIN_HOME/servers/server_name/logs/access.log`
`DOMAIN_HOME/servers/server_name/data/ldap/log/EmbeddedLDAP.log`
`DOMAIN_HOME/servers/server_name/data/ldap/log/EmbeddedLDAPAccess.log` |
| 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.
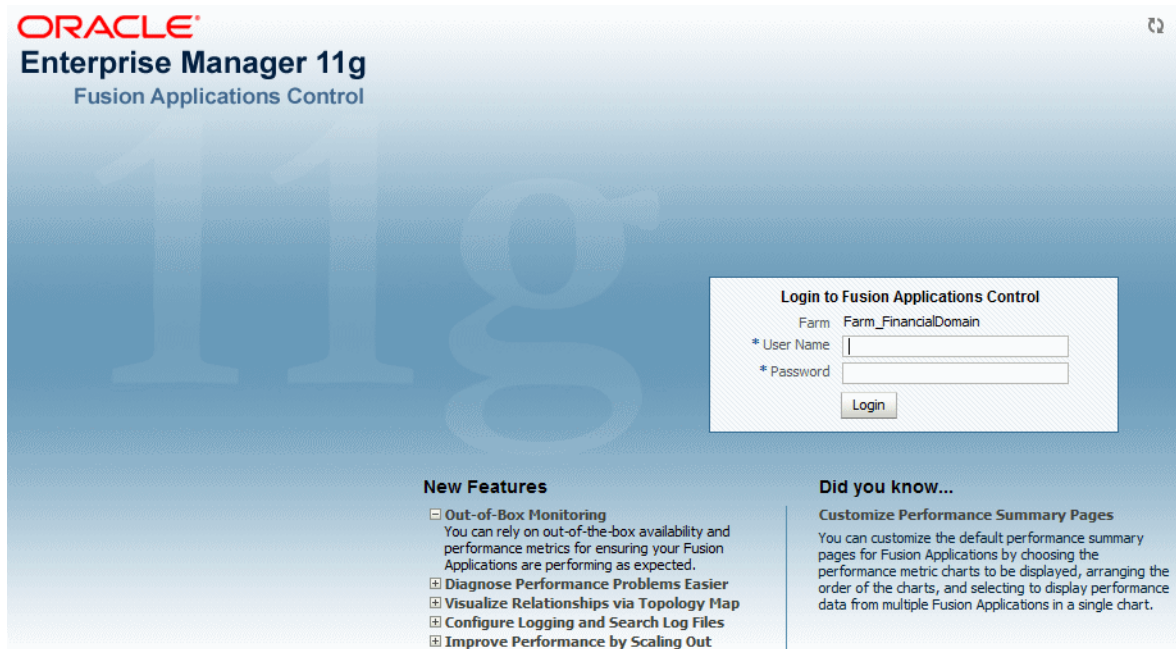
2. Enter Oracle Fusion Middleware administrative user name and password, and then click **Login.**

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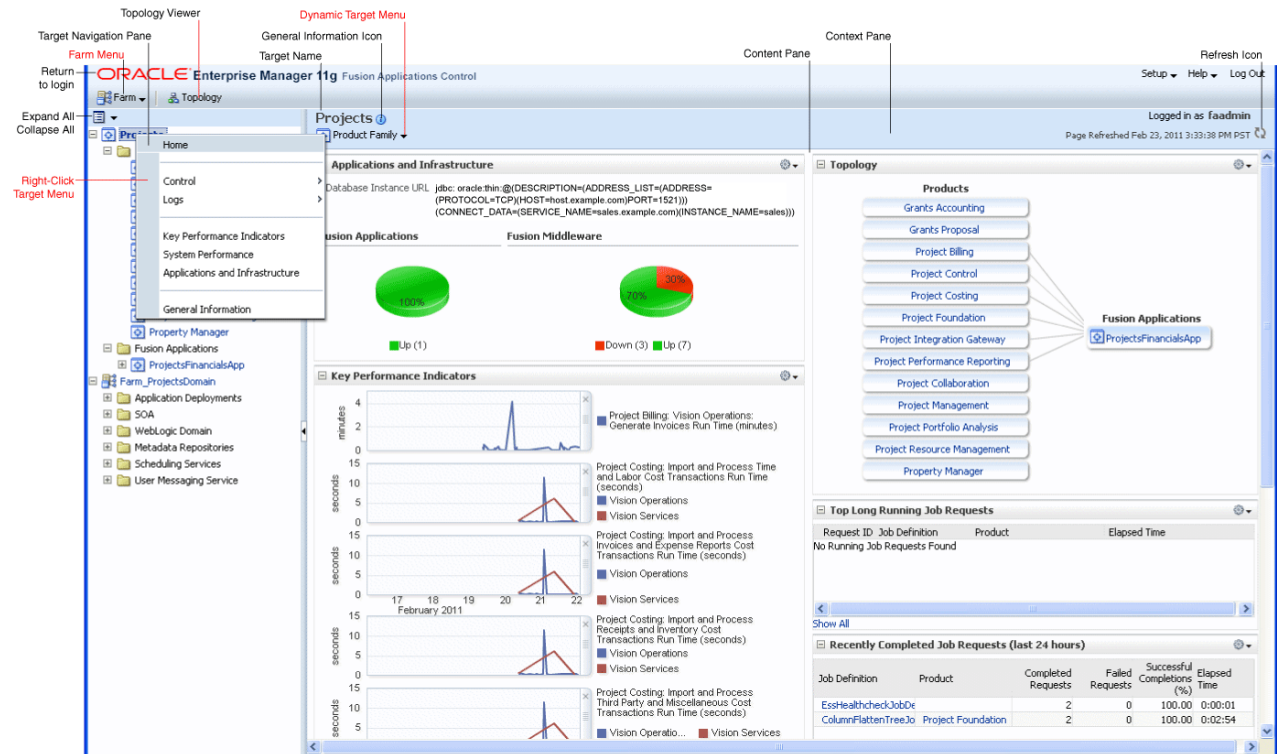
In the preceding figure, the following items are called out:

- **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:
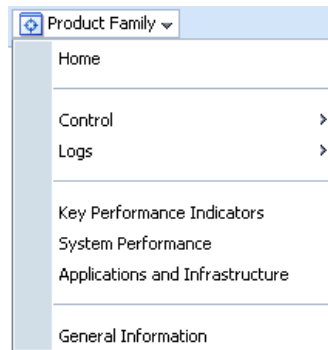
   The following figure shows the **Product** menu:
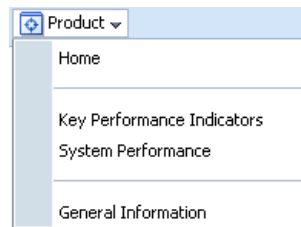
   The **Product Family** and **Product** menu displays the options described in Table 2–6.

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 9.6, and Section 9.7.
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.

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 9.5 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:
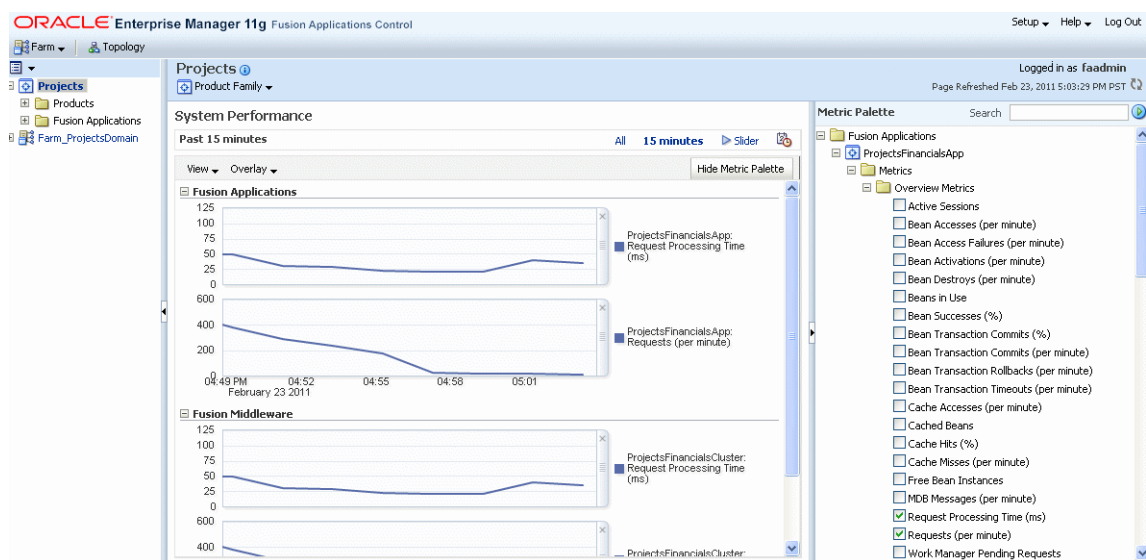
3. Select additional metrics to add them to the Performance Summary.

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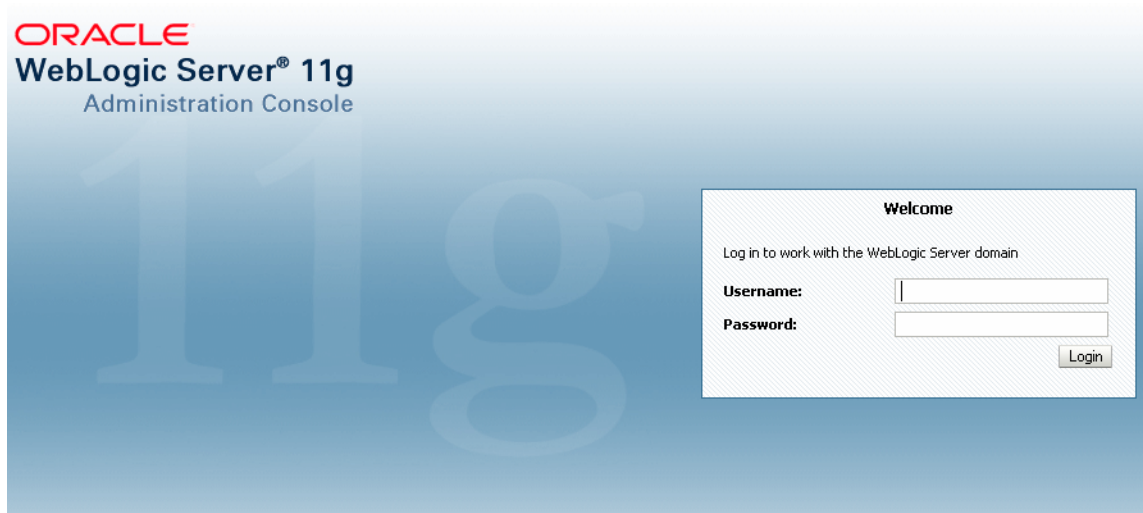
2. Enter Oracle Fusion Middleware administrative user name and password, and then click **Login.**

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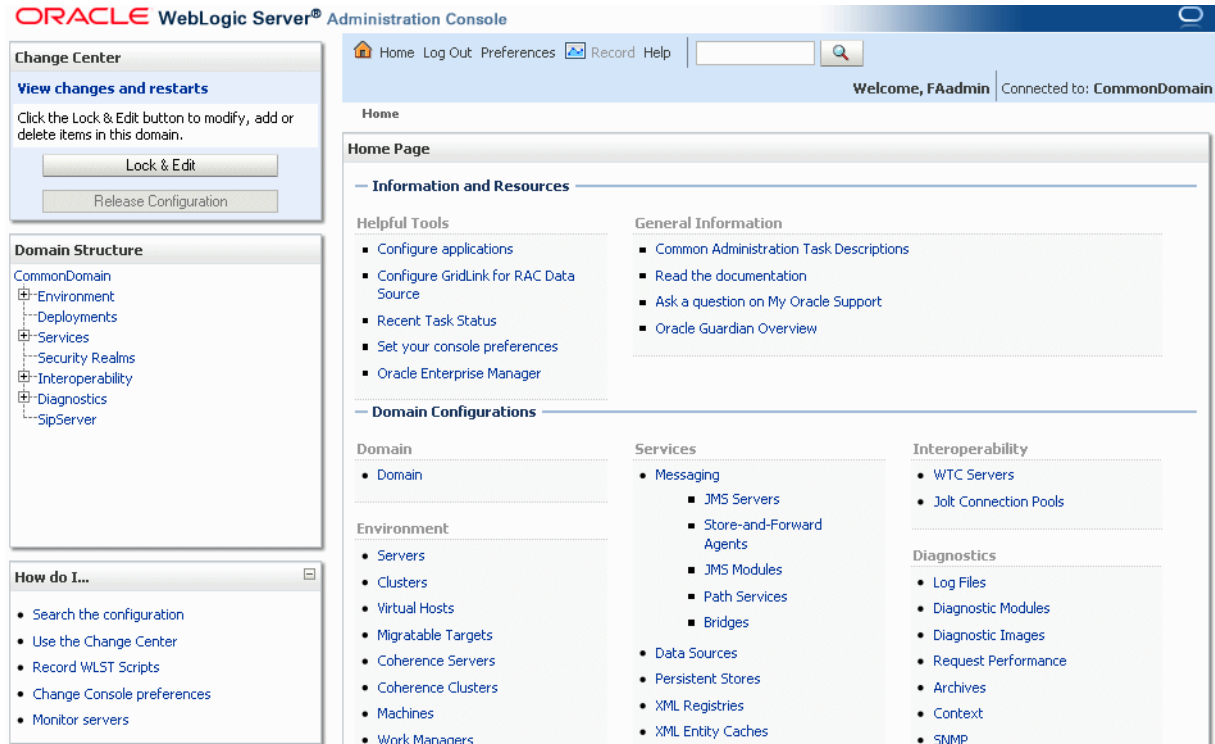
Alternatively, you can access the Administration Console from Fusion Applications Control, from the home pages of targets such as the Administration Server or Managed Servers.

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 Grid Control

Grid Control enables you to monitor runtime and historical data for the Oracle database and the Oracle Fusion Middleware components across the entire Oracle Fusion Applications environment.

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

The section contains the following topics:

- Using Grid Control
- Using Grid Control Help

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

- Section 10.5, "Performing Proactive Performance Monitoring with Grid Control"
- Chapter 11, "Diagnosing the Oracle Fusion Applications Middle Tier Performance"

2.7.1 Using Grid Control

To use Grid Control, you need to separately install Oracle Enterprise Manager 11*g* Grid Control with the Application Management Pack for Oracle Fusion Applications. For more information about installing Grid Control, see the *Oracle Enterprise Manager Grid Control Basic Installation Guide*.
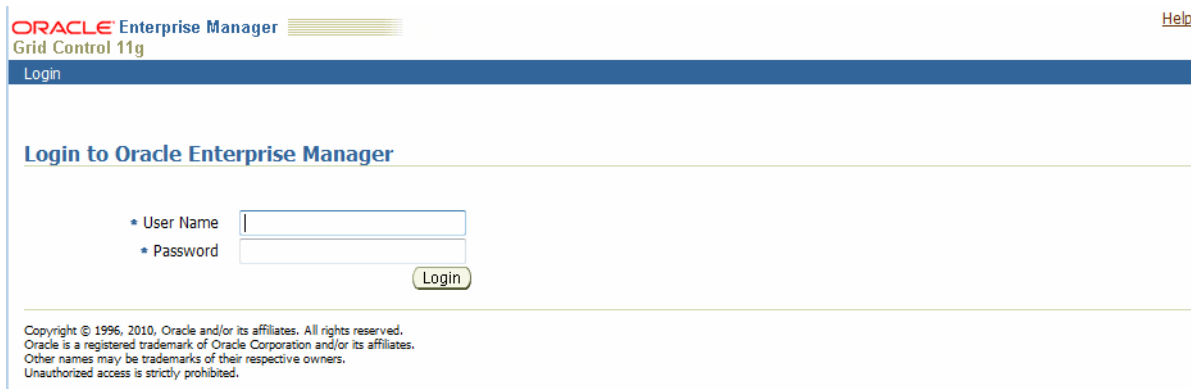
To display Grid Control:

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.

ORACLE Enterprise Manager
Grid Control 11g   Help

Login

Login to Oracle Enterprise Manager

* User Name [   ]
* Password   [   ]
                                    (Login)

Copyright © 1996, 2010, Oracle and/or its affiliates. All rights reserved.
Oracle is a registered trademark of Oracle Corporation and/or its affiliates.
Other names may be trademarks of their respective owners.
Unauthorized access is strictly prohibited.

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 2.7.2 and click the **Help** link at the top of the page.

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

For more information, see the "Enabling or Disabling Licenses" section in the *Oracle Enterprise Manager Grid Control Advanced Installation and Configuration Guide*.

4. Discover the middle tier components.

In order to manage and monitor Fusion Middleware components, you must discover the target Oracle WebLogic Server domains they reside on. For information on discovering middleware targets, see the "Discovering Middleware Targets" chapter in the *Oracle Enterprise Manager Getting Started with Fusion Middleware Management*.

2.7.2 Navigating the Main Tabs of Grid Control

Figure 2–3 shows the home page of a Grid Control. It displays 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.

Figure 2–3   Home Page of Grid Control
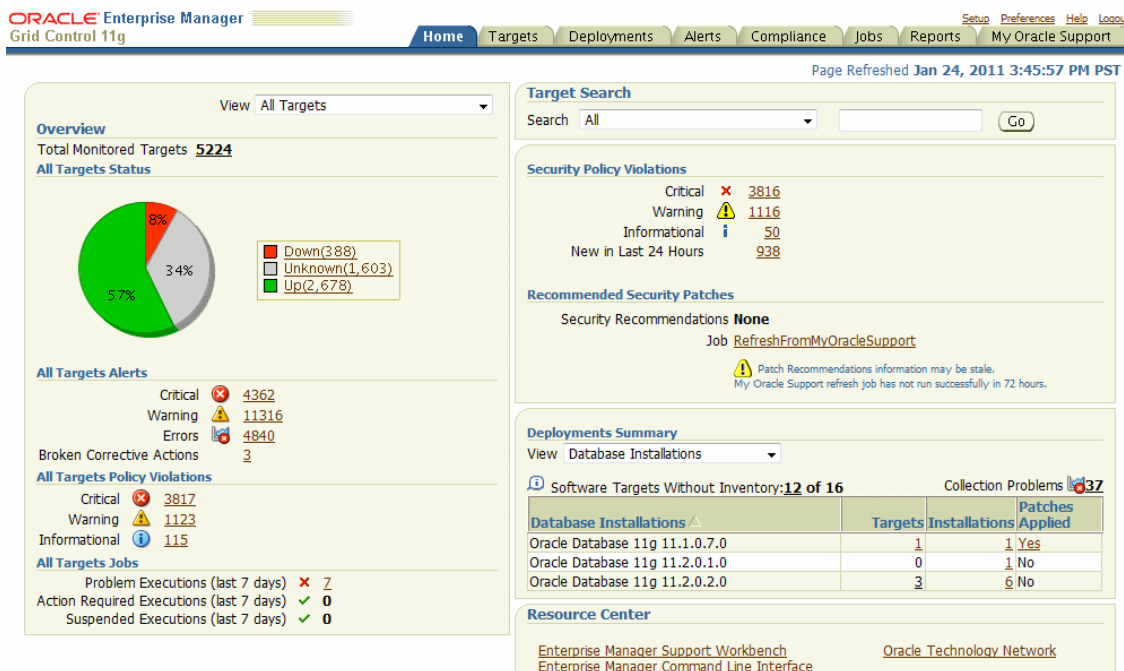
In addition to the Home tab, you can access the following tabs from the Home page:

- **Targets** tab enables you to select the type of target you want to monitor. For more information about the **Middleware** subtab, see Section 10.5.1.

- **Alerts** tab enables you to identify problems. the **Targets Down** subtab provides a good starting point for analyzing the cause of the problem.

- **Compliance** tab enables you to view the policies for specific targets and obtain a summary of all the policy violations in your enterprise.

- **Jobs** tab 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** tab provides enables you to manage all report definitions to which you have been granted access.

- **My Oracle Support** tab provides access to software and patches.

2.7.3 Using Grid 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 Applicationsalso 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.

3

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

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

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. In most situations, you can use Fusion Applications Control, Oracle WebLogic Server Administration Console, or the Oracle WebLogic Scripting Tool (WLST) or Oracle Process Manager and Notification Server (OPMN) commands to start or stop Oracle Fusion Middleware components.

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, the WebLogic Server Administration Console, or Fusion Applications Control.	Section 3.3.6
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. Oracle WebLogic Server provides several ways to start and stop server instances. The method that you choose depends on whether you prefer using a command-line interface, and on whether you are using Node Manager to manage the server's life cycle.	Section 3.2 "Starting an Administration Server with a Startup Script" section in *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server* "Starting an Administration Server from the Windows Start Menu" section in *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server* "Starting an Administration Server with the java weblogic.Server Command" section in *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server* "Starting an Administration Server Using WLST and Node Manager" section in *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server* "Starting an Administration Server Using WLST Without Node Manager" section in *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server*
Node Manager	**Node Manager** is an Oracle WebLogic Server utility that enables you to start, shut down, and restart the Administration Server and the Managed Server instances. On Windows, Node Manager is configured to automatically start by default.	"Node Manager Overview" section in the *Oracle Fusion Middleware Node Manager Administrator's Guide for Oracle WebLogic Server* Task 3, "Start Node Manager"
Managed Server for an application	Managed Servers host business applications, application components, Web services, and their associated resources. Oracle WebLogic Server provides several ways to start and stop server instances. The method that you choose depends on whether you prefer using the Fusion Applications Control, Oracle WebLogic Server Administration Console, or a command-line interface, and on whether you are using Node Manager to manage the server's life cycle.	Section 3.2 "Starting Managed Servers with a Startup Script" section in *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server* "Starting Managed Servers with the java weblogic.Server Command" section in the *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server* "Starting and Stopping Managed Servers Using Fusion Middleware Control" section in *Oracle Fusion Middleware Administrator's Guide* "Start Managed Servers in a cluster" and "Shutdown servers in a cluster" in the *Oracle Fusion Middleware Oracle WebLogic Server Administration Console Online Help*

Table 3–1   (Cont.) Starting and Stopping

Component	Description	Documentation
Oracle Business Intelligence Suite	You can use Fusion Applications Control or the `opmnctl` command to start and stop Oracle Business Intelligence system components.	"Starting and Stopping Oracle Business Intelligence" chapter in the *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*
Oracle Enterprise Scheduler Service	You can use Fusion Applications Control to start and stop the Oracle Enterprise Scheduler Service.	Section 5.6
Oracle HTTP Server	You can use Fusion Applications Control or the `opmnctl` command to start and stop Oracle HTTP Server.	"Starting, Stopping, and Restarting Oracle HTTP Server" section in the *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server*
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, Node Manager, or a startup script.	Section 3.3.5.3 for Fusion Applications Control "Using Node Manager to Start Managed Servers in a WebLogic Domain or Cluster" section in *Oracle Fusion Middleware Oracle WebLogic Scripting Tool* "Starting Managed Servers with a Startup Script" section in *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server*
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 *Oracle Database 2 Day DBA*

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 `fastartstop` 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 Administration Server, use the following script from the Oracle WebLogic Server domain home directory, which is based on WLST commands:

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

where *DOMAIN_HOME* is located in the following locations:

```
(UNIX) APPLICATIONS_CONFIG/instance/domains/host/domain_name
(Windows) APPLICATIONS_CONFIG\instance\domains\host\domain_name
```

When the server successfully completes its startup process, it writes the following message to standard out (which, by default, is the command window):

```
<Notice> <WebLogicServer> <BEA-000360> <Server started in RUNNING mode>
```

Another alternative to starting the Administration Servers and Managed Servers for all the product families is the `fastartstop` utility. See Section 3.3.3.

See the following documentation resources to learn more about other methods for starting and stopping the Administration Servers:

- "Starting an Administration Server from the Windows Start Menu" section in the *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server*

- "Starting an Administration Server with the java weblogic.Server Command" section in *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server*

- "Starting an Administration Server Using WLST and Node Manager" section in *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server*

- "Starting an Administration Server Using WLST Without Node Manager" section in the *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server*

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 Oracle WebLogic Server Managed Servers, use the following script from the Oracle WebLogic Server domain home directory, which is based on WLST commands:

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

where `managed_server_name` specifies the name of the Managed Server and `admin_url` specifies the listen address (host name, IP address, or DNS name) and port number of the domain's Administration Server. See Section 2.4 for more information on locating the Administration Server port. When prompted, enter your user name and password.

When the server successfully completes its startup process, it writes the following message to standard out (which, by default, is the command window):

```
<Notice> <WebLogicServer> <BEA-000360> <Server started in RUNNING mode>
```

Another alternative to starting the Administration Servers and Managed Servers for all the product families is the `fastartstop` utility. See Section 3.3.3.

See the following documentation resources to learn more about other methods for starting and stopping the Managed Servers:

- "Starting Managed Servers with the java weblogic.Server Command" section in the *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server*

- "Starting and Stopping Managed Servers Using Fusion Middleware Control" section in *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server*

- "Start and Stop Servers" and various startup and shutdown procedures in the Cluster section of the *Administration Console Online Help*.

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 the Administration Servers"

- Task 3, "Stop the Managed Servers"

- Task 4, "Stop Oracle Business Intelligence"

- Task 5, "Stop the Oracle Identity Management Suite"

- 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 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.

Use the following script from the Oracle WebLogic Server domain home directory directory, which is based on WLST commands:

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

where `DOMAIN_HOME` is located in the following locations:

```
(UNIX) APPLICATIONS_CONFIG/instance/domains/host/domain_name
(Windows) APPLICATIONS_CONFIG\instance\domains\host\domain_name
```

When the server successfully completes its shutdown process, it writes the following message to standard out (which, by default, is the command window):

```
<Notice> <WebLogicServer> <BEA-000365> <Server state changed to SHUTTING DOWN>
```

Another alternative to stopping the Administration Servers and Managed Servers for all the product families is the `fastartstop` utility. See Section 3.3.3.

See the following documentation resources to learn more about other methods for stopping the Administration Servers:

- "Starting an Administration Server from the Windows Start Menu" section in the *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server*

- "Starting an Administration Server with the java weblogic.Server Command" section in *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server*

- "Starting an Administration Server Using WLST and Node Manager" section in *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server*

- "Starting an Administration Server Using WLST Without Node Manager" section in the *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server*

Task 3 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, use the following script from the Oracle WebLogic Server domain home directory, which is based on WLST commands:

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

where `managed_server_name` specifies the name of the Managed Server and `admin_url` specifies the listen address (host name, IP address, or DNS name) and port number of the domain's Administration Server. See Section 2.4 for more information on locating the Administration Server port. When prompted, enter your user name and password.

When the server successfully completes its shutdown process, it writes the following message to standard out (which, by default, is the command window):

```
<Notice> <WebLogicServer> <BEA-000365> <Server state changed to SHUTTING DOWN>
```

Another alternative to stopping Administration Servers and Managed Servers for all the product families is the `fastartstop` utility. See Section 3.3.3.

See the following documentation resources to learn more about other methods for stopping the Managed Servers:

- "Starting Managed Servers with the java weblogic.Server Command" section in the *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server*

- "Starting and Stopping Managed Servers Using Fusion Middleware Control" section in *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server*

- "Start and Stop Servers" and various startup and shutdown procedures in the Cluster section of the *Administration Console Online Help*.

Task 4 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 5 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 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` utility uses a registry file named `SetupDetails.xml` in the following `fusionapps` Middleware subdirectories to determine the entities to start and stop:

```
(UNIX) FA_MW_HOME/admin/
(Windows) FA_MW_HOME\admin
```

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 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 Business Intelligence, 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

The following shows the syntax of the `fastartstop` utility:

```
(UNIX) fastartstop.sh -start|stop -all|-domains domain_name, ...domain_nameN
-username user_name -password password -appsbase APPLICATIONS_BASE
(Windows) fastartstop.cmd -start|stop -all|-domains domain_name, ...domain_nameN
-username user_name -appsbase APPLICATIONS_BASE
```

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

Table 3–2   Options for the fastartstop Script

Options	Description
`-appsbase`	Specify the path for the Oracle Fusion Applications installation.
`-all|-domains`	Use `-all` for to start or stop all domains, and use `-domains` for specific domains. If using the `-domains` option, delineate the domains names with a comma and space (`domain_name`, `domain_name`).
`-start|stop`	Specify the path for the Oracle Fusion Applications installation.
`-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 `SSUtil.log` in the following `fusionapps` Middleware subdirectories:

```
(UNIX) FA_MW_HOME/admin/lcm/ad/bin
(Windows) FA_MW_HOME\admin\lcm\ad\bin
```

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

```
<?xml version = '1.0' encoding = 'UTF-8'?>
```

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

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 fastartstop Utility"

- Task 5, "Start Oracle Business Intelligence"

- Task 6, "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

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

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

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

Task 4 Run the fastartstop Utility

From the `CommonDomain` domain in the environment, run `fastartstop` to start all the domains or specific domains:

```
(UNIX) fastartstop.sh -start -all|-domains domain_name, ...domain_nameN
-username user_name -password password -appsbase APPLICATIONS_BASE
(Windows) fastartstop.cmd -start -all|-domains domain_name, ...domain_nameN
-username user_name -appsbase APPLICATIONS_BASE
```

For more information about the syntax, see Table 3–2.

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**.

3.3.3.4 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"
- Task 4, "Stop Oracle Business Intelligence"
- Task 5, "Stop the Oracle Identity Management Suite"
- Task 6, "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 domain 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

From the `CommonDomain` domain in the environment, run `fastartstop` to stop all the domains or specific domains:

```
(UNIX) fastartstop.sh -stop -all|-domains domain_name, ...domain_nameN
-username user_name -password password -appsbase APPLICATIONS_BASE
(Windows) fastartstop.cmd -stop -all|-domains domain_name, ...domain_nameN
-username user_name -appsbase APPLICATIONS_BASE
```

For more information about the syntax, see Table 3–2.

Task 4 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 5 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 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*.

3.3.4 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 Oracle WebLogic Server Clusters with a Startup Script
- 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 Oracle WebLogic Server Clusters with a Startup Script

See "Starting Managed Servers with a Startup Script" section in *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server* to start the Managed Servers, one at a time, in the cluster.

3.3.5.2 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*.

3.3.5.3 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.

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

3.3.6 Starting and Stopping Specific Applications

You can start and stop applications using the WLST command line, the WebLogic Server Administration Console, or Fusion Applications Control.

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
- Starting and Stopping all Application Instances for a Cluster Using Fusion Applications Control

3.3.6.1 Starting and Stopping Java EE Applications Using WLST

To start or stop applications with the WLST command line, use the following commands from the `fusionapps` Middleware directory:

```
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

To start or stop an application using Fusion Applications Control:

1. From the navigation pane, expand the product family, then **Fusion Applications**, and then the cluster application.

2. Select the application instance you want to stop.

3. From the **Fusion J2EE Application** menu, choose **Control,** then **Start Up** or **Shut Down.**

3.3.6.3 Starting and Stopping all Application Instances for a Cluster Using Fusion Applications Control

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

1. From the navigation pane, expand the product family, then **Fusion Applications**, and then the cluster application.

2. 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.

You can change the port numbers for some Oracle Fusion Middleware components from Oracle Fusion Applications, Oracle WebLogic Server Administration Console, or the command line.

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 *Oracle Fusion Middleware Administrator's Guide*
Oracle Business Intelligence Suite	"Using Fusion Middleware Control to Scale System Components" section in the *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*
Oracle HTTP Server	"Manage Ports" section in the *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server*
Oracle Enterprise Content Management Suite	"Modifying Server Configuration Parameters for Oracle UCM Content Server" section in the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory* for information about changing the for Oracle Universal Content Management Content Server
Oracle Internet Directory	"Configuring Server Properties" section or the "Setting System Configuration Attributes by Using ldapmodify" section in the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*
Oracle Virtual Directory	"Configuring Oracle Virtual Directory to Listen on Privileged Ports" section in the *Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory*
Node Manager	"Overview of Node Manager Configuration" section in the *Oracle Fusion Middleware Node Manager Administrator's Guide for Oracle WebLogic Server*

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	In an Oracle Fusion Applications installation. The super user is described in the next row of this table.	Section 4.8.5.2
Super User	The super user has the following privileges: ■ 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. You can change the credentials using Fusion Applications Control or Oracle Platform Security Services (OPSS) scripts. 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 *Oracle Fusion Applications Installation Guide* for information about this page. Installation establishes the same username and password for all the domains.	Section 4.8.5.2 to change the administrative password "Managing the Credential Store" section in the *Oracle Fusion Middleware Application Security Guide*
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 *Oracle Fusion Middleware Application Security Guide*

Table 3–4   (Cont.) Changing Passwords

Password Account	Description	Documentation for Changing the Passwords
Application 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 *Oracle Fusion Applications Security Hardening Guide*
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 *Oracle Fusion Applications Installation Guide* for information about using this page.	"Specify Node Manager Username and Password" section in the *Oracle Fusion Middleware Node Manager Administrator's Guide for Oracle WebLogic Server*
BISystemUser	The `BISystemUser` account provides access to the Oracle Business Intelligence system components.	"Default Users and Passwords" section in the *Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition*

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 14.

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 DIAGNOSTIC+TUNING (default) or DIAGNOSTIC to enable Automatic Database Diagnostic Monitor (ADDM).
STATISTICS_LEVEL	Set this parameter to TYPICAL (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 legs

- 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.	*Oracle Fusion Applications Patching Guide*
Oracle Fusion Middleware	OPatch patches any Oracle Fusion Middleware component, except Oracle WebLogic Server. For Oracle WebLogic Server, use Smart Update.	*Oracle Fusion Middleware Patching Guide* for information about patching middleware artifacts *Oracle Smart Update Installing Patches and Maintenance Packs* for information about patching Oracle WebLogic Server
Oracle Database	Oracle Universal Installer and OPatch manage the patching for Oracle database and third-party software.	*Oracle Universal Installer and OPatch User's Guide for Windows and UNIX*

4

Securing Oracle Fusion Applications

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*.

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
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
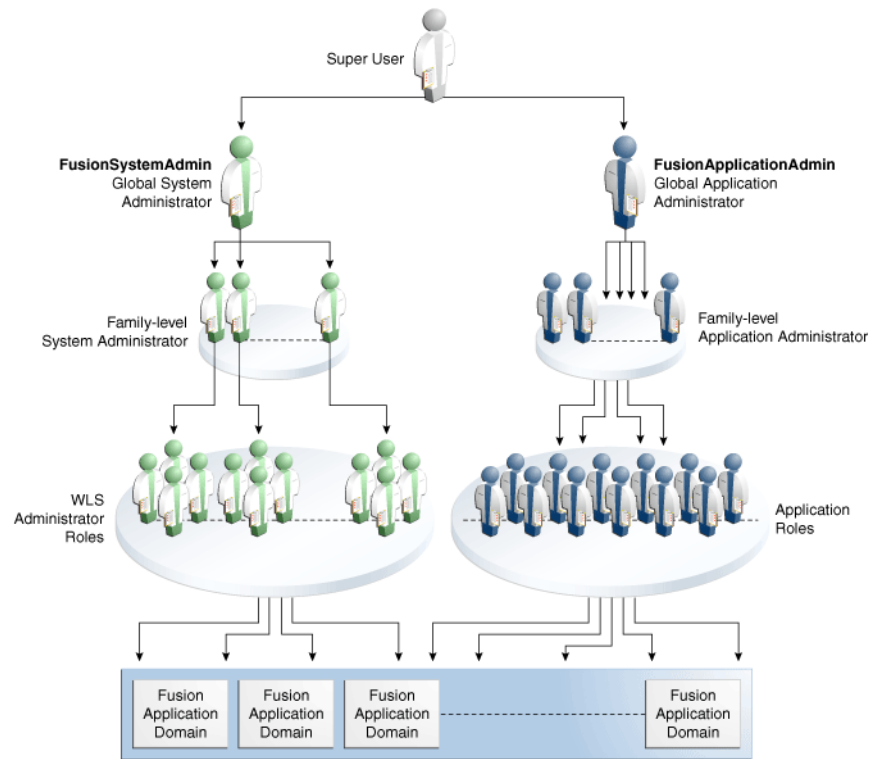
4.3.1.2 Two Types of Users During Provisioning

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, Flow Designer 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 (OVDAuthenticator) 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*

- "User, Account, and Entitlement Provisioning" section in *Oracle Fusion Middleware Integration Overview for Oracle Identity Management Suite*.

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 some of the most frequently uses of Oracle Authorization Policy Manager, see the "Some Frequently Used Operations" 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 "WLST Security Commands" section 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/adflogginghandler.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.

Note 1: The files `Common-Model.jar` and `DataSecurity-Model.jar` are expected to be in the paths indicated above, but, depending on the environment, the 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"
```

Note 2: `DSDataMigrator` 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>
```

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>
```

- `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.

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)*:

- Mapping Application Roles to an External Role

- 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 Application Security Artifacts" 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 _DATA.
Job	Enterprise role mapped to application role. Typically, the name of this role has the suffix _JOB.
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 "Basic Security Artifacts" 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
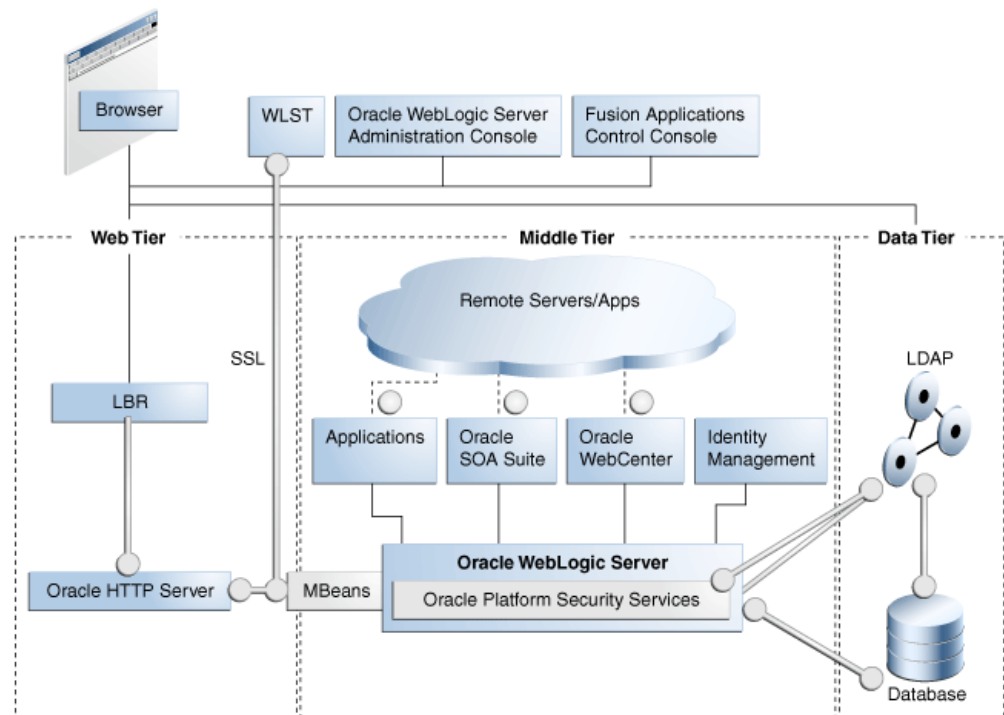
In the figure, the vertical broken lines represent firewalls. The circles represents listeners that can be SSL-enabled for secure communication. As the figure shows, all critical communication paths can be protected with SSL regardless of the tier(s) involved.

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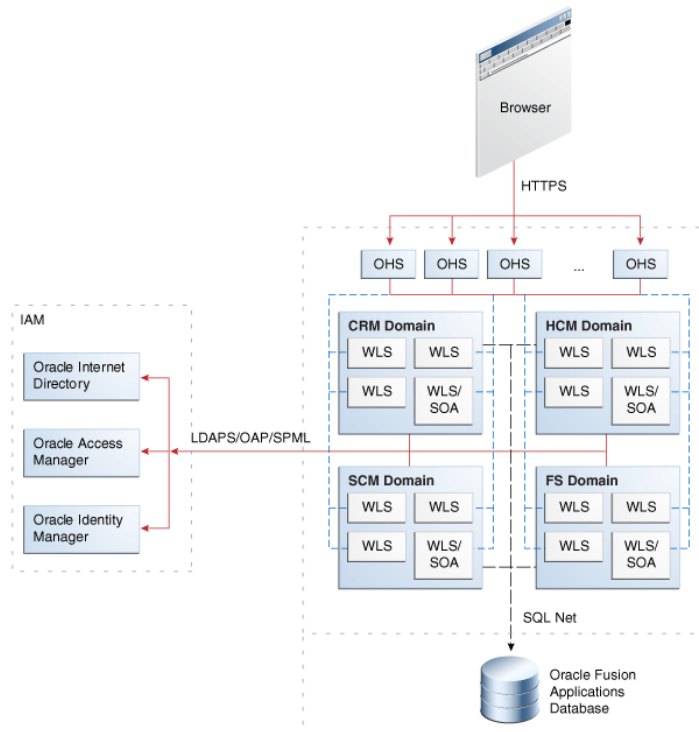
> **Notes:**
>
> ■   The diagram shows some representative domains.
>
> ■   OAP is the Oracle Access Protocol used by Oracle Access Manager. LDAPS is the LDAP-over-SSL protocol.

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/CommondDomain_
      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, 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*.

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

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 in the *Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition*.

To locate the guide:

- Point your browser to the Oracle Business Intelligence site at `http://www.oracle.com/technology/documentation/bi_ee.html`.
- Locate your version of Oracle Business Intelligence Enterprise Edition, and click **View Library**.
- At the Documentation Library page, select the Documentation tab.
- Select the Security Guide.

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 the "Configuring SSL for Oracle ECM Applications" section in the *Oracle Fusion Middleware Installation Guide for Oracle Enterprise Content Management Suite*.

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://mywlsserver.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 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 *Oracle Fusion Applications Security Hardening 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

- 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

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 of 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.

- ■ 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 10*g* Release 2 or Oracle Database 11*g* Release 1
- Oracle Enterprise Manager Grid Control 10*g* Release 4
- 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. Refer to the Oracle Fusion Applications system requirements and supported platforms documentation for details.

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.

d. Click **Datafiles**.

e. On the Datafiles page, enter `temp` in the **Object Name** search box. Press **Go**. Two temp spaces are displayed:

f. Edit each temp space in turn by selecting its radio button and clicking **Edit**. Change the value for each space as shown here:

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 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.

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.

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:

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

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:

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.

> **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.

> **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.

3. 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**.

4. 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 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.

4.9.5 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

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 *Oracle Fusion Middleware Security and Administrator's Guide for Web Services*

Table 4–6   (Cont.) Common LPA Operations

Operation	Reference
Attach "no behavior" policies for Web Services and Web Service Clients	"Disabling a Globally Attached Policy" section in the *Oracle Fusion Middleware Security and Administrator's Guide for Web Services*
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 *Oracle Fusion Middleware Security and Administrator's Guide for Web Services*
Troubleshoot Web service security	"Diagnosing Problems" chapter in the *Oracle Fusion Middleware Security and Administrator's Guide for Web Services*

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 *Oracle Fusion Middleware Security and Administrator's Guide for Web Services*
Create GPAs using Fusion Middleware Control and WLST	"Creating and Managing Policy Sets" section in the *Oracle Fusion Middleware Security and Administrator's Guide for Web Services*
Disable GPAs	"Enabling and Disabling a Policy Set" section in the *Oracle Fusion Middleware Security and Administrator's Guide for Web Services*
Determine what policies are enforced when both LPA and GPA are defined	"Calculating the Effective Set of Policies" section in the *Oracle Fusion Middleware Security and Administrator's Guide for Web Services*
View policies attached to a Web service	"Viewing the Policies That are Attached to a Web Service" section in the *Oracle Fusion Middleware Security and Administrator's Guide for Web Services*
Validate a policy set	"Validating a Policy Set" section in the *Oracle Fusion Middleware Security and Administrator's Guide for Web Services*

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 IDCSS 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 IDCSS 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 New 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.

The obtain this file, the administrator perform the following steps:

1. Run 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. Edit 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.JpsXmlEnterprise RoleImpl`.

3. Deliver the edited `jazn-data.xml` to the developer, who copies it to the directory `<jdevapphome>/src/META-INF`.

Then, within Oracle JDeveloper, the developer perform the following steps:

1. Define a data source to point to the installed database security policy store.

2. Use 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. Create 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 with Oracle JDeveloper are completed, the developer hands the file `jazn-data.xml` back to the system administrator.

4.13.3 Migrating to a Staging Environment

In this phase, the administrator performs the following steps:

1. Reconcile GUIDs in the Oracle Fusion Data Security policy store with GUIDs in the file `jazn-data.xml` by running the command `DSDataMigrator` as described in Section 4.4.3. This operation does not modify the file `jazn-data.xml` but updates GUIDs in data security policies.

   > **Note:** This step is needed only if the developer has customized data security.

2. If the staging environment already has a stripe for the application and that stripe is to be preserved, merge 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.

3. Pack the file `jazn-data.xml` with the application and deploy the application to the staging environment.

   > **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 New Environment

In this phase, the administrator performs the following steps:

1. Migrate the staging policy store to the new environment policy store. It is assumed that the store has undergone some changes since the customization was tested in the staging environment, and that those changes must be preserved. The administrator merges the customized, pillar-level `jazn-data.xml` with the new policy store using Oracle Authorization Policy Manager as described in the "Upgrading Oracle FusionAuthorization Policy Manager" chapter of the *Oracle Fusion Middleware Authorization Policy Manager Administrator's Guide*.

2. Reconcile 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.

3. Deploy the application with the customized `jazn-data.xml` to the new environment.

5

Managing Oracle Enterprise Scheduler Service and Jobs

This chapter describes how to configure and manage Oracle Enterprise Scheduler Service and Oracle Enterprise Scheduler jobs.

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
- 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
- Managing Job Metadata
- Managing Work Assignments and Workshifts
- Managing Oracle ADF Connections
- Managing Application and Role Policies
- Managing Oracle Enterprise Scheduler Web Services
- Managing Logging for Oracle Enterprise Scheduler

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 (work shifts), 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, work shifts and purge policies.

Finally, 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.

5.2 Navigating to the Oracle Enterprise Scheduler Administration Pages

To navigate to Oracle Enterprise Scheduler administrative tasks:

1. From the navigation pane, expand the farm and then **Scheduling Services**, and select the Oracle Enterprise Scheduler component.
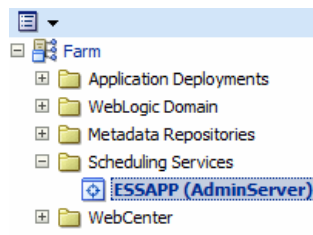
The Scheduling Service home page displays. See Section 5.3 for further information about the contents of the home page.

2. Select the **Scheduling Service** menu.
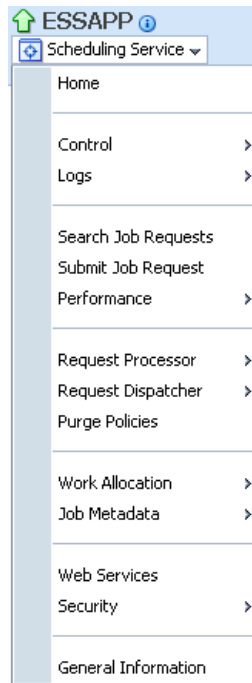
The **Scheduling Service** menu displays the following options.

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: ■ **Start Up:** This option runs the Oracle Enterprise Scheduler server. For more information, see Section 5.6. ■ **Shut Down:** This option shuts down the Oracle Enterprise Scheduler server. For more information, see Section 5.6.
Logs	The **View Log Message** option displays the Log Messages page for viewing the contents of event log files. See "Managing Log Files and Diagnostic Data" chapter in *Oracle Fusion Middleware Administrator's Guide* for more information. The log files are stored in the following directories: ■ *DOMAIN_HOME*/servers/*SERVER_HOME*/logs ■ *DOMAIN_HOME*/servers/*SERVER_HOME*/logs/apps ■ *DOMAIN_HOME*/servers/*SERVER_HOME*/logs/owsm/msglogging/diagnostic.log ■ *DOMAIN_HOME*/servers/*SERVER_HOME*/sysman/log/emoms.log
Search Job Requests	Allows you to search for job requests according to particular parameters. For more information, see Section 5.7.2.
Submit Job Request	Allows you to submit an execution request for a particular job. For more information, see Section 5.7.1.

Menu Option	Description
Performance	Displays the following options: ■ **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	Displays the following options: ■ **Start/Stop**: This option causes the request processor to start or stop running. For more information, see Section 5.6. ■ **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	Displays the following options: ■ **Start/Stop**: This option causes the request dispatcher to start or stop running. For more information, see Section 5.6. ■ **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	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 5.9.
Work Allocation	Displays the following options: ■ **Work Assignments**: This option displays the Work Assignments page which allows you to create, edit and delete work assignments. For more information, see Section 5.11.1. ■ **Workshifts**: This option displays the Workshifts page which allows you to create, edit and delete workshifts. For more information, see Section 5.11.2. ■ **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 5.11.3.
Job Metadata	Displays the following options: ■ **Job Definitions**: This option displays the Job Definitions page, which allows you to manage job definitions. See Section 5.10.1 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 5.10.2. ■ **Incompatibilities**: This option displays the Incompatibilities page, which allows you to create and edit incompatibilities for a given application. For more information, see Section 5.10.3.
ESS Connections	Displays a page that enables defining connections to Oracle ADF. For more information, see Section 5.12.

Menu Option	Description
Web Services	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 5.14.
Security	Displays the following options:
	■ **Application Policies**: This option displays the authorization policies that an application relies upon for controlling access to its resources. See Section 5.13 for further information.
	■ **Application Roles**: This option displays the roles used by security aware applications specific to the application. See Section 5.13 for further information.
General Information	Displays general details about the Oracle Enterprise Scheduler instance. For more information, see Section 5.5.

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
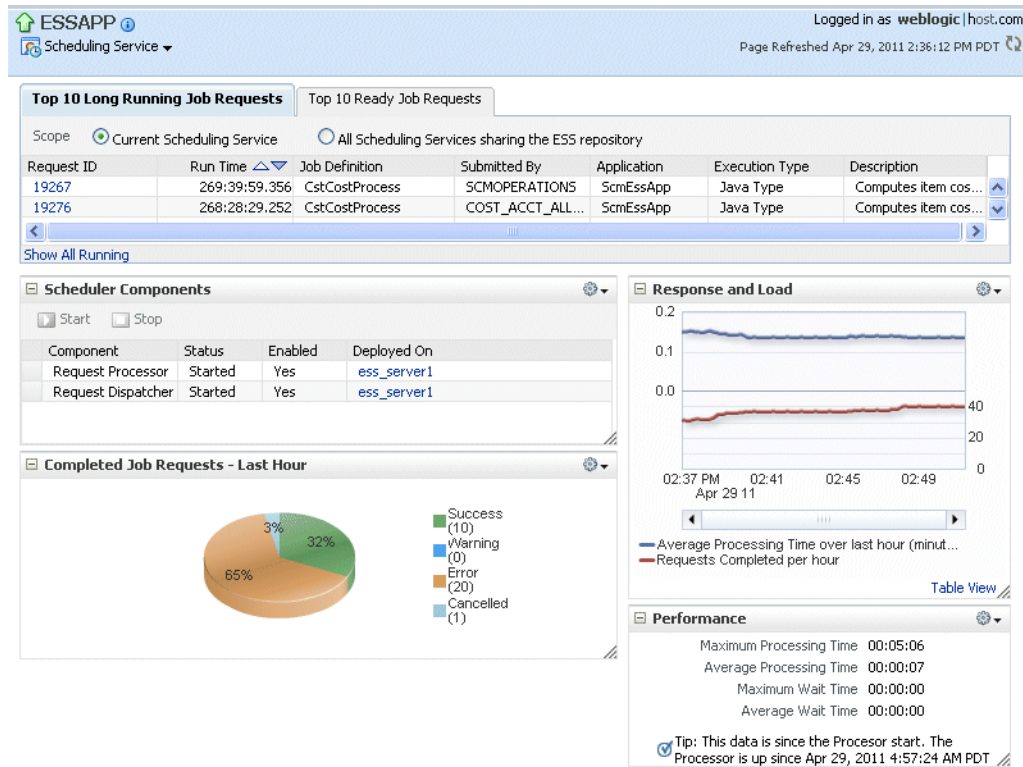
These pages contain the following regions:

■ Top Ten Long Running Requests and Top Ten Ready Job Requests Regions

- Scheduler Components Region
- Completed Requests Region
- Response and Load Region
- Performance Region

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.

2. **Configure the request processor**. Configure the Oracle Enterprise Scheduler component which receives and manages job requests.

3. **Configure the request dispatcher**. Configure the Oracle Enterprise Scheduler component that polls the request processor for job requests and dispatches jobs.

This section contains the following topics:

- Configuring an Oracle Enterprise Scheduler Cluster

- Configuring a Request Processor

- Configuring a Request Dispatcher

5.4.1 Configuring an Oracle Enterprise Scheduler Cluster

An Oracle Enterprise Scheduler cluster is created as part of an Oracle Fusion Applications installation. You may want to expand the cluster to handle a larger load. You can add new cluster nodes to an Oracle Enterprise Scheduler cluster from the Oracle WebLogic Server Console.

When scaling out an Oracle Enterprise Scheduler cluster with an additional managed server, it may be undesirable for the new server to immediately begin processing requests in the default work assignment. This could be the case if all other servers have work assignments bound in standard mode. If one or more of the running servers is using the default work assignment, however, then the present work allocation is compatible with a server using the default work assignment.

Oracle Enterprise Scheduler provides protection for a new server so that work allocations can be configured before the server begins processing jobs. When a newly added server starts for the first time, Oracle Enterprise Scheduler determines whether the default work assignment can be used. If it cannot, Oracle Enterprise Scheduler binds a pre-seeded internal work assignment (ESSInternalWA) which includes a health check job. You can use the health check job to check the server, and unbind the internal work assignment and then bind your own work assignment as required.

> **Note:** When determining whether the default work assignment can
> be used, Oracle Enterprise Scheduler considers all running servers in
> the group without regard to the applications that are deployed to the
> servers. Servers that are down are not taken into account.

For more information about work assignments, see Section 5.11.1.

For more information about expanding a cluster using Oracle WebLogic Server
Console, see *Oracle Fusion Middleware Oracle WebLogic Server Administration Console
Online Help*.

5.4.2 Configuring a Request Processor

A work assignment must be bound to a request processor before it can process job
requests. A binding can be made in one of two modes: standard (the default) or
exclusive.

Standard binding mode means the request processor can process job requests as
defined by the specialization rules when an active workshift is defined. If a job request
is specialized to two different work assignments, either of those work assignments or
the default work assignment can process the job request.

When using exclusive binding mode, job requests specialized to the work assignment
are processed exclusively by that work assignment when it is active. These job requests
are excluded from all other work assignments, including the default work assignment.
If the work assignment does not have an active workshift, then the job request can be
processed by another work assignment.

Requirements for binding a work assignment are as follows:

1. The work assignment must be enabled, meaning its active flag must be set.

2. The work assignment must have at least one workshift.

3. Each workshift in the work assignment must have a thread allocation of at least
   one.

4. If the workshift includes a schedule, the following must be true:

   ■ The schedule must be active, meaning it is not expired.

   ■ The workshift duration must be at least one.

5. A work assignment can be bound to a particular server at most one time.

6. A work assignment can be bound to any number of servers in the group, but all
   bindings must use the same mode. Within a group, a work assignment cannot be
   bound in standard mode on one server and exclusive mode on another server.

For more information about work assignments, see Section 5.11.

To configure a request processor:

1. From the navigation pane, expand the **Scheduling Services** folder and select the
   Oracle Enterprise Scheduler application.

2. From the Scheduling Services menu, select **Request Processor** and then select
   Configure.

3. Under the Enable Request Processor field, enable or disable the request processor.
   Select the **Enable Request Processor** checkbox to turn on request processing.

4. Optionally, under the Work Assignment Bindings section, bind a work assignment to the request processor.

   - **Add**: To bind a work assignment to the request processor, click **Add**.

     From the Select Work Assignment popup window, select a work assignment from the list and click **OK**.

   - **Remove**: To delete a work assignment binding from the job request processor, select a work assignment from the list of bindings and click **Remove**.

   - **Exclusive Binding**: Select to bind a work assignment exclusively to the request processor. The binding applies even if the work assignment is not enabled in the work assignment definition, or if the server is down. If a work assignment is bound to multiple servers, all the servers to which it is bound must use the same binding mode.

5. Optionally, under the Advanced Configuration section, configure the number of threads and the starvation threshold in minutes.

   - **Thread Count**: In the text field, enter the number of threads you want the processor to handle.

   - **Starvation threshold (min)**: In the text field, enter the starvation threshold in minutes. The starvation threshold specifies the wait time for job requests.

6. Click **Apply** to save your changes.

5.4.3 Configuring a Request Dispatcher

Use the Configure Request Dispatcher page to enable or disable the job request dispatcher. You can also configure the polling interval for the request dispatcher.

Requests remain in waiting state in the Oracle Enterprise Scheduler repository, until the request dispatcher polls the repository for requests that are ready to run. After polling the repository, the dispatcher sets all requests to ready state. The request processor takes over control of the job requests once they have been placed in ready state.

The default maximum polling interval is 15 seconds.

To configure a request dispatcher:

1. From the navigation pane, expand the **Scheduling Services** folder and select the Oracle Enterprise Scheduler application.

2. From the Scheduling Services menu, select **Request Dispatcher** and then select **Configure**.

3. Under the Enable Request Dispatcher field, enable or disable the request processor.

   Select the **Enable Request Dispatcher** checkbox to turn on request dispatching.

4. Optionally, under the Advanced Configuration section, configure the polling interval for the request dispatcher. In the text **Maximum Poll Interval** field, enter the minimum interval at which the dispatcher polls for requests.

5. Click **Apply** to save your changes.

5.5 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 **General 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.

5.6 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.

This section contains the following topics:

- Starting and Stopping an Oracle Enterprise Scheduler Service Instance

- Starting and Stopping a Request Processor or Dispatcher

5.6.1 Starting and Stopping an Oracle Enterprise Scheduler Service Instance

The Control menu for the scheduling service allows you to start an Oracle Enterprise Scheduler instance and consequently begin managing scheduled job requests.

   > **Note:** Stopping an Oracle Enterprise Scheduler instance or component is not recommended.

To start up an instance of Oracle Enterprise Scheduler:

1. From the Scheduling Service menu, select **Control**.

    2. Select **Start Up**.

To shut down an instance of Oracle Enterprise Scheduler:

1. From the Scheduling Service menu, select **Control**.

2. Select **Shut Down**.

5.6.2 Starting and Stopping a Request Processor or Dispatcher

You can start or stop a configured request processor or dispatcher from the Scheduling Service menu.

> **Note:** Stopping an Oracle Enterprise Scheduler instance or component is not recommended.

To stop a request processor or dispatcher:

1. From the navigation pane, expand the **Scheduling Services** folder and select the Oracle Enterprise Scheduler application.

2. Start the request processor or request dispatcher as follows.

- To start the request processor: From the Scheduling Services menu, select **Request Processor** and then select **Start**.

- To start the request dispatcher: From the Scheduling Services menu, select **Request Dispatcher** and then select **Start**.

3. Stop the request processor or request dispatcher as follows.

- To stop the request processor: From the Scheduling Services menu, select **Request Processor** and then select **Stop**.

- To stop the request dispatcher: From the Scheduling Services menu, select **Request Dispatcher** and then select **Stop**.

4. When prompted as to whether or not to stop the processor or dispatcher, click **OK**.

5.7 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

- Viewing Job Request Details

- Holding and Resuming Oracle Enterprise Scheduler Service Job Requests

- Cancelling Oracle Enterprise Scheduler Job Requests

5.7.1 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

5.7.1.1 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 5.10.1.2.

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 5–9 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 5.7.1.2

5.7.1.2 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 5.10.1.2.

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.

5.7.2 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

5.7.2.1 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.

5.7.2.2 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 5.7.2.1.

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.

5.7.2.3 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.**

5.7.3 Viewing Job Request Details

Once you submit a search for job requests, you can view detailed information about an individual job request by clicking the job request ID or the request parent ID. If the job is in an error state, an information box displays details regarding the error that occurred.

You can take the following actions on an individual job request:

- Display log information for the job request.

- Display the job set and all child job requests.

- Submit a job or job set request with parameters similar to those of a given job request.

- Recover an incomplete job request.

To view job request details:

1. Search for the relevant job requests as described in

2. In the table displaying the job request search results, select the job request whose details you want to view.

3. To view job request details, click the job request ID. Alternatively, click the parent ID associated with the job request to view the details of the job set with which the job is associated.

   In the job request details page, you can take any of the following actions.

   - **Request Log**: Select **Action** and then select Request Log to display log information for the job request.

   - **Request Tree**: Select **Action** and then select Request Tree to display the parent job set along with all related child jobs.

   - **Cancel**: Select **Action** and then select Cancel to cancel the job request.

   - **Hold**: Select **Action** and then select Hold to temporarily suspend the job request.

   - **Resume**: Select **Action** and then select Resume to resume a suspended job request.

   - **Recover Stuck Request**: Select **Action** and then select Recover Stuck Request to recover an incomplete job request that cannot continue.

   - **Change Priority**: Select **Action** and then select Change Priority to raise or lower the priority of the job request. Job requests with a higher priority will be dispatched prior to job requests with lower priority.

   - **Change Schedule**: Select **Action** and then select Change Schedule to assign a different schedule to the job request.

   - **Submit Like**: Select **Action** and then select Submit Like to submit a job or job set with parameters similar to this one.

5.7.4 Holding and Resuming Oracle Enterprise Scheduler Service Job Requests

You may hold or resume a job request following its submission.

To hold and resume a job request:

1. Search for a job request.

2. In the table of job request search results, select the job request you want to hold.

3. Click the **Hold** button.

4. To resume the paused job request, select the job request and click **Resume**.

5.7.5 Cancelling Oracle Enterprise Scheduler Job Requests

You may cancel an executable job request following its submission. Executable job requests include job sets and job requests submitted without schedules. A job request can be canceled when it is not in a state of execution.

When initiating a job request cancellation, the resulting cancellation state depends on the current state of the job request. A job request that is not executing is set directly to `CANCELLED` state. A job request in `RUNNING`, `PAUSED` or `COMPLETED` states are placed in `CANCELLING` state following the initiation of request cancellation.

The final state of a request depends on the processing stage the job request at the point when cancellation has been initiated.

Table 5–3 displays the cancellation states applied to each executable request depending on its state when initiating cancellation.

Table 5–3   Executable Requests and Cancellation States

Job State When Initiating Cancellation	New Cancellation State
WAIT	CANCELLED
READY	CANCELLED
HOLD	CANCELLED
BLOCKED	CANCELLED
PENDING_VALIDATION	CANCELLED
ERROR_AUTO_RETRY	CANCELLED
PAUSED	CANCELLING
RUNNING	CANCELLING
COMPLETED	CANCELLING
Any terminal state	No state change
CANCELLING	No state change
ERROR_MANUAL_RECOVERY	No state change

An executable request is processed in three major stages: pre-processor, executable and post-processor. Requests can be cancelled at any of these stages. If the job request cannot complete all three stages, it transitions to `CANCELLED`.

This section contains the following topics:

- Initiating Cancellation During Pre-Process Handler Execution
- Initiating Cancellation During Synchronous Java Job Execution
- Initiating Cancellation During Asynchronous Java Job Execution
- Initiating Cancellation During PL/SQL Job Execution
- Initiating Cancellation During Spawned Job Execution
- Cancelling Parent Job Requests
- Cancelling a Child Job Request

5.7.5.1 Initiating Cancellation During Pre-Process Handler Execution

When initiating cancellation during the pre-processor stage, the state to which the job request transitions depends on the status returned by the pre-processor.

Table 5–4 displays the job state following cancellation initiation.

Table 5–4   Pre-processor States and States Following Cancellation Initiation

State Returned by the Pre-Processor	State Following Cancellation Initiation
PROCEED	CANCELLED
DELAY	CANCELLED
WARN	CANCELLED
SYSTEM_ERROR	ERROR
BIZ_ERROR	ERROR

5.7.5.2 Initiating Cancellation During Synchronous Java Job Execution

When initiating cancellation during the job execution stage, the state to which the request transitions depends on the manner of completion.

Table 5–5 displays the job state following cancellation initiation.

Table 5–5   Returned Synchronous Java Job States and States Following Cancellation Initiation

State Returned by the Job	State Following Cancellation Initiation
Normal return (success)	CANCELLED
ExecutionWarningException	CANCELLED
ExecutionPausedException	CANCELLED
ExecutionBizErrorException	ERROR
ExecutionErrorException	ERROR

5.7.5.3 Initiating Cancellation During Asynchronous Java Job Execution

When initiating cancellation during the job execution stage, the state to which the request transitions depends on the manner of completion.

Table 5–6 displays the job state following cancellation initiation.

Table 5–6   Returned Asynchronous Java Job States and States Following Cancellation Initiation

State Returned by the Job	State Following Cancellation Initiation
SUCCESS	CANCELLED
PAUSE	CANCELLED
WARNING	CANCELLED
ERROR	ERROR
CANCEL	CANCELLED
UPDATE	No state change. Wait for further updates.
BIZ_ERROR	ERROR
ERROR_MANUAL_RECOVERY	ERROR_MANUAL_RECOVERY

5.7.5.4 Initiating Cancellation During PL/SQL Job Execution

If the job executable is executing, Oracle Enterprise Scheduler attempts to kill the running RDBMS Scheduler job. If the job is successfully killed, the request transitions to `CANCELLED` state. If the RDBMS Scheduler job completes before it can be killed, the state to which the job request transitions depends on the result of the job execution.

Table 5–7 displays the job state following cancellation initiation.

Table 5–7   Returned PL/SQL Job States and States Following Cancellation Initiation

State Returned by the Job	State Following Cancellation Initiation
SUCCESS	CANCELLED
WARNING	CANCELLED
PAUSE	CANCELLED
BIZ_ERROR	ERROR
SYSTEM_ERROR	ERROR

5.7.5.5 Initiating Cancellation During Spawned Job Execution

If the spawned job is executing, Oracle Enterprise Scheduler attempts to kill the running process. If the process is successfully killed, the request transitions to a `CANCELLED` state. If the process completes before it can be killed, the state to which the request transitions depends on the result of the process execution.

Table 5–8 displays the job state following cancellation initiation.

Table 5–8   Returned Spawned Job States and States Following Cancellation Initiation

State Returned by the Job	State Following Cancellation
SUCCESS	CANCELLED
WARNING	CANCELLED
PAUSE	CANCELLED
BIZ_ERROR	CANCELLED
SYSTEM_ERROR	CANCELLED

5.7.5.6 Cancelling Parent Job Requests

A parent job request that is not executable must be in `WAIT` state to be cancelled. It then proceeds to `CANCELLING`. The cancellation operation propagates to all eligible child job requests. For example, job set steps are cancelled when the job set is canceled, sub-requests are cancelled when the parent request is cancelled and recurring executable job requests are cancelled when the recurring parent is cancelled. If a child job request is executable, it is subject to the rules described in the preceding sections. When all the child requests have completed, the parent request transitions to `CANCELLED` state.

5.7.5.7 Cancelling a Child Job Request

A child job request may be an executable or a parent job request. When cancelling a sub-request that completes in `CANCELLED` or another terminal state, the parent job request resumes as usual as long as the parent job request has not been cancelled as well. The state of the sub-request does not affect the state of the parent request.

When cancelling a step within a job set, the job set transitions to a `CANCELLED` state when the job set step transitions to a `CANCELLED` state. However, the job set may not revert to `CANCELLED` state if another job set step results in an error.

To cancel a job request:

1. Search for a job request.

2. In the table of job request search results, select the job request you want to cancel.

3. Click the **Cancel** button.

5.8 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

5.8.1 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.

5.8.2 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.

5.9 Managing Purge Policies

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.

This section contains the following topics:

- Setting Up a New Purge Policy

- Physically Purging Job Requests from the Database

- Viewing Purge Policies

- Updating a Purge Policy

- Deleting a Purge Policy

5.9.1 Setting Up a New Purge Policy

A purge policy determines which job requests are to be purged and which retained. Defining a purge policy involves:

- Selecting the jobs to be purged: Selection criteria include the related application or product, a particular job definition or job type, the job submitter or a maximum number of requests.

- Specifying retention criteria: Decide how long job requests are to be retained depending on their status.

- Specifying purge frequency: Decide how often you want the purge policy to run.

Note: After a purge policy has run, it is necessary to physically delete purged job requests from the database. For more information, see Section 5.9.2.

To set up a new purge policy:

1. From the navigation pane, expand the **Scheduling Services** folder and select the Oracle Enterprise Scheduler application.

2. From the Scheduling Services menu, select **Purge Policies**.

3. To configure a new policy, click **Setup New**.

   In the Setup Purge Policy page, configure the purge policy.

4. In the Description text field, enter a description for the purge policy.

5. In the Request Criteria for Purge section, configure the characteristics of the job requests to be purged.

   - **Application**: From the dropdown list, select the application for which you want to create a purge policy.

   - **Product**: Enter the name of the product.

   - **Job Definition**: Click the browse button next to the Job Definition text field.

     In the Select Job Definition window, enter the name of the job definition in the Job Definition Name text field.

     Click the search button, and select the relevant job definition from those that display. Click **OK**.

   - **Execution Type**: From the dropdown list, select the job type required: `All`, `JAVA_TYPE`, `SQL_TYPE` or `PROCESS_TYPE`.

   - **Submitted by**: Enter the name of the job submitter.

6. In the Retention Criteria for Purge section, configure the characteristics of the job requests to be retained.

   - **Default Retention Period (in days)**: Enter the default period, in days, during which requests are to be retained.

   - **Retention Period - Success (in days)**: Enter the period, in days, during which successful job requests are retained.

   - **Retention Period - Error (in days)**: Enter the period, in days, during which errored job requests are retained.

   - **Retention Period - Warning (in days)**: Enter the period, in days, during which job requests ending in a warning status are retained.

   - **Retention Period - Cancelled (in days)**: Enter the period, in days, during which cancelled job requests are retained.

7. In the Advanced section, select **Maximum requests to be processed**. Enter the maximum number of job requests to be processed.

8. In the Schedule section, set a schedule for the job request purge policy.

- To run the purge policy only one time, select **Once**. Click the calendar icon to select the date and time you want the purge policy to run.

- To set a schedule for the purge policy, select **Use Existing Schedule**. Click the browse button to search for a schedule.

  In the Select Schedule window, select a schedule and click **OK**.

9. Click **OK** to save the purge policy.

5.9.2 Physically Purging Job Requests from the Database

Job request data is kept in the Oracle Enterprise Scheduler runtime store as records in the runtime store tables. These job requests records are kept in the runtime store until they are physically purged by a database administrator using SQL purge scripts. You must logically delete a job request before physically purging it.

Use the method `esspurge.purge_requests` on the Oracle Enterprise Scheduler schema to delete purged job requests from the database. In an Oracle Fusion Applications environment, the schema is typically called `FUSION_ORA_ESS`.

The `ESSPURGE` package contains a stored procedure which you can use to purge completed job requests. The package is normally loaded when the Oracle Enterprise Scheduler schema is created or updated. The stored procedure is shown in Example 5–1.

Example 5–1   ESSPURGE Stored Procedure for Purging Completed Job Requests

```
--- Purges job requests that have completed.
---
--- p_older_than : Purge only job requests that have completed after this time.
--- p_max_count : The maximum number of job requests to purge.
--- p_max_runtime : The maximum time, in minutes, the purge should run.
---    If null (default), this defaults to one day which effectively means
---    there is no time limit.
---
procedure purge_requests
( p_older_than in timestamp,
  p_max_count in number,
  p_max_runtime integer default null );
```

The basic syntax of `esspurge.purge_requests` is shown in Example 5–2, where `FUSION_ORA_ESS` is the name of the Oracle Enterprise Scheduler schema and `password` is the password.

Example 5–2   Purging Job Requests from the Database

```
sqlplus FUSION_ORA_ESS/password
set serveroutput on size unlimited
set timing on
execute esspurge.purge_requests(systimestamp, 1000000);
```

Additional examples are shown in the following list.

- To purge job requests completed earlier than the current time, at a maximum of 50000 job requests, execute the following command:

  ```
  execute esspurge.purge_requests(systimestamp, 50000);
  ```

- To purge job requests completed 30 days earlier, at a maximum of 50000 job requests, execute the following command:

```
execute esspurge.purge_requests(systimestamp - 30, 50000);
```

- To purge job requests that completed before June 01, 2010 at 15:00:00 (UTC), at a maximum of 50000 job requests, execute the following command:

```
execute esspurge.purge_requests(TIMESTAMP '2010-06-01 15:00:00 -00:00', 50000);
```

5.9.3 Viewing Purge Policies

You can view all purge policies defined for the scheduling service.

To display purge policies for the scheduling service:

1. From the navigation pane, expand the **Scheduling Services** folder and select the Oracle Enterprise Scheduler application.

2. From the Scheduling Services menu, select **Purge Policies**.

5.9.4 Updating a Purge Policy

You can edit the request or retention criteria of an existing purge policy, as well as the purge policy schedule.

To update a purge policy:

1. From the navigation pane, expand the **Scheduling Services** folder and select the Oracle Enterprise Scheduler application.

2. From the Scheduling Services menu, select **Purge Policies**.

3. From the Purge Policies table, select the policy you want to update and click **Update**.

4. Edit the purge policy accordingly and click **OK**.

5.9.5 Deleting a Purge Policy

You can delete a given purge policy.

To delete a purge policy:

1. From the navigation pane, expand the **Scheduling Services** folder and select the Oracle Enterprise Scheduler application.

2. From the Scheduling Services menu, select **Purge Policies**.

3. From the Purge Policies table, select the policy you want to delete and click **Delete**.

5.10 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

5.10.1 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

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

5.10.1.1 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.

5.10.1.2 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 5.10.1.2.1, Section 5.10.1.2.2, Section 5.10.1.2.3 and Section 5.10.1.2.4.

7. Click **Save and Close**.

5.10.1.2.1 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 5.10.1.2.3 and Section 5.10.1.2.4.

5. Click **Save and Close** or **Save and Create Another**.

5.10.1.2.2 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**.

5.10.1.2.3 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**.

5.10.1.2.4 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 5.10.1.2.3.

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.

5.10.2 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

This section contains the following topics:

- Viewing Job Sets

- Creating or Editing a Job Set

- Deleting a Job Set

5.10.2.1 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.

5.10.2.2 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 5.10.1.2.1.

   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 5–9.

   In the **Initial Value** text field, enter the value you want to assign to the system property.

Table 5–9   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.
	For more information about business errors, see Section 5.7.2.2
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 oracle.as.scheduler.Executable 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.

Table 5–9   (Cont.) System Properties

System Property	Description
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 EFFECTIVE_APPLICATION system property. This property can only be specified via metadata and cannot be specified as a submission parameter.
SYS_environmentVariables	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.
	This property is optional for a Process job type. It is not used for other job types.
SYS_executePast	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.
	Valid values for this property are:
	■ 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.
	If this property is not specified, the system defaults to TRUE.
SYS_extensionListener	This property is for internal use only.
SYS_external_id	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.
SYS_groupName	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.
SYS_inputList	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.
	Oracle Enterprise Scheduler imposes no format on the value of this property.
SYS_inputWorkDirectory	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.
SYS_listener	This property specifies the event listener class associated with the request. This should be the name of a Java class that implements the oracle.as.scheduler.EventListener interface.
SYS_locale	This property specifies a language code to apply to the job.
SYS_logWorkDirectory	This property specifies the logging working directory.

Table 5–9 (Cont.) System Properties

System Property	Description
SYS_outputList	This property specifies output from a request.
	The output of a serial job set is the OUTPUT_LIST of the last step. The output of a parallel job set is the concatenation of the OUTPUT_LIST of all the steps, in no guaranteed order, with oracle.as.scheduler.SystemProperty.OUTPUT_LIST_DELIMITER as a separator.
SYS_outputWorkDirectory	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.
SYS_postProcess	This property specifies the post-process callout handler class. This should be the name of a Java class that implements the oracle.as.scheduler.PostProcessHandler interface.
SYS_preProcess	This property specifies the pre-process callout handler class. This should be the name of a Java class that implements the oracle.as.scheduler.PreProcessHandler interface.
SYS_priority	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.
SYS_procedureName	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.
SYS_product	This property specifies the product within the application that submitted the request.
SYS_redirectedOutputFile	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.
SYS_reprocessDelay	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.
SYS_request_timeout	This property enables the job request to time out.
SYS_requestCategory	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.
SYS_requestedProcessor	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 oracle.as.scheduler.WorkAssignment (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.

Table 5–9   (Cont.) System Properties

System Property	Description
SYS_requestExpiration	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.
SYS_retries	This property specifies the retry limit for a failed request. If request execution fails, the request will 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.
SYS_runasApplicationID	This property enables elevating access privileges for completing a scheduled job.
	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.
SYS_selectState	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 SELECT_STATE must be set to true. If SELECT_STATE 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.
SYS_sqlJobClass	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.
	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.
	This property is optional for a SQL job type. It is not used for other job types.
SYS_submittingApplication	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.
SYS_substitutionHandlers	This property has been deprecated.
SYS_successExitCode	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.
	This property is optional for a process job type. It is not used for other job types.

Table 5–9   (Cont.) System Properties

System Property	Description
SYS_userFileDir	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.
	Oracle Enterprise Scheduler supports a configuration parameter that specifies a file directory where requests may store files. At request submission, a USER_FILE_DIR property is automatically added for the request if the configuration parameter is currently set and USER_FILE_DIR 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.
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.
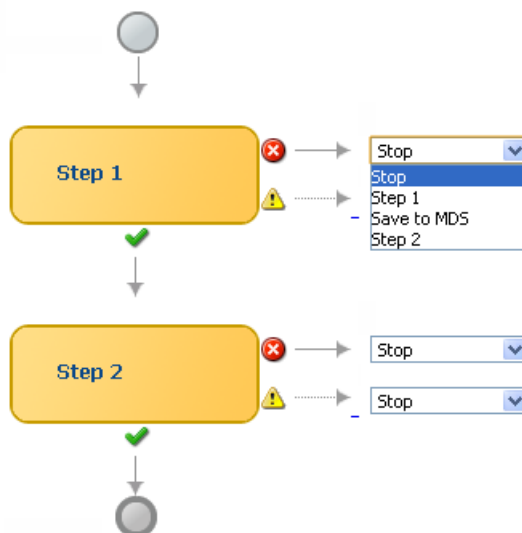
7. Continue to define steps as required for the job set.

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.

5.10.2.3 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 .

2. In the Results table, select the job set you want to delete and click the **Delete** icon.

5.10.3 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

5.10.3.1 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.
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.

5. To view the details of a specific incompatibility, click the name of the relevant incompatibility.

5.10.3.2 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:

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. 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 5.10.2.2.

8. Configure access control for the incompatibility, as described in step 9 of Section 5.10.2.2.

9. Click **OK** to save the incompatibility.

5.10.3.3 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 5.10.3.1.

2. In the Results table, select the incompatibility you want to delete and click the **Delete** icon.

5.11 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.

This section contains the following topics:

- Managing Work Assignments
- Managing Workshifts
- Managing Schedules

5.11.1 Managing Work Assignments

Work assignments provide the following controls for processing job requests:

- At runtime, work assignments can be attached to a request processor and can limit the period during which job requests of a certain type can be processed.

- At runtime, a request processor can be controlled and a work assignment can limit the resources that are available to process all job requests.

Determining Active Work Assignments

A bound work assignment is active if it is enabled, has an active workshift and the enabled flag is set on the work assignment definition. A workshift is active if it has an allocation greater than zero and includes a current schedule (where the current time is within a time window defined by the schedule and duration), or the workshift is a 24x7 workshift. If there are no bound work assignments on the server, the default work assignment will be active using all threads and having no limits on asynchronous workers.

Only one workshift can be active for a given work assignment at any point in time. If a work assignment has more than one active workshift, Oracle Enterprise Scheduler chooses the most specific of these to be the actual active workshift. The most specific workshift is the one that ends soonest, or, given two workshifts that have ended at the same time, the workshift that started most recently.

Determining Work Assignment Thread Allocation

An active work assignment is assigned the thread allocation specified by the active workshift unless the total number of threads needed to satisfy the allocations of all active work assignments exceeds the configured thread count. In this case, Oracle Enterprise Scheduler weighs the thread allocation against the percentage of threads allotted to the workshift out of the total number of thread allocations for all work assignments.

For example, suppose work assignment 1 has a thread allocation of 70, work assignment 2 has a thread allocation of 30, and there are 20 processor threads configured. The total desired allocation is 100, so the weight for work assignment 1 is 70% while the weight for work assignment 2 is 30%. Oracle Enterprise Scheduler allocates 14 threads to work assignment 1 and six threads to work assignment 2.

If the default work assignment is active, the number of threads allocated to the work assignment is equal to the configured thread count.

> **Note:** Each active work assignment is assigned at least one thread.

Processing Active Work Assignments

After determining work assignments and thread allocations, Oracle Enterprise Scheduler begins a thread pool for each active work assignment. The threads are responsible for processing job requests that are specialized to the work assignment, except for requests that are specialized to an exclusive work assignment. The exclusion is effective for all work assignments, including the default work assignment. If an exclusive work assignment is bound to any server in the group, no other work assignment can process job requests specialized to that exclusive work assignment.

> **Note:** All work assignments bound in exclusive mode are excluded, including disabled work assignments. Exclusive bindings apply even if the server to which they are bound is unavailable. You must unbind an exclusive work assignment if you do not want it to be excluded.

This section contains the following topics:

- Creating or Editing a Work Assignment
- Deleting a Work Assignment

5.11.1.1 Creating or Editing a Work Assignment

A work assignment includes two primary components that define request processor constraints:

- Specialization rules: Define restrictions for job request processing on a request processor.
- Workshifts: Specify temporal windows for processing job requests and thus describe the schedule for when job requests can be processed on a request processor

This combination of work assignment controls, including specialization rules and workshifts gives you the ability to select the kinds of job requests to be processed and decide how to allocate the request processor resources. For example, you can define two workshifts: a day shift and a night shift to allocate processing for these periods. The day shift could have more resources allocated for a peak usage period while the night shift could provide a different mix for its resource allocation.

By default, no work assignments are bound to a request processor and the request processor processes any ready job request. The default behavior is similar to using a request processor with no specialization rules and a workshift of 24/7 duration, all configured threads are used and there are no limits on the number of asynchronous jobs.

Table 5–10 shows the properties that you can define for specialization rules.

Table 5–10   Specialization Properties Available for Specialization Rules

Specialization Property	Description
Application	Specifies the name of the application associated with a job request.
Product	Specifies the name of the product within an application.

Table 5–10   (Cont.) Specialization Properties Available for Specialization Rules

Specialization Property	Description
Submitted By	Specifies a user who submitted a job request.
Job Definition	Specifies a specific job request name.
Request Category	Specifies a label defined by the system administrator allowing administrators to group job requests for specific needs.

The following operators can be used to join the conditions of a rule: AND, OR (both of type binary), CONTAINS and NOT (unary).

Example 5–3 shows sample specialization rules that can be used in a work assignment.

Example 5–3   Sample Specialization Rules

```
application = 'EssDemoApp' AND (definition = 'JobDefinition://mypackage/Job_
essdemo1' OR definition = 'JobDefinition://mypackage/LongRunningJob')

requestCategory ='Priority'

user = 'sam'
(requestCategory ='LongRunning') AND NOT (definition =
'JobDefinition://mypackage/LongRunningJob')
```

If a job request is specialized to two different work assignments, either of those work assignments may process the job request depending on resource availability. Similarly, if the same work assignment is bound to two different servers, either server may process the job request.

To create or edit a work assignment:

1. From the navigation pane, expand the **Scheduling Services** folder and select the Oracle Enterprise Scheduler application.

2. From the Scheduling Services menu, select **Work Allocation** and then select Work Assignments.

    The Work Assignments page displays.

3. Click the **Create** or **Edit** buttons to create or edit a work assignment.

    The Create Work Assignment page displays.

4. Enter a name and description for the work assignment in the relevant text fields.

5. Select the **Enabled** checkbox to enable the work assignment.

6. Click **Create Specialization** to add conditions under which the work assignment runs.

7. In the Workshifts section, click the **Add** icon to add one or more workshifts to the work assignment.

8. Click **OK** to save the work assignment.

5.11.1.2 Deleting a Work Assignment

You can delete a work assignment from the list of work assignments. Before deleting a work assignment, make sure that it is not bound to a request processor.

To delete a work assignment:

1. From the navigation pane, expand the **Scheduling Services** folder and select the Oracle Enterprise Scheduler application.

2. From the Scheduling Services menu, select **Work Allocation** and then select Work Assignments.

    The Work Assignments page displays.

3. Select the work assignment you want to delete and click **Delete**.

5.11.2 Managing Workshifts

A workshift indicates the operating, active times for a request processor. Specifically, a workshift defines a sequence of temporal windows in which resources or threads, are made available for processing job requests. When a work assignment is bound to a request processor, one or more workshifts are associated with the work assignment. At runtime, Oracle Enterprise Scheduler determines the resource allocation for the workshifts within the work assignment.

A workshift defines the following resources:

- Thread allocation,

- Asynchronous job limits for asynchronous Java and PL/SQL jobs.

Thread allocation specifies the number of threads that can be used by the request processor. These threads are used to perform local tasks, including processing synchronous jobs, initiating and finalizing asynchronous jobs, pre- and post-processing of job requests and updating events. When the workshift in a work assignment is active, each request processor for the work assignment can use the specified number of threads. For example, suppose a work assignment includes a 24x7 workshift with a thread allocation of 15. If that work assignment were bound to three request processors, each request processor could use 15 threads, for a total of 45 threads across all three servers.

Asynchronous, PL/SQL and asynchronous Java jobs use a global resource that must be shared across the entire system. The workshift can specify limits on the number of PL/SQL jobs and asynchronous Java jobs that can be active for the work assignment. This limit applies across all request processors for the work assignment in the entire system. For example, suppose a work assignment includes a 24x7 workshift with a PL/SQL job limit of ten. If that work assignment were bound to three request processors, all three request processors would share the ten PL/SQL asynchronous workers, for a maximum of ten PL/SQL jobs active for that work assignment.

When deciding on thread allocation and asynchronous job limits for a workshift, take note of the types of jobs specialized to the work assignment where the workshift is to be used.

This section contains the following topics:

- Creating or Editing a Workshift

- Deleting a Workshift

5.11.2.1 Creating or Editing a Workshift

Creating a workshift involves the following:

- Schedule: Associate a schedule with the workshift.

- Duration: Enter a duration for the workshift.

- Thread allocation: Specify the number of threads you want to allocate to the workshift.

To create or edit a workshift:

1. From the navigation pane, expand the **Scheduling Services** folder and select the Oracle Enterprise Scheduler application.

2. From the Scheduling Services menu, select **Work Allocation** and then select Workshifts.

   The Workshifts page displays.

3. Click the **Create** or **Edit** buttons to create or edit a workshift.

   The Create Workshift page displays.

4. Enter the following information for the workshift.

   - **Name**: Enter a name for the workshift.

   - **Description**: Enter a description for the workshift.

5. In the Active Period section, select one of the following activation periods for the workshift.

   - **Active 24x7**: Select to always enable the workshift. Selecting this option removes the Duration text field.

   - **Use existing schedule**: Select to enable the workshift using a schedule you created previously. Click the **Browse** button to display the Select Schedule window and search for the schedule using the Name and Package fields.

     In the Duration text field, enter the duration of the workshift in minutes.

   - **Specify schedule**: Select to create a schedule for the workshift. Enter a name and description for the schedule in the text fields provided.

     From the Time Zone dropdown list, select a time zone for the schedule.

     In the Start region, select **Immediately** or **Later**. If you choose Later, select the date and time using the calendar icon.

     In the Repeating region, use the Repeat dropdown list to select the frequency at which the workshift schedule will repeat: **Do not repeat**, **Every N minutes/hours/days/weeks**, **Specific days of the week**, **Specific days of the month**. Enter the number of minutes, hours, days or weeks. Alternatively, select the days of the week or dates of the month on which the schedule is to run.

     In the End Date region, select **No End Date** or **Specified End Date**. When selecting Specified End Date, use the calendar icon to set an end date in the **Date** text field.

     In the Duration text field, enter the duration of the workshift in minutes.

6. Expand the Advanced region. Specify the thread allocation and the asynchronous job limits.

   a. In the Thread Allocation field, enter the number of threads to be allocated to the processor when the workshift is active. If the total thread allocation of the active workshifts exceeds the number of available threads, then the threads are allocated proportionately to this workshift.

    b. In the Asynchronous Job Limits region, enter the number of asynchronous Java and PL/SQL jobs that can be active for the work assignment. This limit applies across all processors to which the work assignment is bound.

7. Click **OK** to save the workshift.

5.11.2.2 Deleting a Workshift

You can delete a workshift from the list of workshifts.

To delete a workshift:

1. From the navigation pane, expand the **Scheduling Services** folder and select the Oracle Enterprise Scheduler application.

2. From the Scheduling Services menu, select **Work Allocation** and then select Workshifts.

    The Workshifts page displays.

3. Select the workshift you want to delete and click **Delete**.

5.11.3 Managing Schedules

Schedules are used to configure the start and end times as well as the frequency of work allocations, purge policies and job requests.

This section contains the following topics:

- Creating or Editing a Schedule

- Deleting a Schedule

5.11.3.1 Creating or Editing a Schedule

A job schedule specifies a predefined time or recurrence for scheduling and executing a job request, work allocation or purge policy. Schedules are defined independent of jobs and can be associated with one or more jobs.

To create or edit a schedule:

1. From the navigation pane, expand the **Scheduling Services** folder and select the Oracle Enterprise Scheduler application.

2. From the Scheduling Services menu, select **Work Allocation** and then select Schedules.

    The Schedules page displays.

3. Click **Create** to define a new schedule, or **Edit** to modify an existing schedule.

    The Create Schedules page displays.

4. Enter the following information:

- **Name**: Enter a name for the schedule.

- **Description**: Optionally, add descriptive text for the schedule.

- **Time Zone**: From the dropdown list, select the time zone in which the schedule will run.

5. Configure the beginning of the schedule. In the Start field, select **Immediately** or **Later**.

    If you select Later, use the date picker to select a date and time for the beginning of the schedule.

6. In the Repeat field, configure the frequency of the schedule.

   a. From the **Repeat** dropdown list, select the unit of time by which you want to repeat the schedule. Options include: Do Not Repeat, By the Minute, Hourly, Daily, Weekly or Weeks.

   b. In the Frequency text field, enter the number of units for the schedule frequency. Enter the number of minutes, hours, weeks or times per day.

      If selecting Weekly, use the Day of the Week dropdown list to select the day or days of the week you want to run the schedule.

   c. In the End date field, select an end date for the schedule. Select **No End Date** or **Specified End Date**. If selecting **Specified End Date**, use the date picker to select the end time and date.

7. Click **OK** to save the schedule.

5.11.3.2 Deleting a Schedule

You can delete a schedule from the list of schedules.

To delete a schedule:

1. From the navigation pane, expand the **Scheduling Services** folder and select the Oracle Enterprise Scheduler application.

2. From the Scheduling Services menu, select **Work Allocation** and then select Schedules.

   The Schedules page displays.

3. Select the schedule you want to delete and click the **Delete** button.

5.12 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 Service, 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:

- **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.

- **RequestFileDirectory**: Configure this property to indicate the path of the folder to which job request logs are to be written. This property is used for job requests submitted by Oracle ADF applications.

- **SAMLTokenPolicyURI**: Set this property to the URI of the SAML policy used to secure job requests submitted by Oracle ADF applications.

- **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.

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 **ESS 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.

5.13 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.

5.14 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*.

5.15 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.

This section contains the following topics:

- Finding Log Information for a Job Request
- Viewing Job Request Logs
- Viewing Log Messages for Scheduled Job Requests
- Setting Oracle Enterprise Scheduler Log Levels
- Tracing Oracle Enterprise Scheduler Jobs
- 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.

5.15.1 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 17.5.3

5.15.2 Viewing Job Request Logs

You can view the log for a particular job request from the Request Details pane.

To view job request logs:

1. Search for the relevant job requests as described in Section 5.7.2.1

2. In the table displaying the job request search results, select the job request whose log you want to view.

3. To view job request details, click the job request ID. Alternatively, click the parent ID associated with the job request to view the details of the job set with which the job is associated.

4. In the job request details page, select **Action** and then select Request Log to display log information for the job request.

5.15.3 Viewing Log Messages for Scheduled Job Requests

You can use the Oracle Enterprise Scheduler Log Messages page to view logging information regarding scheduled job requests.

For more information about logging, see Chapter 13.

To view log messages for scheduled job requests with Fusion Applications Control:

1. From the navigation pane, expand the farm, and then **Scheduling Services**.

2. Select the **ESSAPP** application for the appropriate Managed Server.

3. In the Scheduling Service home page, from the **Scheduling Service** menu, choose **Logs** and then select **View Log Messages**.

4. Search for the relevant log messages using the Date Range, Message Types and Message fields. You can optionally add additional search fields. For more information about logging, see Chapter 13.

5. You can use the ECID to retrieve more information about a given job request. The Execution Context ID (ECID) is a global unique identifier of the execution of a particular request in which the originating component participates. You can use the ECID to correlate error messages from different components

6. By default, when view the logs for a request, Fusion Applications Control displays only messages logged in the scope of the Oracle Enterprise Scheduler cluster. If the Oracle Enterprise Scheduler application is not deployed to a cluster, Fusion Applications Control displays the messages logged in the managed server scope. However, Oracle Enterprise Scheduler propagates the ECID associated with the job request across sub-systems such as SOA, ADF, and so on.

   To view messages logged by other sub-systems, broaden the target scope to display messages logged across the domain or farm. Click **Broaden Target Scope** and select the target whose log messages you want to view.

5.15.4 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 5–11 shows the Oracle Enterprise Scheduler logger names, log levels and a description for each level.

Table 5–11   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.
	CONFIG	Application endpoint activation and de-activation for Oracle Enterprise Scheduler resource adapter.
`oracle.apps.fnd.applcp.srs`	WARNING	Any problems that occur during UI rendering to submission.
	INFO	Messages related to job fetch and submission API calls.
	FINEST	Details tracing messages for the scheduled job request submission UI.
`oracle.apps.fnd.applcp.monitor`	WARNING	Any problems that occur during UI rendering to operations in the UI.
	INFO	Messages related to job request fetch and various API calls.
	FINEST	Details tracing messages for the job request monitoring UI.
`oracle.apps.fnd.applcp`	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 17.2. 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 Section 5.15.2.

For information about setting the log levels for Oracle WebLogic Server, see Section 5.15.5. 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.

5.15.5 Tracing Oracle Enterprise Scheduler Jobs

Enabling tracing for Oracle Enterprise Scheduler jobs provides additional information regarding job requests, which can then be relayed to Oracle technical support to help solve any Oracle Enterprise Scheduler-related issues.

This section contains the following topics:

- Navigate to the Specify Loggers Region
- Define the Loggers
- Attach the Loggers
- Download the Trace File

Before You Begin

Before enabling tracing, note that tracing log messages may produce a considerable amount of output. By default, log message traces go to the Oracle WebLogic Server console stdout. If the stdout is re-directed to a file, it may occupy a significant amount of disk space and possibly affect server performance.

For more information about Oracle Diagnostic Logging levels, see Section 13.7.3.

To enable tracing for Oracle Enterprise Scheduler jobs:

Task 1 Navigate to the Specify Loggers Region

Navigate to the Log Configuration pane and expand the Specify Loggers region.

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 Levels tab, from the View dropdown list select **Loggers with Persistent Log Level State**. Expand the Specify Loggers region toward the bottom of the pane.

Task 2 Define the Loggers

Create the following loggers:
`oracle.as.scheduler.security.internal.SecurityExecutorFactory`
and `oracle.as.scheduler.security.internal.SecurityHelper`.

1. In the Name text field, enter the name of the logger.

2. From the **Oracle Diagnostic Logging Level (Java Level)** dropdown list, select the desired logging level.

3. Repeat for the second logger name.

Task 3 Attach the Loggers

Attach the loggers to the Oracle Diagnostic Logging Handler.

1. In the Log Configuration pane, click the **Log Files** tab.

2. From the list of handlers, select **odl-handler** and click **Edit Configuration**.

3. From the Loggers to Associate dropdown list, select the loggers you created and click **OK**.

Task 4 Download the Trace File

Download the trace file.

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 View Log Messages.

    The Log Messages pane displays.

3. In the Search field, expand the Selected Targets region and select the relevant Oracle Enterprise Scheduler target.

4. Select the relevant date range, message type and so on, and click **Search**.

5. From the list of log messages that displays, select the relevant message.

6. In the Log File column, click the name of the log file.

    The View Log File pane displays.

7. Next to the name of the log file, click **Download** to download the file.

5.15.6 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 Universal Content Management. 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:

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`.

6

Managing Report Delivery Servers

This chapter provides information you need to configure and manage report delivery servers.

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 the Oracle Enterprise Scheduler Service
- Through the Reports and Analytics pane
- From an application page

Requests submitted through Oracle Enterprise Scheduler Service 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 the Oracle Enterprise Scheduler Service 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
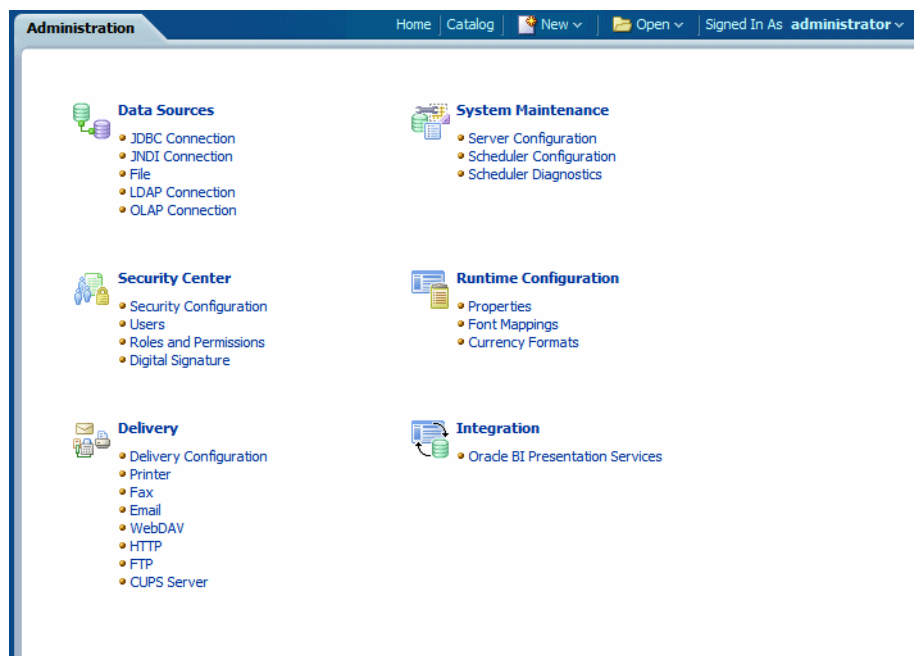
6.3 Configuring Report Delivery Servers

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 Service 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

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 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 Fusion Middleware 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 Scheduler (Oracle Enterprise Scheduler Service).

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

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 Oracle Fusion Middleware Control

For information on using Oracle Fusion Middleware 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)*.

7

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 tags, which supports crawling searchable objects that contain Oracle WebCenter 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
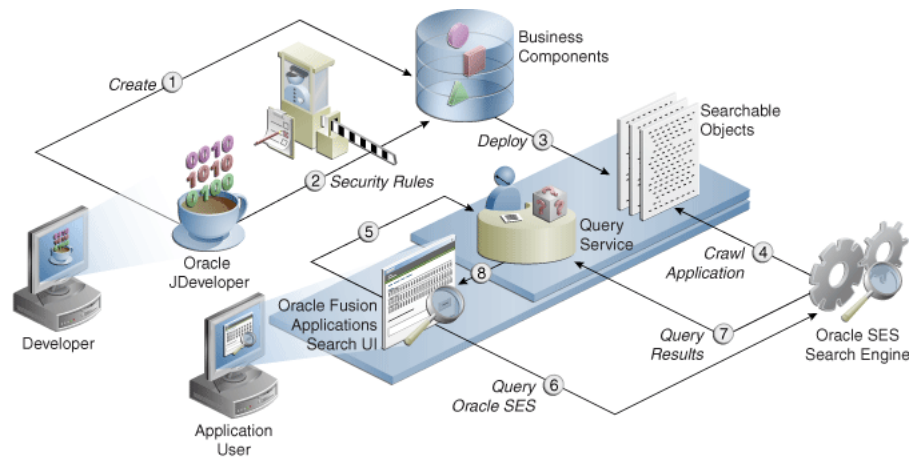
The process flow includes the following steps:

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 (`ecsf.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.
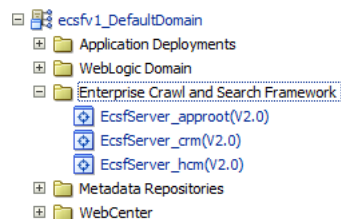
The Enterprise Crawl and Search Framework Configuration Settings page displays. For information about the Configuration Settings page, see Section 7.3.2.

2. Expand the **Enterprise Crawl and Search Framework** target menu.
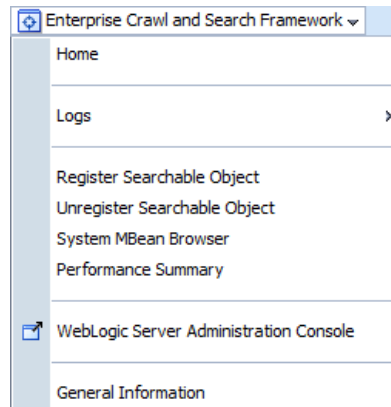
The menu displays the options described in Table 7–1.

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 *Oracle Fusion Middleware Administrator's Guide*.
	The log files are stored in the following directory:
	DOMAIN_HOME/sysman/log/emoms.log
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

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

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

5. 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:

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 activate, and then select the **Searchable Objects** tab, shown in 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 **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

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 .

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

9. 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:

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 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.

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 **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
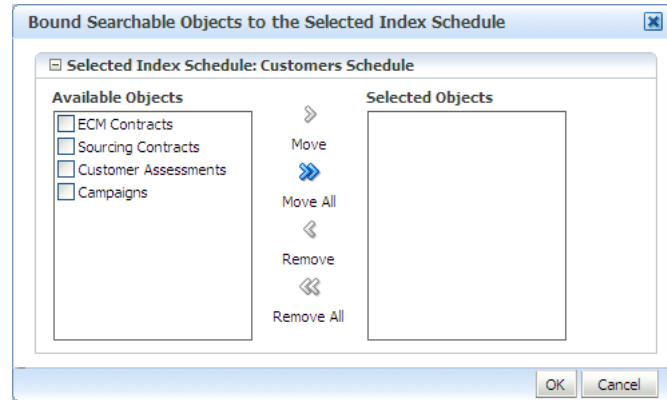
7. Click **OK**.

   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:

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 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 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

Administering Search

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

5. Click the **Import** button.

The Import External Categories dialog, shown in Figure 7–9, appears.

Figure 7–9   Import External Categories Dialog
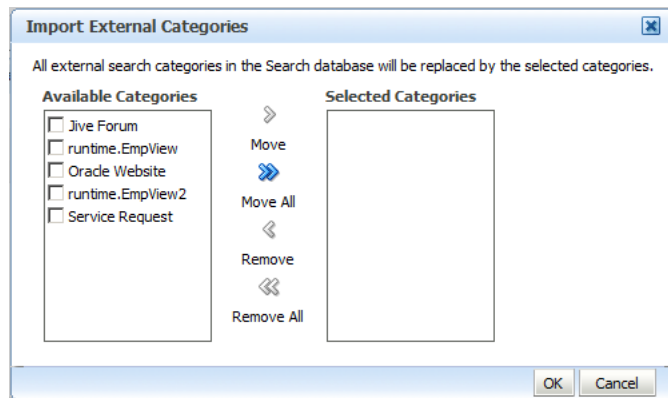
6. In the Available Categories column, select the checkbox of the external search categories you want to import and click the **Move** icon to shuttle your selection to the Selected Categories column.

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. If you had previously deleted any of the records corresponding to the external search categories, you must delete them again to make them unavailable for querying.
>
> 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
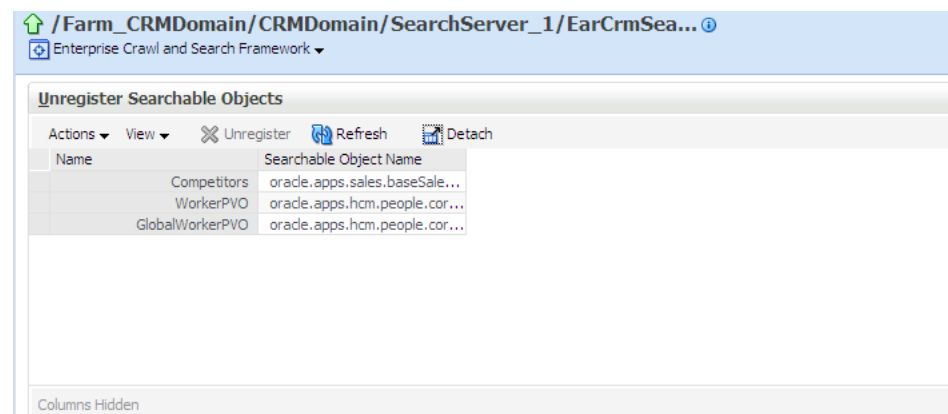
5. Click the **Unregister** icon.

   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.
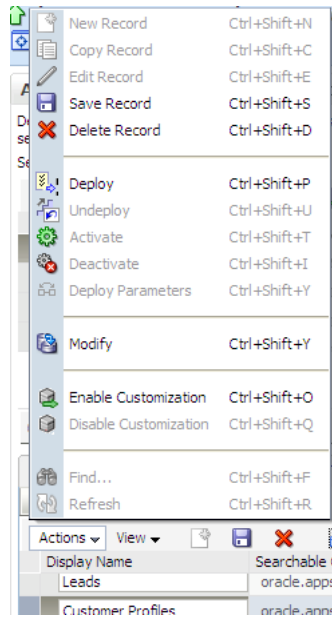
This validates all the modified expressions for title, body, keywords, and action title and sets the customization flag. A green check mark appears in the Customized column to indicate that the searchable object is customized.

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

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
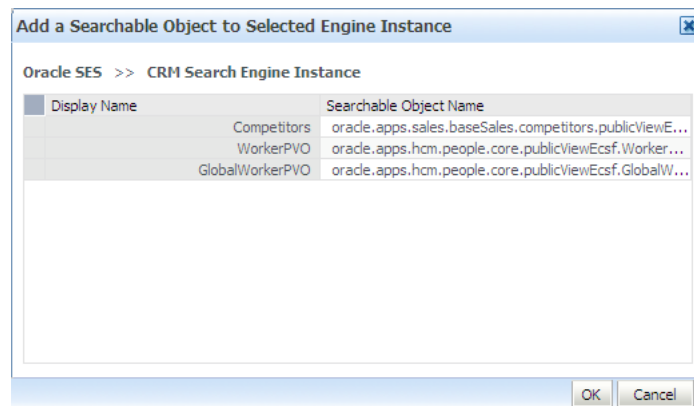
6. Select the desired searchable object from the table, and click **OK**.

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

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
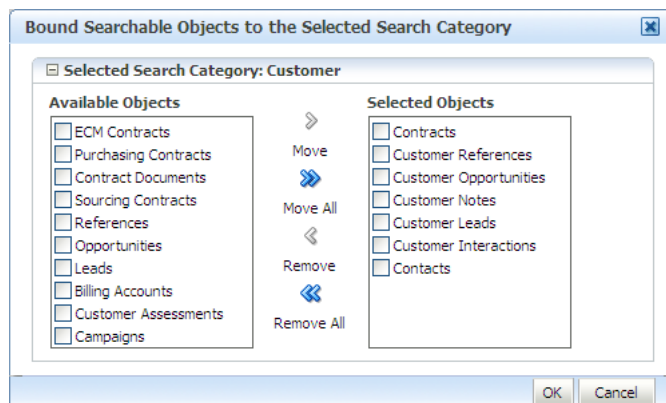
7. Click **OK**.

    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/approot/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/approot/searchfeedservlet`
ECSF_REDIRECT_SERVICE	URL to ECSF feed servlets. This URL is used to redirect drilldowns/actionable results.	`http://example.com:7101/approot/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
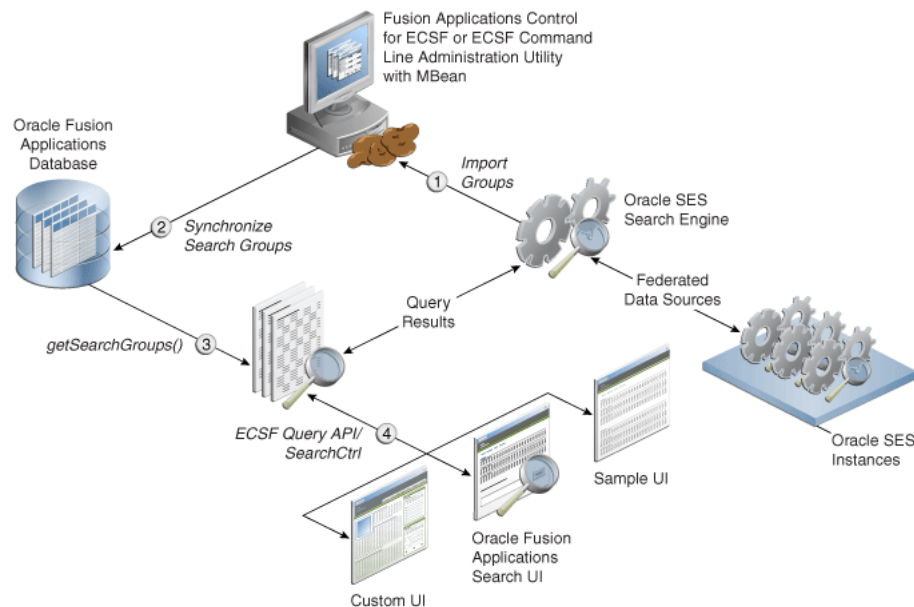
The process flow includes the following steps:

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/ecsf/service/query/common/componentname/}SearchServi
   ce
   ```

   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.

8

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
ChargingAllowed	String	Specifies if charging is allowed.	Default: False **Note**: No server restart is required if you change the value.
ControllerUri	String in URI format	Specifies the Third Party controller (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.
MaximumCallDuration	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 no P-Asserted-Identity header is 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"

9

Monitoring Oracle Fusion Applications

This chapter describes how to manage Oracle Fusion web applications using Oracle Enterprise Manager Fusion Applications Control (Fusion Applications Control).

This chapter includes the following topics:

- Introduction to Monitoring Oracle Fusion Applications
- The Oracle Fusion Applications Product Family Home Page and Sections
- The Oracle Fusion Applications Product Home Page and Sections
- 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

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 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 Oracle Fusion Applications Product Family Home Page and Sections

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 Oracle Product Family Home Page
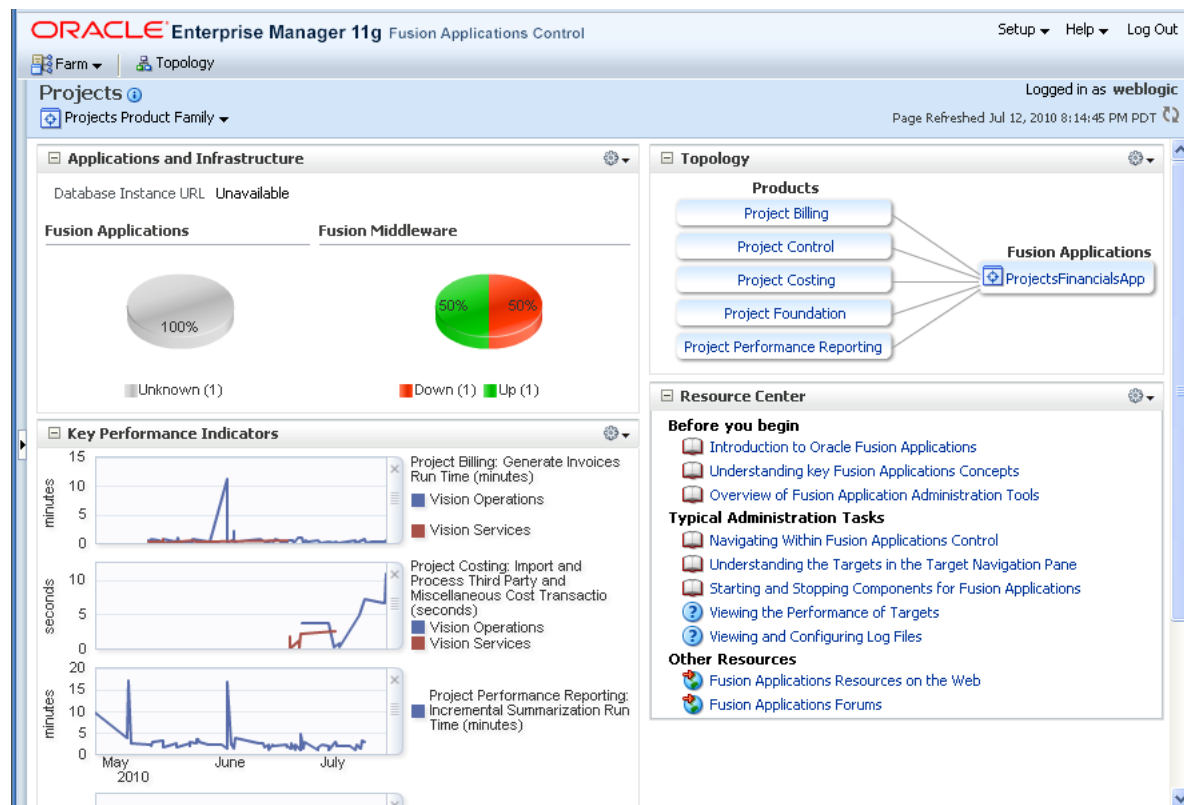
The Product Family home page displays the following sections:

- 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 9.7

9.2.3 Topology Section

The Topology section displays the components of the product family or product. For example, the Topology section in the Product Family home page displays the product family or families under the Oracle Fusion Applications heading, and the related products under the Products heading. In the Product home page, the Topology section displays the same information, with the irrelevant products and Oracle Fusion 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
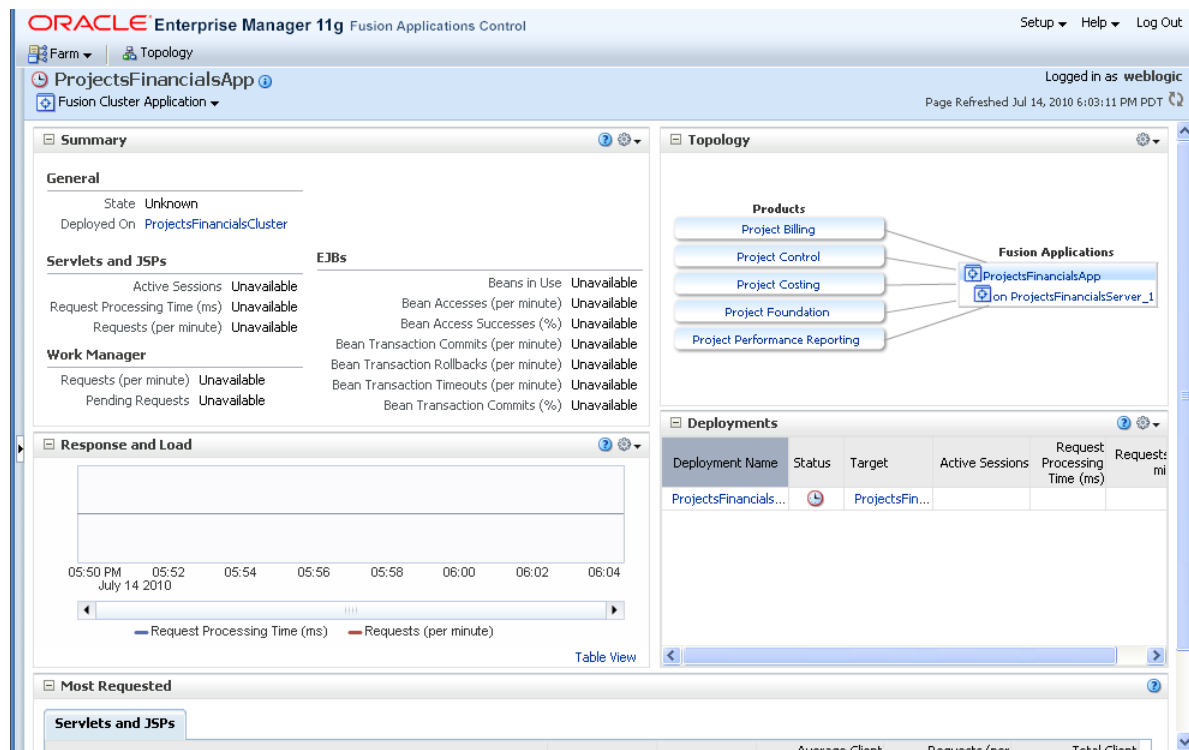
For more information about viewing topology information, see Section 10.2.3

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.

9.3 The Oracle Fusion Applications Product Home Page and Sections

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   Oracle Fusion Applications Product Home Page
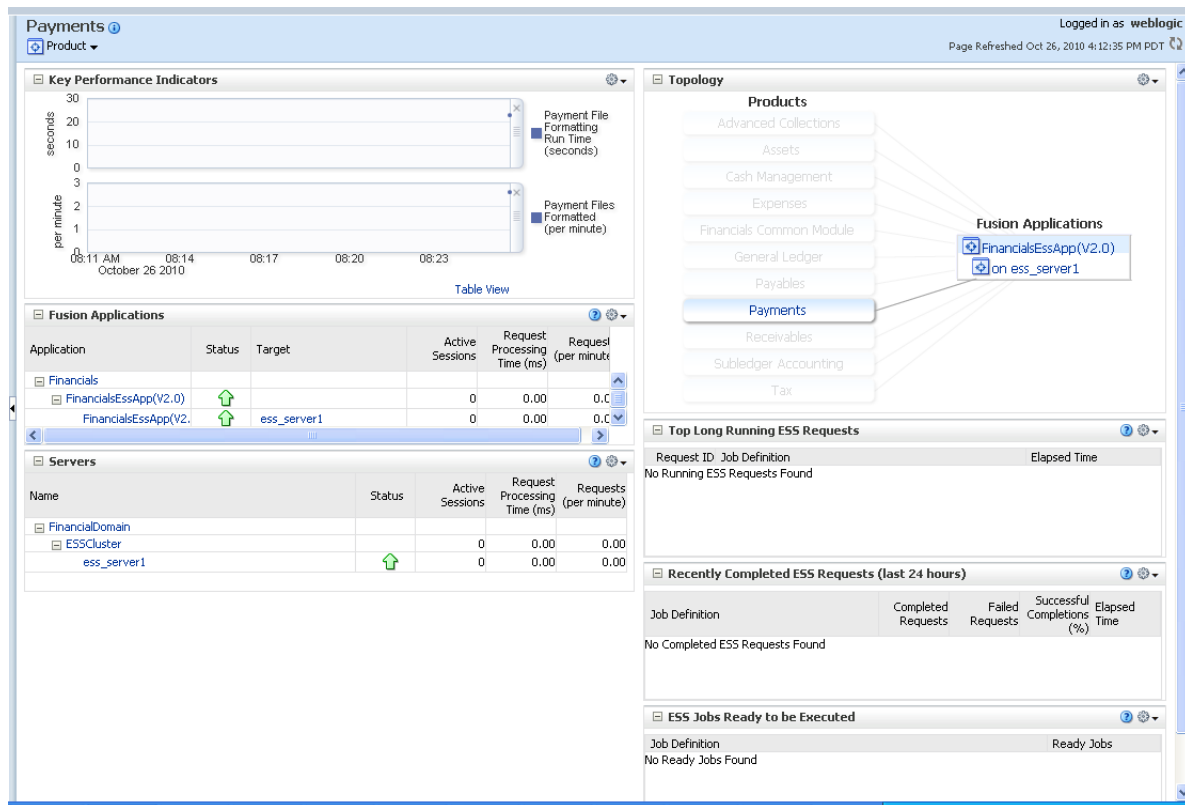
The Oracle Fusion Applications Product Home page displays the following sections:

- 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 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:

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. 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 9.5.

3. 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.
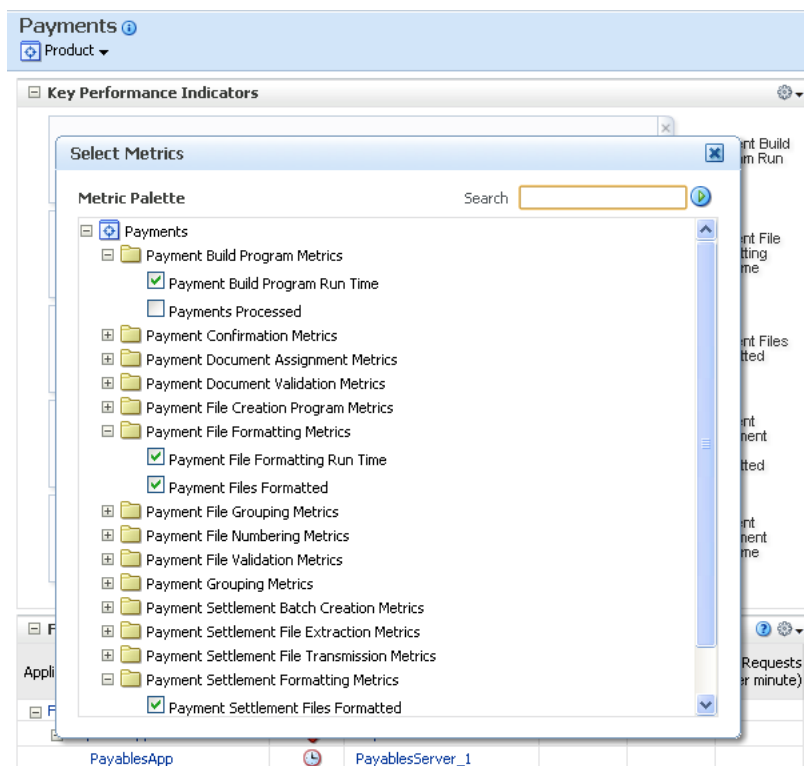
In the Select Metrics dialog, you can search for a particular metric in the Search field and select or remove the metrics you want to display or hide in the **Key Performance Indicators** section.

For more information about monitoring and modifying performance metrics, see Section 9.6.

- **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**.

4. 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.

9.5 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 in Fusion Applications Control:

1. From the navigation pane, select the product family.

2. From the **Product Family** menu, 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.
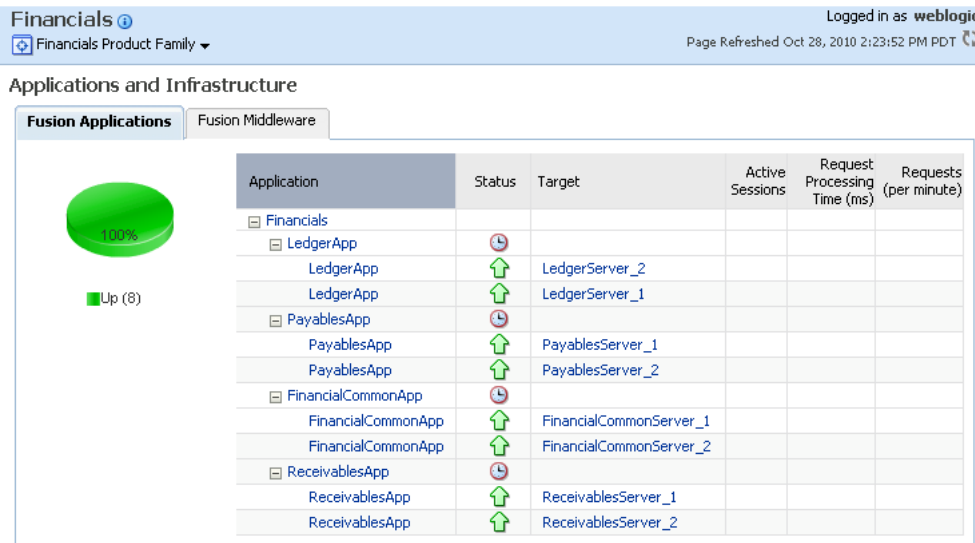
9.6 Monitoring Business Performance Metrics for a Product Family

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:

1. From the **Product Family** menu, 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** pane.

2. Display the metrics pane by clicking **Show Metric Palette**.

3. 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.

4. 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.

5. 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.

For information about viewing and modifying the display of product family performance indicators, see Section 2.5.7.

9.7 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:

1. From the navigation pane, select the product family name, and then expand the **Products** folder.

2. In the **Products** folder, click the product name.

   The product home page displays. You can view the key metrics for the product in the **Key Performance Indicators** pane.

3. To modify the display of metrics in the Key Performance Indicators pane, follow the procedure in Section 9.6, beginning with Step 2.

10

Monitoring the Oracle Fusion Applications Middle Tier

This chapter describes how to monitor the middle tier.

This chapter includes the following topics:

- Introduction to Monitoring the Middle Tier
- Monitoring and Viewing Performance Status of Oracle Fusion Middleware with Fusion Applications Control
- Monitoring System Performance Metrics for a Product Family or a Product with Fusion Applications Control
- Monitoring Performance for Key Oracle Fusion Middleware Components with Fusion Applications Control
- Performing Proactive Performance Monitoring with Grid Control
- Tuning Middle Tiers for Oracle Fusion Applications

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

10.1 Introduction to Monitoring the Middle Tier

Oracle Fusion Applications 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.

You can use Oracle Enterprise Manager Fusion Applications Control(Fusion Applications Control) to monitor the real-time data. 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 use Fusion Applications Control to monitor a specific set of metrics for performance tuning information.

If you are interested in viewing historical data, use Oracle Enterprise Manager Grid Control (Grid Control). See Section 10.5 for more specific information about usage of Grid Control.

10.2 Monitoring and Viewing Performance Status of Oracle Fusion Middleware with Fusion Applications Control

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:

- Verifying the General Status of Oracle Fusion Middleware Components

- Viewing the Detailed Status of Oracle Fusion Middleware Components

- Viewing the Routing Topology

- Viewing the Performance of an Oracle WebLogic Server

- Monitoring the System Performance of Oracle Fusion Applications Deployed to a Cluster

10.2.1 Verifying the General Status of Oracle Fusion Middleware Components

You can view the overall status of the underlying Oracle Fusion Middleware components from the home page of the product family using Fusion Applications Control.

To view the general status with Fusion Applications Control:

1. From the navigation pane, select the product family.

2. View the **Applications and Infrastructure** section, as shown in Figure 10–1.

Figure 10–1   Applications and Infrastructure Section on Product Family Home Page
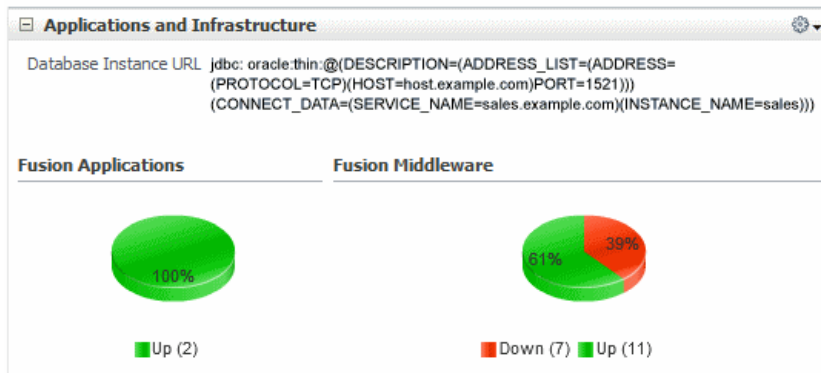
The Fusion Middleware pie chars display the percentage of Oracle Fusion Middleware components that are currently up or down.

10.2.2 Viewing the Detailed Status of Oracle Fusion Middleware Components

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, from Fusion Applications Control.

To view detailed status with Fusion Applications Control:

1. From the navigation pane, select the product family.

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.

Figure 10–2   Fusion Middleware Tab on Applications and Infrastructure Tab
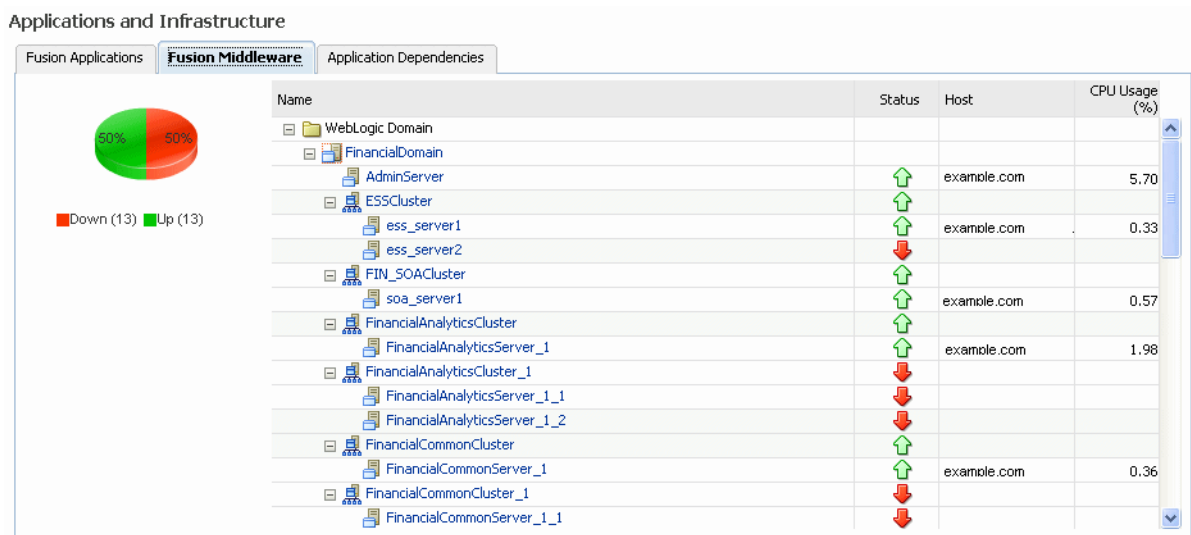
4. Select a specific component to gain more insight.

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

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.
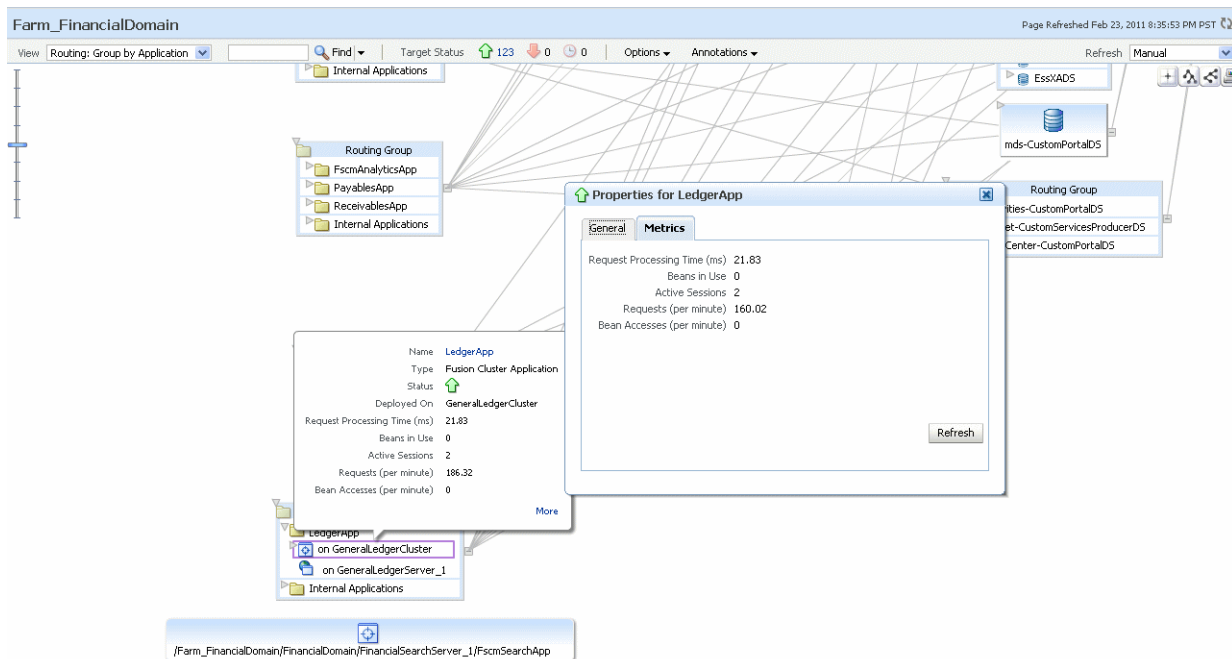
4. From the **Options** menu, you can save or print the image, expand or collapse all of the nodes, change the orientation of the topology to be left to right or top to bottom, or show and hide the routing navigator.

   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 workaround 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.

10.2.4 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. 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

- Monitoring the Performance of Oracle Fusion Application Instances

10.2.5.1 Monitoring the Performance of Oracle Fusion Applications Cluster Applications

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:
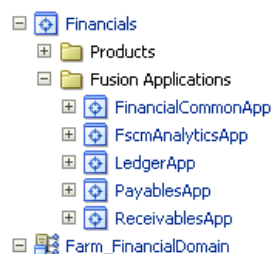
2. Monitor the overall performance of the applications on a cluster by clicking the cluster application name, for example, **PayablesApp**.

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 Application Instances

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 Cluster Application** menu, choose **Performance Summary**.

The Performance Summary page displays. For more information about using this page, see Section 2.5.7.

10.3 Monitoring System Performance Metrics for a Product Family or a Product with Fusion Applications Control

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 the product family or metrics specific to a particular product, to find out more information about the problem.

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, as shown in the following figure. 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.

10.4 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	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 *Oracle Fusion Middleware Administrator's Guide* for instructions on accessing the server home page. Specifically, review the Summary section of the page.
	Performance metrics: "Monitoring Service Levels" section in the *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*
Oracle Enterprise Content Management Suite, Content Server	Status of an Oracle WebLogic Server Managed Server for Content Server: "Monitoring an Oracle WebLogic Server Administration or Managed Server" section in the *Oracle Fusion Middleware Administrator's Guide* for instructions on accessing the server home page. Specifically, review the Summary section of the page.
	Performance Metrics: "Viewing Performance Information for Oracle Content Server" section in the *Oracle Fusion Middleware System Administrator's Guide for Oracle Content Server*
Oracle Enterprise Crawl and Search Framework	Performance metrics: Section 10.4.1
Oracle Enterprise Scheduler Service	Status of an Oracle WebLogic Server Managed Server for Oracle Enterprise Scheduler: "Monitoring an Oracle WebLogic Server Administration or Managed Server" section in the *Oracle Fusion Middleware Administrator's Guide* for instructions on accessing the server home page. Specifically, review the Summary section of the page.
	Performance metrics: Section 10.4.2.4 and Section 10.4.4.3

Table 10–1   (Cont.) Monitoring Oracle Fusion Middleware Components

Component	Documentation
Oracle Identity Management, Oracle Access Manager	Status of an Oracle WebLogic Server Managed Server for Oracle Access Manager: "Monitoring an Oracle WebLogic Server Administration or Managed Server" section in the *Oracle Fusion Middleware Administrator's Guide* for instructions on accessing the server home page. Specifically, review the Summary section of the page.
	Performance metrics: "Viewing Performance in Fusion Middleware Control" section in the *Oracle Fusion Middleware Administrator's Guide for Oracle Access Manager with Oracle Security Token Service*
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 *Oracle Fusion Middleware Administrator's Guide* 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 *Oracle Fusion Middleware Administrator's Guide* for instructions on accessing the server home page. Specifically, review the Summary section of the page.
	Performance Metrics: "Monitoring Oracle Internet Directory" in the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*
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 *Oracle Fusion Middleware Administrator's Guide* 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 *Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory*
Oracle SOA Suite	Performance metrics: "Monitoring Processing Requests" section in the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite* and Section 10.4.4.2
Oracle WebCenter	Status of an Oracle WebLogic Server Managed Server for Oracle WebCenter: "Monitoring an Oracle WebLogic Server Administration or Managed Server" section in the *Oracle Fusion Middleware Administrator's Guide* 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 *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter*
Oracle Web Tier, Oracle HTTP Server	Performance metrics: "Monitoring Oracle HTTP Server Performance" section in the *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server*
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 Section 5.11.1.1, "Creating or Editing a Work Assignment."

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 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 (%) Active 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 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 Work Manager Requests (per minute) Work Manager Stuck 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 **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)	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	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)	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)	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	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)	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)	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	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)	1. From the navigation pane, expand **Scheduling Services**.
Wait Time	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	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 Grid 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.

Grid 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 Grid 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 10–3 shows the **Middleware** page for Oracle WebLogic Server domains in the Oracle Fusion Applications environment.

Figure 10–3   Targets Page: Middleware Tab of Grid Control

To access the pages for Middleware page in Grid Control:

1. From the home page, click the **Targets** tab.

2. From the **Targets** tab, click the **Middleware** subtab.

    The Middleware page displays.

3. From the **Search** list, select a specific target type to monitor. In Figure 10–3, the target selected is **Oracle WebLogic Server Domain**.

For more information on 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 Grid Control:

1. Click **Setup** to access the Setup Overview page.

2. From the left navigation pane, select **Monitoring Templates** to access the Monitoring Templates main page.

3. Select a template from the list and click **Create Like**. Optionally, you can click **Create** to create a new template without any data from another template.

4. Fill out the tabs accordingly.

In the **Metric Threshold** tab, select the metrics to add to the template. Use the tables in the section for tips on recommended metrics to add to the template.

- Target Type: Host
- Target Type: Oracle WebLogic Server
- Target Type: Oracle HTTP Server
- Target: SOA Composite
- Target: SOA Infrastructure

Use Table 10–9 to help add metrics for the Host target type to a monitoring template.

Table 10–9   Target Type: Host

Metric Name	Metric Column Name
CPU Usage	CPU System Time (%)
	CPU Utilization (%)
Load	CPU in I/O Wait (%)
	Run Queue Length (1 minute average
	Memory Utilization (%)
	Swap Utilization
Network Interface	Network Interface Total I/O Rate (MB/sec) - Public
	Network Interface Total I/O Rate (MB/sec) - Private
Disk Activity	Average Disk I/O Service Time (ms)
	Average Disk I/O Wait Time (ms)

Use Table 10–10 to help add metrics for the Oracle WebLogic Server target type to a monitoring template.

Table 10–10   Target Type: Oracle WebLogic Server

Metric Name	Metric Column Name	Warning Threshold	Critical Threshold
Server Overview Metrics	Active Sessions	250	300
	Open JDBC Connections	150	200
	Request Processing Time (ms)	10000 (10sec)	15000 (15sec)
JVM Metrics	Active Threads	1000	1250
	CPU Usage (%)	70%	80%
	Heap Usage (MB)	1.8GB	2.0GB
JVM Threads	JVM Threads - Deadlocked Threads	2	5
Response	Status	Down	>30Mins
Server Datasource Metrics	Server Datasource - Connections in Use	150	200
	Server Datasource - Connection Requests Waiting	5	10
Server Work Manager Metrics	Server Work Manager - Stuck Threads	2	4

Use Table 10–11 to help add metrics for the Oracle HTTP Server target type to a monitoring template.

Table 10–11   Target Type: Oracle HTTP Server

Metric Name	Metric Column Name
OHS Server Metrics	Request Processing Time (seconds)
	Response Data Processed (KB per response)
OHS Virtual Host Metrics	Request Processing Time for a Virtual Host (seconds)
	Response Data Processed for a Virtual Host (KB per response)
OHS Response Code Metric	HTTP 4xx errors
	HTTP 5xx errors
Resource Usage	CPU Usage (%)
	Memory Usage (%)

Use Table 10–12 to help add metrics for the SOA Composite target type to a monitoring template.

Table 10–12   Target: SOA Composite

Metric Name	Metric Column Name
Mediator Case	Invocation count throughput in last 5 minutes
SOA Composite - Response Metrics	Composite Status
SOA Composite - Component Detail Metrics	Component: Business Faults
	Component: Error Rate (%)
SOA Composite - Services/References Detail Metrics	Service/Reference: Average Incoming Messages Processing Time (ms)
	Service/Reference: Average Outbound Messages Processing Time (ms)
	SOA Composite: Error Rate (%)
	SOA Composite: Synchronous Response Time (ms)
	SOA Composite: Total Business

Use Table 10–13 to help add metrics for the SOA Infrastructure target type to a monitoring template.

Table 10–13   Target: SOA Infrastructure

Metric Name	Metric Column Name
SOA Infra Response	Up Down Status
SOA Infrastructure - Message Detail Metrics	Errors (minute)
SOA Infrastructure - Service Engine Detail Metrics	Service Engine: Error Rate (%)

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 Grid Control:

1. Click **Setup** to access the Setup Overview page.

2. From the left navigation pane, select **Blackouts** to access the Blackouts main page.

3. From the table, click **Create**.

    Grid 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 Grid Control:

1. Click **Preferences**.

2. In the left navigation pane, select **Rules**.

3. In the Notification Rules page, click Create to create a new rule.

4. When entering data in the **General** tab, from the **Target Type** list, choose the target type, or use one of the **Apply rule** options to apply to rule to all target types on which the corrective actions will run.

5. On the Notification Rule Metrics page, click **Add** to go to the Add Metrics page.

6. On the Add/Edit Metrics page, select the metrics against which the corrective actions will run, then scroll down the page to **Corrective Action States**.

7. Select the states for which you would like to receive e-mail notifications. For example, if you clicked on **Problem** for **Corrective Actions for Critical**, this means if the metric triggers a critical alert and a corrective action is executed, then a notification will be sent only if there is a problem with the corrective action execution. Click **Continue**.

8. On the Notification Rule Methods page, click on the **Send Me E-mail** option. Be sure that your e-mail addresses and notification schedule have been set up (both of these are located in the **Preferences** section of the Grid Control console).

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 Middleware Fusion Developer's Guide for Oracle Application Development Framework* for details about parameters

■ *Oracle Fusion Middleware Performance and Tuning Guide* to identify and tune the top tuning areas for middle tier in Oracle Fusion Middleware

■ "Managing Log Files and Diagnostic Data" in the *Oracle Fusion Middleware Administrator's Guide*

11

Diagnosing the Oracle Fusion Applications Middle Tier Performance

This chapter describes how to diagnose Java applications in the middle tier using Oracle Enterprise Manager Grid Control (Grid Control):

This chapter includes the following topics:

- Introduction to Java Diagnostics in the Middle Tier
- 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

11.1 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 Grid Control (Grid 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.

For a better understanding of JVM diagnostics, see the chapters in "JVM Diagnostics" part in the *Oracle Enterprise Manager Getting Started with Fusion Middleware Management*.

11.2 Viewing JVM Diagnostics Summary Information

The JVM Diagnostics Summary page in Grid Control is the best starting place for diagnosing performance problems in Java applications. It provides diagnostics summary details for a selected period for an entire Oracle WebLogic Server domain and each of its Managed Servers and Administration Server.

Figure 11–1 shows the JVM Diagnostics Summary page for an Oracle WebLogic Server domains in the Oracle Fusion Applications environment.

Figure 11–1   JVM Diagnostics Summary Page
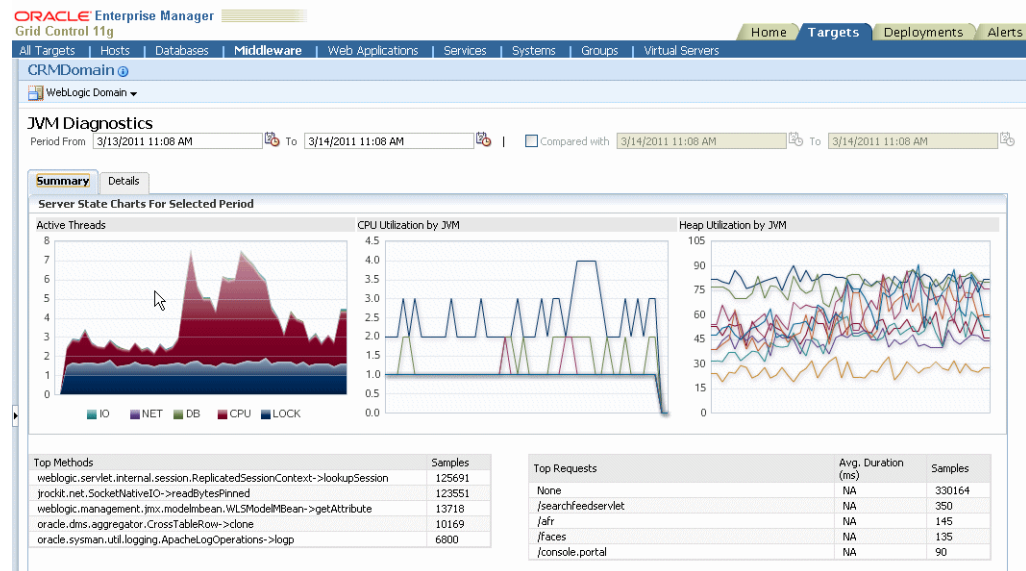
The following charts display on this page:

■ **Server State Charts For Selected Period:** 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:** 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 percentage of CPU used by the system and processes in the time period.

– **Heap Utilization:** This chart shows the percentage of computer memory used by the Java heap in the time period.

– **Garbage Collections:** 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.

■ **Top Methods:** This section shows the most expensive Java methods in the selected time period.

■ **Top Requests:** This section shows the top page requests in the selected time period.

The **Compared with** feature enables you to compare the diagnostics across two specified periods of time.

To access the JVM Diagnostics Summary page in Grid Control:

1. Click the **Targets** tab.

2. Click the **Middleware** secondary tab.

3. From the **Search** list, select **Oracle WebLogic Server Domain**, and then click **Go**.

4. Click on a domain or one of its servers which has JVMD agents deployed.

   The WebLogic Server Domain Home page displays.

5. From the **WebLogic Domain** menu, choose **JVM Diagnostics** > **Summary**.

6. Click the **Summary** tab, which is shown in Figure 11–1.

7. You can click the **Details** tab for a a more detailed view of the diagnostic data, as shown in the following figure:
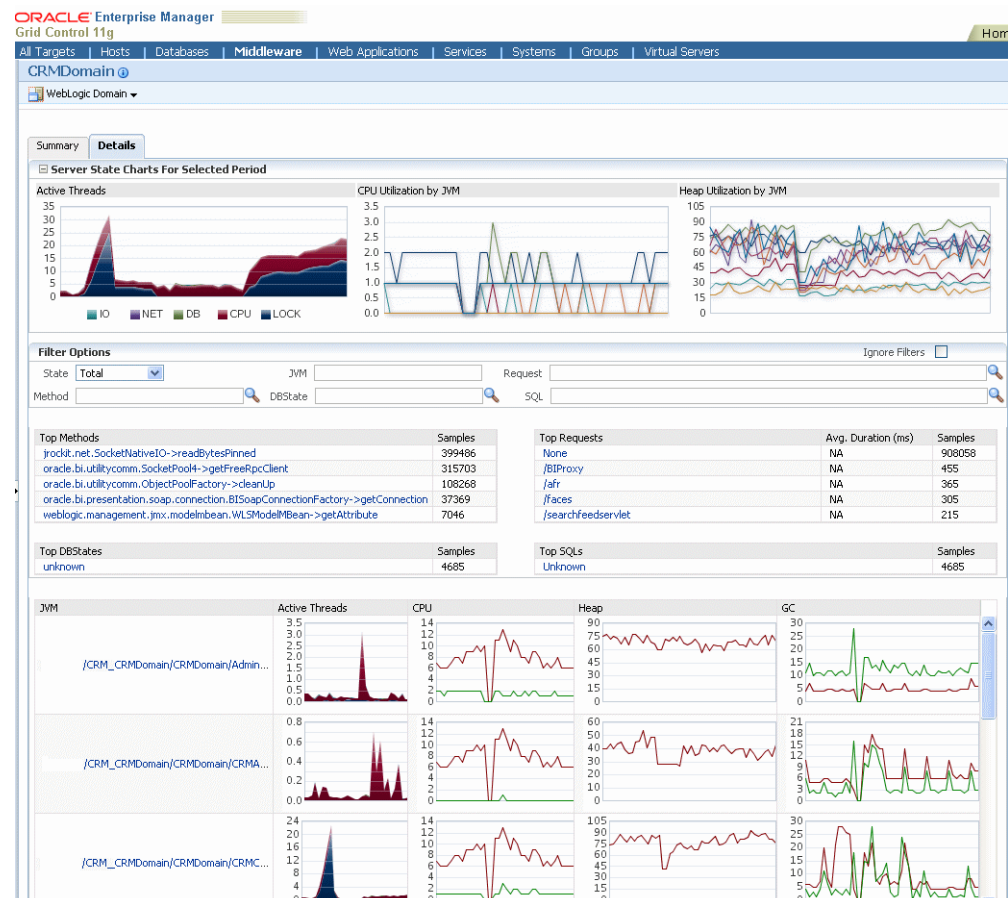
11.3 Finding the Top Java Methods

If you have a slow-running application, locate the Java method causing the potential issue.

To find the top Java methods with Grid Control:

1. Click the **Targets** tab.

2. Click the **Middleware** secondary tab.

3. From the **Search** list, select **Oracle WebLogic Server Domain**, and then click **Go**.

4. Click on a domain which has JVMD agents deployed.

    The WebLogic Server Domain Home page displays.

5. From the **WebLogic Domain** menu, choose **JVM Diagnostics** > **Summary**.

6. Click the **Details** tab.

7. In the **Top Methods** section, review the list of methods ordered by their cost (the number of samples).

8. In the **Top Methods** section, click on a method to add it as a filter.

    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)

9. Click on the plus sign next to the method to investigate further.

    The Method Stack Tree dialog displays with a call stack that you can traverse.

10. Click **OK** to close the Method Stack Tree dialog.

11. In **Filter Options**, clear out the **Method** text box filter and click anywhere to remove the filter.

11.4 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 Grid Control:

1. Click the **Targets** tab.

2. Click the **Middleware** secondary tab.

3. From the **Search** list, select **Oracle WebLogic Server Domain**, and then click **Go**.

4. Click on a domain which has JVMD agents deployed.

    The WebLogic Server Domain Home page displays.

5. From the **WebLogic Domain** menu, choose **JVM Diagnostics** > **Summary**.

6. Click the **Details** tab.

7. In the **Top SQLs** section, review the list of SQL calls ordered by their cost (the number of samples).

8. 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 DBStates** chart)

9. In **Filter Options**, clear out the **SQL** field and click anywhere to remove the filter.

11.5 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 Grid Control:

1. Click the **Targets** tab.

2. Click the **Middleware** secondary tab.

3. From the **Search** list, select **Oracle WebLogic Server Domain**, and then click **Go**.

4. Click on a domain which has JVMD agents deployed, or click on a specific server having problems.

5. From the **WebLogic Domain** menu or **WebLogic Server** menu, choose **JVM Diagnostics** > **Threads** > **Real-Time Analysis**.

6. In the **JVMs** section, click on a thread in the upper section show details in the **Threads** section.

7. In the **Threads** section, look for a thread having the prefix **[STUCK THREAD]** and click on it.

8. Look at the fields **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).

9. Look for the **Thread State** and the **Locks** in the details view.

10. If the stuck thread is in the **DB Wait** state then click on the hyperlink and go directly to the database session to see what that thread is doing in the database, or use the technique described in Section 11.6.

11.6 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 Grid Control:

1. Click the **Targets** tab.

2. Click the **Middleware** secondary tab.

3. From the **Search** list, select **Oracle WebLogic Server Domain**, and then click **Go**.

4. Click on a domain which has JVMD agents deployed, or click on a specific server having problems.

5. From the **WebLogic Domain** menu, choose **JVM Diagnostics** > **Summary**.

6. Click the **Details** tab.

7. In the **Top SQLs** section, click on a SQL call to view the charts for that call.

8. In **Filter Options**, in the **SQL** field, copy the ID of the SQL call.

9. Click on **Databases** in secondary navigator.

10. Select the relevant target database.

11. Click the **Performance** tab.

12. In the **Additional Instance Monitoring Links** section, click **Search SQL**.

13. Paste the ID of the SQL call into the relevant field with any other choices you may need and then click **Search**.

14. Analyze the SQL.

11.7 Analyzing Potential Memory Leaks

To find and analyze memory leaks, you can use Grid 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, "Take a Heap Snapshot"

- Task 2, "Upload a Heap Snapshot"

- Task 3, "Analyze Heap Snapshots"

Task 1 Take a Heap Snapshot

To create a snapshot of the heap for later loading and examination for leaks:

1. Click the **Targets** tab.

2. Click the **Middleware** secondary tab.

3. From the **Search** list, select **Oracle WebLogic Server Domain**, and then click **Go**.

   Click on a domain which has JVMD agents deployed, or click on a specific server having problems.

4. From the **WebLogic Domain** menu, choose **JVM Diagnostics** > **Heap** > **Take Snapshot Now**.

   The dump Java heap page displays with the name of the file, `heapdump#.txt` in `/tmp` on Unix and `C:\` on Windows. This file lists all the heap objects and their relationships that you can load and examine for leaks.

   To modify the default location of the `heapdump#.txt` file:

   a. From the **WebLogic Domain** menu, choose **JVM Diagnostics** > **Setup**.

   b. Click the **JVMs** tab.

   c. Select the edit icon for a JVM ID.

   d. In the Edit JVM Information page, modify the **Heap Dump Dir** field.

   e. Click **Save.**

Task 2 Upload a Heap Snapshot

To create a snapshot of the heap for later loading and examination for leaks:

1. Click the **Targets** tab.

2. Click the **Middleware** secondary tab.

3. From the **Search** list, select **Oracle WebLogic Server Domain**, and then click **Go**.

   Click on a domain which has JVMD agents deployed, or click on a specific server having problems.

4. From the **WebLogic Domain** menu, choose **JVM Diagnostics** > **Setup**.

5. Click the **Downloads** tab.

6. For **Load Heap**, click on the **Download** icon.

7. Transfer the `loadheap.zip` file to a computer that has Oracle Database client installed.

8. Unzip `loadheap.zip` to some known location, for example, `~/jvmd`.

   Steps 1 - 8 are one-time steps. The next heap dump that needs to be loaded does not require the above steps to be repeated in order to upload a heap snapshot.

9. Transfer the heap dump files to the same host, for example, `/scratch/jvmd/heapdump22.txt`.

10. In the known location, go to the `loadheap` directory. For example:

    ```
    cd ~/jvmd/loadheap
    ```

11. Run the `loadpheap` utlity for the file you created:

    ```
    (UNIX) loadheap.sh HEAP_FILE_NAME
    (Windows) loadheap.cmd HEAP_FILE_NAME
    ```

    where *HEAP_FILE_NAME* is the heap dump file you created in Task 1, "Take a Heap Snapshot".

    ```
    ./loadheap.sh /scratch/jvmd/heapdump22.txt
    ```

    While running the command, you will be asked to enter in the connect information to your repository (host name, SID, username and password). This command can be time consuming process. After the command is finished, scroll up to examine if there has been any error.

    If you want to compare heaps, you run `loadheap` at least twice to get a pair of heaps loaded and compared.

Task 3 Analyze Heap Snapshots

To analyze heap snapshots:

1. Click the **Targets** tab.

2. Click the **Middleware** secondary tab.

3. From the **Search** list, select **Oracle WebLogic Server Domain**, and then click **Go**.

   Click on a domain which has JVMD agents deployed, or click on a specific server having problems.

4. From the **WebLogic Domain** menu, choose **JVM Diagnostics** > **Heap** > **Saved Heap Snapshots**.

The **Available Heaps** table lists all the snapshots that have been loaded into the repository.

5. Examine one of the heaps by clicking on the number in the **%** column.

   The Heap Usage page displays.

6. Click the **Roots** tab to see the objects reachable by roots, which are objects that are directly reachable from the JVM itself.

   The Show Heap Usage by Roots page displays.

7. Click a root to drill down and view the objects that consume a lot of memory. They are most likely to be in the **Dictionary** root.

   The Top 40 Objects in Threadlocals page displays.

8. Click the **Compare Heaps** tab 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.

9. Select the second heap by clicking on the number in the **%** column.

10. Compare both the heaps. Compare the number of objects and the occupied memory size in each heap dump. This measure indicates the objects that are growing over the period of time when the snapshots were taken.

11. Drill down into the row which had the largest delta (**Adjusted**) in order to find the biggest memory leak.

12

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 plans 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.
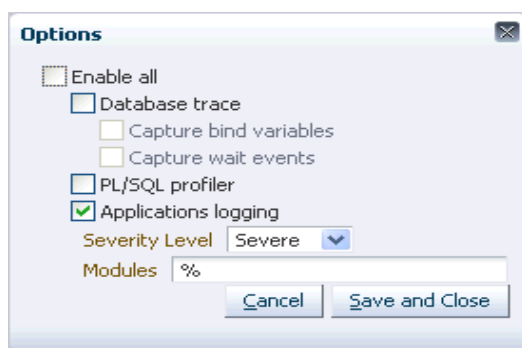
5. Click **Save and Close**.

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*

13

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 17.

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.

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.

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 17. For more information about monitoring log files, see Section 13.6, Section 13.6.2, 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 17.

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

A standard tool that you can use to complete these tasks is the Oracle Enterprise Manager Fusion Applications Control user interface. 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. It is recommended that you use Fusion Applications Control for viewing and searching Oracle Fusion Applications log files, but it is also possible to view and search log files using Oracle Enterprise Manager Fusion Middleware Control.

> **Note:** Oracle Enterprise Manager Fusion Applications Control is 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 17.

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.

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

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

- 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 Oracle Fusion applications that use standard logging functionality, you can view the log files from Oracle Enterprise Manager 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 Managed Server for a particular Oracle Fusion application, sometimes called an Oracle Fusion application instance. When you select this target type and view log files, you view the log for the selected individual server.

The Fusion Applications Control screen 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 Oracle Enterprise Manager 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. 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

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)**.

   > **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 Managed Server for a particular Oracle Fusion application as your target, the content pane displays just the log messages for the selected Oracle Fusion application 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. 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. No wildcard characters are necessary.

   > **Note:** If you search using more than one criterion, the search results display only log entries that match all of the specified criteria (logical AND).

5. 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.`

6. Click **Search.**

13.6.3 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="path_to_log_directory/server_
name-diagnostic.log"/>
    <property name="maxFileSize" value="104857600"/>
    <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_ROLE_IDS,
    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"/>
  </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 17.2.9 and Section 17.5.2.

- 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 17.2.10.

- 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 17.2.11.

13.6.4 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 17.5.2. For information about viewing nonstandard log files for Oracle Fusion General Ledger, see Section 17.2.10. 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 Oracle Enterprise Manager 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 Oracle Enterprise Manager 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 17.2.1.

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** screen 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 17.2.

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_ FILENAME (FND: Log File for C)	C only	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.	/temp/Cdiagnostic.log	Site, User	diagnostic.log
		It is recommended that you set the value of this profile option 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.			
		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.			
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).	1000, 900, 800, 700, 500, 400, 300	Site, User	1000
		For more information about log levels, see Section 13.7.3.			
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

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_ 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	`%`
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	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. By default, the `SITE` level value for this profile option is set to `APPLLOG_ DIR/diagnostic.log`, where `APPLLOG_DIR` is a directory object that was defined using the Oracle Fusion Applications Repository Creation Utility during installation. If you would like the log file name to indicate that the log messages are generated from PL/SQL, you can change the `AFLOG_PLSQL_FILENAME` profile option value to a value such as `APPLLOG_ DIR/plsqldiagnostic.log`. For more information about directory objects, see the section about the Create Directory command in the *Oracle Database SQL Language Reference*.	`APPLLOG_ DIR/mylog.log`	Site, User	`APPLLOG _ DIR/dia gnostic .log`

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 (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
Log Levels for Reporting Failures and Normal Events:			

ERROR:1	SEVERE	1000	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:
			
			`The process cannot be started now because another process is using the data.`
			
			`The participant's rate could not be found. Contact your benefits administrator to validate the standard rate for this compensation object.`

WARNING:1	WARNING	900	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. For example, the following error message might be associated with this log level: `You must change your password before date.`
NOTIFICATION:1	INFO	800	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.
Log Levels for Tracing and Reporting Progress:			
NOTIFICATION:16	CONFIG	700	Configuration properties and environment settings. This is a finer level of granularity for reporting normal events.

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	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.
			For example, the following error message might be associated with this log level:
			`User has chosen to Cancel the Purge operation.`
TRACE:16	FINER	400	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.
			For example, the following error message might be associated with this log level:
			`Entering validateLogin.`
TRACE:32	FINEST	300	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.
			For example, the following error message might be associated with this log level:
			`Copying string from buffer xyz to buffer zzz.`

> **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 17.

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 Oracle Enterprise Manager 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 system 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 17.

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 Oracle Enterprise Manager 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 17.2.

You can use either Oracle Enterprise Manager 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.

g. If you want PL/SQL log records to be written to a log file location other than the default, enter the path and file name in the **File Name** field in the **PL/SQL Settings** section.

h. If you want C log records to be written to a log file location other than the default, enter the path and file name in the **File Name** field in the **C Settings** section.

i. Click **OK.**

Skip Step 5.

5. If you clicked **Site-Level**, complete the following substeps:

a. 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.

b. Make sure that **Logging Enabled** is selected.

c. In the **Module Filters** field, indicate if you want Oracle Fusion Applications to log only for a particular code module—by default, log entries are written for all Oracle Fusion Applications code modules that use standard logging implementation for Oracle Fusion Applications.

d. 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.

e. 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.

f. If you want PL/SQL log records to be written to a log file location other than the default, enter the path and file name in the **File Name** field in the **PL/SQL Settings** section.

g. If you want C log records to be written to a log file location other than the default, enter the path and file name in the **File Name** field in the **C Settings** section.

h. Click **Apply.**

13.7.6 Using Additional Settings to Configure Logs for Selected Components

Some functionality areas in Oracle Fusion Applications use nonstandard logging mechanisms that are disabled by default. For these areas, 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. Some other functionality areas use nonstandard logging mechanisms that are usually set to gather minimal amounts of information. You may want to change those settings when troubleshooting.

This section contains the following topics:

■ Configuring Additional Log Settings for Oracle Fusion Incentive Compensation Batch Jobs

■ Configuring Additional Log Settings for Oracle Fusion General Ledger

■ Configuring Additional Log Settings for Oracle Fusion Receivables AutoInvoice

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 17.2.9.

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 17.2.10.

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 17.2.11.

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 Oracle Enterprise Manager Fusion Applications Control

By default, you can use Oracle Enterprise Manager 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 13.6.3 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 `plsqlfndlog.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** entry in the navigation tree, then expand the **Fusion Applications** entry, then select the Oracle Fusion cluster application **PayablesApp**.

b. In the context pane, from the dynamic target menu, choose **Logs > View Log Messages** to display the aggregated log messages for the selected target.

For an example in which the selected target is an Oracle Fusion Applications cluster application, the aggregated log messages include log messages from all of the Managed Servers in the cluster.

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.

To display to the Diagnostic Dashboard:

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

13.10 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 or by Categorization Tag

- 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, using an account that has one of the following job roles:

   - `Diagnostic Regular User`

   - `Diagnostic Advanced User`

- `Diagnostic Administrator`

For 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.**

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, or click the alert icon in the **Input Status** column to display the Input Parameters dialog box.

- For any root level test that displays the message `Required 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.

> **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 Input Specified,` 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 Oracle Enterprise Manager 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

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 substeps 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* 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, navigate to the location of the `diagctl` executable under the `fusionapps` Oracle Fusion Middleware home directory:

   ```
   (UNIX) MW_HOME/atgpf/bin/diagctl.sh
   (Windows) MW_HOME\atgpf\bin\diagctl.cmd
   ```

5. 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 or by Categorization Tag

You can search for diagnostic tests either by specifying part of the name of the test or by specifying a categorization tag that is associated with the test.

In general, searching for diagnostic tests is done as a portion of the process of running diagnostic tests. For more information, see Section 13.10.1.

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. 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. 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**.

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 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) MW_HOME/atgpf/bin/diagctl.sh
(Windows) 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 Section 5.7.5.

– 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. 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) MW_HOME/atgpf/bin/diagctl.sh
   (Windows) 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, as follows:

■   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. 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.

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 14, "Backing Up and Recovering Oracle Fusion Applications"
- Chapter 15, "Configuring High Availability and Scaling Out Oracle Fusion Applications"
- Chapter 16, "Moving Components for Oracle Fusion Applications Across Environments"

14

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.

This chapter contains the following topics:

- Introduction to Backup and Recovery

- Overview of Backing Up Your Environment

- Overview of Recovering Your Environment

- Backup and Recovery Recommendations for Oracle Fusion Applications

- Performing a Backup

- Creating a Record of Your Oracle Fusion Applications Configuration

- Recovering After Data Loss, Corruption, or Media Failure

- Recovering After Loss of Host

- A Case Study: Recovering Oracle Fusion Customer Relationship Management

14.1 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.

14.2 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

14.2.1 Tools to Use to Back Up Your Environment

To back up your Oracle Fusion Applications environment, you can use:

- 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 14.2.2 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.

14.2.2 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.

14.2.3 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/bea/beahomelist
  (Windows) C:\bea\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.

14.2.4 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 14.2.3. 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 14.2.3. 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 May 20, 2011, name the backup:

  ```
  ApplBase_backup_052011.tar
  ```

14.3 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

14.3.1 Tools to Use to Recover Your Environment

To recover your Oracle Fusion Applications environment, you can use:

- 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.

14.3.2 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.

14.3.3 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 14.7.9 for procedures for recovering them to the same point in time and reconciling differences.

14.4 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 in 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 14.5, Section 14.7, and Section 14.8.

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

14.4.1 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 14.4.13 for information about backing up and recovering Oracle E-Mail and Web Marketing.

Dependencies on Oracle Fusion Middleware Components

Oracle Metadata Service, Oracle Business Intelligence metadata, Oracle Real-Time Decisions, an LDAP provider such as Oracle Internet Directory, Oracle WebCenter (including tagging, group spaces, and the Oracle WebCenter schemas), portlet metadata, Oracle SOA Suite, Oracle ADF user interface, Oracle Essbase, Oracle Universal Content Management, Oracle Enterprise Crawl and Search Framework metadata and data, including Oracle Secure Enterprise Search, Oracle Enterprise Data Quality for Product Data

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 Universal Content Management

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 14.4.13.

Back up the databases containing the Oracle Fusion Applications, the LDAP store, and Oracle Fusion Middleware schemas, including those for Oracle Universal Content Management.

Back up Oracle Universal Content Management, as described in "Backup and Recovery Recommendations for Oracle Universal Content Management" 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 14.4.13.

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 14.7.10.1.

Recover Oracle Universal Content Management, as described in "Backup and Recovery Recommendations for Oracle Universal Content Management" in the *Oracle Fusion Middleware Administrator's Guide*.

Recover the databases containing the Oracle Fusion Applications, Oracle Fusion Middleware schemas, including those for Oracle Universal Content Management.

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 14.7.9.

14.4.2 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 (portlets), Oracle Metadata Service, Oracle Imaging and Process Management, 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 Universal Content Management, as described in "Backup and Recovery Recommendations for Oracle Universal Content Management" 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 Universal Content Management.

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 14.7.

Recover Oracle Universal Content Management, as described in "Backup and Recovery Recommendations for Oracle Universal Content Management" 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 14.7.9.

14.4.3 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 (portlets), Oracle Metadata Service, Oracle Enterprise Scheduler, Oracle Business Intelligence Publisher, an LDAP provider such as Oracle Internet Directory, Oracle Fusion Middleware Extensions for Applications, Oracle Universal Content Management

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 Universal Content Management

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 Universal Content Management, as described in "Backup and Recovery Recommendations for Oracle Universal Content Management" 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 Universal Content Management.

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 14.7.

Recover Oracle Universal Content Management, as described in "Backup and Recovery Recommendations for Oracle Universal Content Management" 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 Universal Content Management 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 14.9.9.

14.4.4 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 Universal Content Management

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 Universal Content Management, as described in "Backup and Recovery Recommendations for Oracle Universal Content Management" 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 Universal Content Management.

Recovery Recommendations

Recover the Managed Server to which the application is deployed.

Recover Oracle Universal Content Management, as described in "Backup and Recovery Recommendations for Oracle Universal Content Management" 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 14.7. For the steps specific to recovering from loss of host, see Section 14.8.6.1.

14.4.5 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 Service, an LDAP provider such as Oracle Internet Directory, Oracle WebCenter (including tagging, group spaces, and the Oracle WebCenter 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 Universal Content Management

Backup Recommendations

Back up the Applications base directory and the Administration Server domain in which Oracle Fusion Project is deployed.

Back up Oracle Universal Content Management, as described in "Backup and Recovery Recommendations for Oracle Universal Content Management" 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 Universal Content Management, as described in "Backup and Recovery Recommendations for Oracle Universal Content Management" 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 14.7.

14.4.6 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 Universal Content Management

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 Universal Content Management, as described in "Backup and Recovery Recommendations for Oracle Universal Content Management" 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 Universal Content Management.

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 14.7.10.2. For the steps specific to recovering from loss of host, see Section 14.8.6.2.

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 Universal Content Management, as described in "Backup and Recovery Recommendations for Oracle Universal Content Management" 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 14.7.9.

14.4.7 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 Universal Content Management

Backup Recommendations
Back up the Applications base directory and the Administration Server domain in which Oracle Fusion Incentive Compensation is deployed.

Back up Oracle Universal Content Management, as described in "Backup and Recovery Recommendations for Oracle Universal Content Management" 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 Universal Content Management.

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 14.7.

Recover Oracle Universal Content Management, as described in "Backup and Recovery Recommendations for Oracle Universal Content Management" 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 14.7.9.

14.4.8 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 (including portlets, tagging, group spaces, and forums), Oracle WebCenter, 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 Universal Content Management

Backup Recommendations

Back up the Applications base directory and the Administration Server domain in which Oracle Fusion Applications Technology is deployed.

Back up Oracle Universal Content Management, as described in "Backup and Recovery Recommendations for Oracle Universal Content Management" 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 Universal Content Management.

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 14.7.

Recover Oracle Universal Content Management, as described in "Backup and Recovery Recommendations for Oracle Universal Content Management" 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 14.7.9.

14.4.9 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 Universal Content Management, Oracle WebCenter (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 Universal Content Management

Backup Recommendations

Back up the Applications base directory and the Administration Server domain in which Oracle Fusion Setup is deployed.

Back up Oracle Universal Content Management, as described in "Backup and Recovery Recommendations for Oracle Universal Content Management" 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 Universal Content Management.

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 14.7.

Recover Oracle Universal Content Management, as described in "Backup and Recovery Recommendations for Oracle Universal Content Management" 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 Universal Content Management.

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 14.7.9.

14.4.10 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 Universal Content Management

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, see Section 14.7.11.2. For the steps specific to recovering from loss of host, see Section 14.8.7.2.

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 14.7.9.

14.4.11 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, see Section 14.7.11.3. For the steps specific to recovering from loss of host, see Section 14.8.7.3.

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 14.7.9.

14.4.12 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, see Section 14.7.11.4. For the steps specific to recovering from loss of host, see Section 14.8.7.4.

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 14.7.9.

14.4.13 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, see Section 14.7.11.5. For the steps specific to recovering from loss of host, see Section 14.8.7.5.

14.5 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 Full Offline Backup
- Performing an Online Backup of Configuration Files

14.5.1 Performing a Full Offline Backup

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 14.2. 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_052011.tar APPLICATIONS_BASE/*
   (Windows) jar  cf ApplBase_backup_052011.tar APPLICATIONS_BASE\*
   ```

3. Back up the Applications configuration directory on all hosts. For example:

   ```
   (UNIX) tar -cpf ApplConfig_backup_052011.tar APPLICATIONS_CONFIG/*
   (Windows) jar  cf ApplConfig_backup_052011.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.

   For example:

   ```
   (UNIX) tar -cpf domain_home_backup_052011.tar
   APPLICATIONS_CONFIG/instance/domains/domain_name/*
   (Windows) jar  cf domain_home_backup_052011.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 14.2.4 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_052011.tar ORACLE_INSTANCE/*
   (Windows) jar  cf instance_home_backup_052011.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_052011.tar
   APPLICATIONS_CONFIG/instance/domains/domain_name/servers/server_name/*
   (Windows) jar  cf man_server1_backup_052011.jar
   APPLICATIONS_CONFIG\instance\domains\domain_name\servers\server_name\*
   ```

7. Back up the OraInventory directory. For example:

   ```
   (UNIX) tar -cpf Inven_home_backup_052011 /scratch/oracle/OraInventory
   ```

```
(Windows) jar cf Inven_home_backup_052011.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 14.6.

14.5.2 Performing an Online Backup of Configuration Files

You should perform a backup of configuration files on a regular basis and at the times described in Section 14.2.4.

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. For example:

   ```
   (UNIX) tar -cpf domain_home_backup_052011.tar
   APPLICATIONS_CONFIG/instance/domains/domain_name/*
   (Windows) jar  cf domain_home_backup_052011.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_052011.tar ORACLE_INSTANCE/*
   ```

```
(Windows) jar  cf instance_home_backup_052011.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*.

14.6 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.

14.7 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.

14.7.1 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_052011.tar
(Windows) jar xtf ApplBase_backup_052011.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.

14.7.2 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_052011.tar
(Windows) jar xtf mw_home_backup_052011.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.

14.7.3 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_052011.tar
(Windows) jar xtf domain_backup_052011.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 14.7.6.

5. If you cannot start a Managed Server, recover it, as described in Section 14.7.7.

14.7.4 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_052011.tar
(Windows) jar xtf Oracle_home_backup_052011.jar
```

2. Restart the Managed Server to which applications are deployed, using the WLST `start` command. For example:

```
wls:/mydomain/serverConfig> start('myserver','Server')
```

14.7.5 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

14.7.5.1 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_052011.tar
    (Windows) jar xtf instance_home_backup_052011.jar
    ```

3. Start all relevant processes. That is, start all processes that are related to that Oracle instance:

    ```
    opmnctl startall
    ```

14.7.5.2 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_052011.tar
    (Windows) jar xtf instance_home_backup_052011.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 -wlserverHome Middleware_Home
    ```

14.7.6 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.

14.7.7 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 14.7.10 and Section 14.7.11.

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 14.7.6.

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_052011.tar
   (Windows) jar xtf mw_home_backup_052011.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.

14.7.8 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 Universal Content Management, 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 14.7.9.

14.7.9 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

14.7.9.1 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.

14.7.9.2 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.

14.7.9.3 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.

14.7.10 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

14.7.10.1 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 14.7.7.

2. Recover the Java servers for Oracle E-Mail and Web Marketing, as described in Section 14.7.11.5.

3. Note that the Oracle Enterprise Data Quality for Product Data 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 14.9.

14.7.10.2 Recovering Oracle Fusion Procurement

To recover Oracle Fusion Procurement:

1. Recover the Managed Server to which the application is deployed, as described in Section 14.7.7.

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.

14.7.11 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

14.7.11.1 Recovering Oracle HTTP Server

To recover Oracle HTTP Server, you recover the Oracle instance that contains Oracle HTTP Server, as described in Section 14.7.5.

14.7.11.2 Recovering Oracle Enterprise Scheduler

To recover Oracle Enterprise Scheduler:

1. Recover the domain directory from the backup file, as described in Section 14.7.3.

2. Recover the Oracle home to the original directory from the backup file, as described in Section 14.7.4.

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 14.7.8.

14.7.11.3 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 14.7.3.

2. Recover the database containing schemas related to Oracle Enterprise Crawl and Search Framework, as described in Section 14.7.8.

14.7.11.4 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 14.7.3.

3. Recover the Oracle home to the original directory from the backup file, as described in Section 14.7.4.

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 14.7.8.

6. Start all relevant processes. That is, start all processes that are related to the domain, as described in Section 3.3.2.1.

14.7.11.5 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 14.7.7.

2. Recover the installation directory for the Bounce Handling Daemon.

14.8 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

14.8.1 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_052011.tar
(Windows) jar xtf ApplBase_backup_052011.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.

14.8.2 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

14.8.2.1 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 14.7.3.

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_052011.tar
      (Windows) jar xtf domain_backup_052011.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()
      ```

14.8.2.2 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 14.7.2 through Section 14.7.5.

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 14.8.5.3.

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/ta
    skhelp/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 14.8.3.2.

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 14.8.5.2.

10. Update Oracle Inventory, as described in Section 14.8.5.4.

11. Edit the targets.xml file for Fusion Middleware Control, as described in Section 14.8.5.1.

14.8.3 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

14.8.3.1 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_052011.tar
   (Windows) jar xtf mw_home_backup_052011.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.

14.8.3.2 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_052011.tar
   (Windows) jar xtf mw_home_backup_052011.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 Hostname 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, which is located at::

   http://download.oracle.com/docs/cd/E21764_01/apirefs.1111/e13952/ta
   skhelp/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
   ```

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 14.8.5.2.

7. Edit the targets.xml file for Fusion Middleware Control, as described in Section 14.8.5.1.

Now you can start and stop the Managed Server on Host C using the Administration Server running on Host A.

14.8.4 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 Universal Content Management, 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 14.7.9.

14.8.5 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

14.8.5.1 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.

14.8.5.2 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/F
usionVirtualHost_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>
```

14.8.5.3 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()
```

14.8.5.4 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/bea/beahomelist
(Windows) C:\bea\beahomelist
```

14.8.6 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

14.8.6.1 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 14.8.3.

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 14.8.3.

2. Recover the Oracle instance for Global Order Promising, as described in Section 14.7.5.

14.8.6.2 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 14.8.3.

To recover Oracle Fusion Procurement to a different host:

1. Recover the Managed Server to which the application is deployed, as described in Section 14.8.3.

2. Ensure that Oracle Business Intelligence Publisher and CUPS IP address and port number reflects the different host.

14.8.7 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

14.8.7.1 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 14.7.5.

To recover Oracle HTTP Server to a different host:

1. Recover the Middleware home, as described in Section 14.7.2.

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 14.8.5.1.

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 14.8.5.2.

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.

14.8.7.2 Recovering Oracle Enterprise Scheduler After Loss of Host

To recover Oracle Enterprise Scheduler, follow the procedure in Section 14.7.11.2.

14.8.7.3 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 14.7.11.3.

14.8.7.4 Recovering Oracle Authorization Policy Manager After Loss of Host

To recover Oracle Authorization Policy Manager, follow the procedure in Section 14.7.11.4.

14.8.7.5 Recovering Oracle E-Mail and Web Marketing After Loss of Host

To recover Oracle E-Mail and Web Marketing, follow the procedure in Section 14.7.11.5.

14.8.7.6 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 Enterprise Deployment Guide*.

14.9 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 14.3.

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

14.9.1 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 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

- Oracle Business Intelligence

- Oracle Essbase

- Oracle Enterprise Scheduler

- Oracle SOA Suite

- Oracle Universal Content Management

- 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, 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.

14.9.2 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, 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. 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_012111.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.

14.9.3 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_012111.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
```

14.9.4 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 3, "Stop the Managed Servers" in Section 3.3.2.2.

2. Stop the Administration Server for the CRMDomain, as described in Task 2, "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_012111.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.

14.9.5 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 CRMPerformanceServer_1 CRMAnalyticsServer_1 CRMSearchServer_1 CRMODIServer_1 CRMSOAServer_1 CustomerServer_1 SalesServer_1 MarketingServer_1 OrderCaptureServer_1	The entire CRMDomain. For example: *APPLICATIONS_CONFIG*/instance/domains/CRMDomain
FinancialDomain	AdminServer FinCommonServer_1 FinAnalyticsServer_1 FinSOAServer_1 LedgerServer_1 ESSServer_1	The entire FinancialDomain. For example: *APPLICATIONS_CONFIG*/instance/domains/FinancialDomain
SCMDomain	AdminServer ProductManagementServer_1 SCMSOAServer_1	The entire SCMDomain. For example: *APPLICATIONS_CONFIG*/instance/domains/SCMDomain
HCMDomain	AdminServer CoreSetupServer_1 FINSOAServer_1 TalentServer_1 ESSServer_1 HCMAnalyticsServer_1	The entire HCMDomain. For example: *APPLICATIONS_CONFIG*/instance/domains/HCMDomain
SetupDomain	AdminServer HomePageServer_1 HelpPortalServer_1 FunctionalSetupServer_1 WCServices_1 SpacesServer_1	The entire SetupDomain. For example: *APPLICATIONS_CONFIG*/instance/domains/SEtupDomain

14.9.6 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 14.7.6.

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_052011.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.

14.9.7 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_052011.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 14.8.7.1.

14.9.8 Recovering an Oracle Fusion Customer Relationship Management Cluster

In this scenario, the cluster CRMPerformanceCluster 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 CRMPerformanceServer_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 14.9.5 for Managed Servers to domain mappings.) Take the following steps:

1. Stop the cluster:

```
stop('CRMPerformanceCluster', 'Cluster')
```

2. Stop the Administration Server, as described in Task 2, "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_052011.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('CRMPerformanceCluster', 'Cluster')
```

> **Note:** You can use this procedure to recover a cluster if the membership is mistakenly deleted.

14.9.9 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 Universal Content Management, 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 14.7.9.

For more information, see the following topics:

- For data loss, corruption, or media failure loss, see Section 14.7.8.

- For loss of host, see Section 14.8.4.

15

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*.

15.1 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*

15.2 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*.

15.3 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

15.3.1 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

15.3.1.1 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.

15.3.1.2 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.

15.3.1.3 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 15–1 shows the Oracle Fusion Applications clients, the protocols they use to connect, and their high availability configuration.

Table 15–1   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

15.3.1.4 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.

15.3.1.5 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.

15.3.1.6 Oracle Fusion Applications Log File Location

For detailed information about Oracle Fusion Applications log file locations, see Section 13.6.3.

15.3.2 Oracle Fusion Applications High Availability Architecture and Failover Considerations

Figure 15–1 shows the high availability architecture of an Oracle Fusion Applications cluster deployment.

Figure 15–1   Oracle Fusion Applications High Availability Cluster Deployment
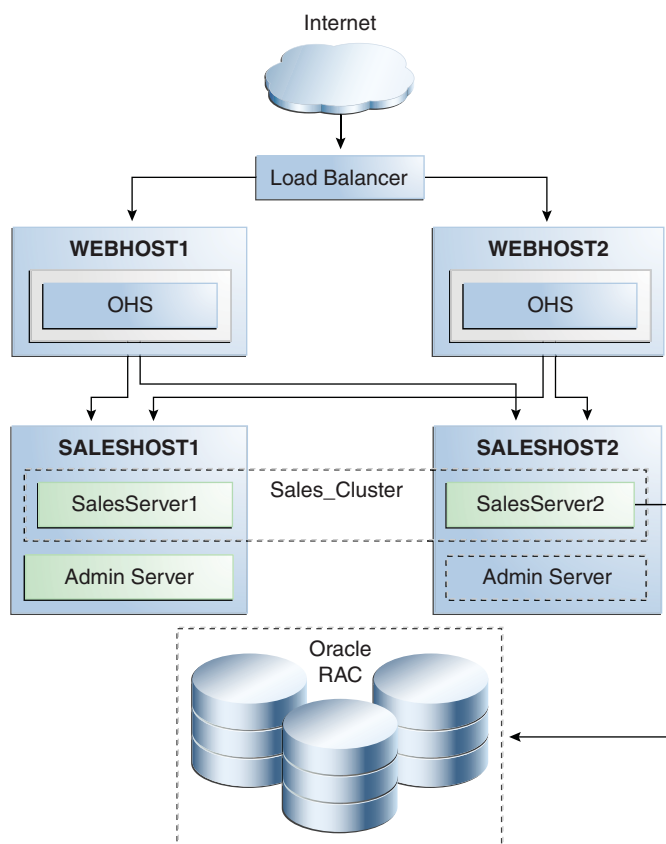
Some of the key characteristics of this architecture are:

- 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 15–1 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

15.3.2.1 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.

15.3.2.2 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.

15.3.2.3 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

15.3.2.3.1 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.

15.3.2.3.2 Machine Failure When a computer fails, the expected behavior is the same as for process failure (described in Section 15.3.2.3.1) except that Node Manager and OPMN do not attempt restarts.

15.3.2.3.3 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.

15.3.2.3.4 Troubleshooting Oracle Fusion Applications If you experience issues with your Oracle Fusion Applications high availability deployment, check the WebLogic Server log. See Section 15.3.1.6 for details.

15.4 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

15.4.1 Oracle Real Application Clusters

Oracle RAC and Oracle Clusterware[1] 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

---

[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.

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*.

For information about online application maintenance and upgrades, see the "Online Application and Maintenance and Upgrades" section of *Oracle Database High Availability Overview*.

15.4.2 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

15.4.2.1 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.

15.4.2.2 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 15.4.3 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*.

15.4.2.3 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*.

15.4.2.4 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 15.4.3 for recommendations when configuring load balancing with WebLogic Server.

15.4.2.5 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).

Section 15.4.3 to configure retry properties with WebLogic Server.

15.4.3 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 15–2.

Table 15–2   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 15–3.

Table 15–3   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*.

15.5 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 performing 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*.

15.5.1 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 for scaling 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**.

The Summary of Servers page appears.

d. Select **SalesServern** in the **Names** column of the table.

The Settings page for the server appears.

e. Click the **SSL** tab.

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\SetupDomain_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 port*n* 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**.

15.5.1.1 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.

15.5.2 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.

> **Note:** Before you scale out the topology, verify that the product binaries reside on the shared disk and that the domain home resides on the local disk. See the *Oracle Fusion Applications Enterprise Deployment Guide* for more information.

Follow these steps for scaling out the topology:

1. On the new computer, mount the existing *APPLICATIONS_BASE* (`/oracle` in Figure 1–6), if needed, which should include the Oracle Fusion Applications installation.

2. Log in to the Oracle WebLogic Administration Console.

3. Create a new computer for the new computer that will be used and add the computer to the domain. Follow these steps:

a. Navigate to *domain_name* and then choose **Environment** and then **Machines**.

For example, navigate to **CRM_Domain** and then choose **Environment** and then **Machines**.

b. In the Change Center, click **Lock & Edit** (if needed).

c. On the Summary of Machines page, click **New** to create a new computer.

d. In the **Name** field, enter the name of the computer (`SALESHOST2`) and choose the appropriate operating system from **Machine OS**.

      e. Click **OK**.

      f. Navigate to *domain_name* and then choose **Environment** and then **Machines**.

        For example, navigate to **CRM_Domain** and then choose **Environment** and then **Machines**.

      g. On the Summary of Machines page, click the name of the newly-created computer.

      h. On the Settings for *new-computer* page, click the Node Manager tab.

      i. On the Node Manager Properties page, in the **Type** field, enter SSL. In the **Listen Address** field, enter the name of the host (`SALESHOST2`). In the **Listen Port** field, enter the listen port (for example, 5556).

      j. Click **Save and Activate Changes**.

4. Use the Oracle WebLogic Server Administration Console to clone `SalesServer1` into a 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.

> **Note:** These steps assume that you are adding a new server to the new computer (`SALESHOST2` in this example), where no Managed Server was running previously.

---

5. For the listen address, assign the host name to use for this new Managed Server.

    Ensure the port number for this Managed Server is not in use on this computer.

6. 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** (if needed).

      d. Add the new server's address and port to the **Cluster Address** field. For example:

```
SALESHOST1:8001,SALESHOST2:8001
```

7. Run the pack command on `SALESHOST1` to create a template pack as follows:

```
(UNIX) SALESHOST1> cd ORACLE_COMMON_HOME/common/bin

SALESHOST1> ./pack.sh -managed=true -domain=APPLICATIONS_BASE/config/
SALESHOST1/CRMDomain/ -template=crmdomaintemplateScale.jar
-template_name=crm_domain_templateScale
```

Run the following command on `SALESHOST1` to copy the template file created to `SALESHOST2`:

```
(UNIX) SALESHOST1> scp crmdomaintemplateScale.jar oracle@SALESHOST2:/
APPLICATIONS_BASE/fusionapps/oracle_common/common/bin

(Windows) SALESHOST1> copy crmdomaintemplateScale.jar S:\fusionapps\
oracle_common\common\bin
```

The Windows command above assumes that `S:\` is a Windows shared drive that maps to the *APPLICATIONS_BASE* directory on `SALESHOST2`.

Run the unpack command on `SALESHOST2` to unpack the template in the Managed Server domain directory as follows:

```
(UNIX) SALESHOST2> cd APPLICATIONS_BASE/fusionapps/oracle_common/common/bin

SALESHOST2> unpack.sh -domain=APPLICATIONS_BASE/config/SALESHOST2/CRMDomain/
-template=crmdomaintemplateScale.jar

(Windows) SALESHOST2> cd APPLICATIONS_BASE\fusionapps\oracle_common\common\bin

SALESHOST2> unpack.sh -domain=APPLICATIONS_BASE\config\SALESHOST2\CRMDomain\
-template=crmdomaintemplateScale.jar
```

If the `config.xml` file is missing from the *APPLICATIONS_BASE*/`config` directory on `SALESHOST2` after the unpack operation, manually copy this file from the config directory on `SALESHOST1` to the `config` directory on `SALESHOST2`. For example:

```
(UNIX) SALESHOST1> cd APPLICATIONS_BASE/config

SALESHOST1> scp config.xml oracle@SALESHOST2:/APPLICATIONS_BASE/config/

(Windows) SALESHOST1> cd APPLICATIONS_BASE/config

SALESHOST1> copy config.xml S:\config\
```

The Windows copy command above assumes that `S:\` is a Windows shared drive that maps to the *APPLICATIONS_BASE* directory on `SALESHOST2`.

8. 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 `SALESHOST2`. 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. In the Oracle Enterprise Manager Console, select **Oracle WebLogic Server Administration Console**.

b. Expand the **Environment** node in the Domain Structure window.

c. Click **Servers**.

d. Select **SalesServer***n* in the **Names** column of the table.

e. Click the **SSL** tab in the Settings page.

f. Click **Advanced**.

g. Set Hostname Verification to **None**.

h. Click **Save**.

9. Verify the out-of-the-box configuration for the Node Manager during provisioning. If the new computer does not have a Node Manager, you must create the Node Manager directories by following these steps:

> **Note:** If the new computer has a Node Manager, skip this step.

a. Create a directory to store the configuration files for the Node Manager on this computer. In the example below, *node_name* is the listen address of the Node Manager, which is usually the physical hostname of the computer:

```
(UNIX) mkdir -p APPLICATIONS_BASE/admin/nodemanager/node_name
mkdir APPLICATIONS_BASE/admin/nodemanager/node_name/bin
mkdir APPLICATIONS_BASE/admin/nodemanager/node_name/common

(Windows) mkdir APPLICATIONS_BASE\admin\nodemanager\node_name
mkdir APPLICATIONS_BASE\admin\nodemanager\node_name\bin
mkdir APPLICATIONS_BASE\admin\nodemanager\node_name\common
```

The /common directory will be the Node Manager Home directory for this computer. The /bin directory will hold the Start script for the Node Manager on this computer.

b. The /common directory is the new Node Manager Home directory and will hold all the Node Manager configuration and log files (by default, the NodeManagerHome is *MW_HOME*/wlserver_10.3/common/nodemanager, where *MW_HOME* is the *APPLICATIONS_BASE*/fusionapps directory). At this point, the only file in that directory should be nodemanager.domains. Copy that file to the new Node Manager home directory:

```
(UNIX) cp MW_HOME/wlserver_10.3/common/nodemanager/* APPLICATIONS_BASE/
admin/nodemanager/node_name/common

(Windows) copy MW_HOME\wlserver_10.3\common\nodemanager\* C:\APPLICATIONS_
BASE\admin\nodemanager\node_name\common
```

The Windows copy command above assumes that C:\ is the Windows drive where the *MW_HOME* and *APPLICATIONS_BASE* directories are located.

c. The start script for Node Manager (startNodeManager.sh) on each computer will be customized. Therefore, copy it to the /bin directory on the new computer:

```
(UNIX) cp MW_HOME/wlserver_10.3/server/bin/startNodeManager.sh
Applications_Base/admin/nodemanager/node_name/bin

(Windows) copy MW_HOME\wlserver_10.3\server\bin\startNodeManager.sh
C:\Applications_Base\admin\nodemanager\node_name\bin
```

The Windows copy command above assumes that C:\ is the Windows drive where the *MW_HOME* and *APPLICATIONS_BASE* directories are located.

d. Edit the start script (startNodeManager.sh) that was copied to the new computer. Change this line:

```
NODEMGR_HOME="${WL_HOME}/common/nodemanager"
```

to this:

```
NODEMGR_HOME="$APPLICATIONS_BASE/admin/nodemanager/node_name/common"
```

Spell out the full directory name instead of typing *APPLICATIONS_BASE*.

e. Use this modified start script to start the Node Manager on the new computer:

```
$ cd APPLICATIONS_BASE/admin/nodemanager/node_name/bin
$ ./startNodeManager.sh
```

The *APPLICATIONS_BASE*/admin/nodemanager/*node_name*/common directory should now have configuration and log files for this Node Manager.

10. If the new computer already had Node Manager, ensure that the entry for the newly created domain that appears in *MW_HOME*/wlserver_ 10.3/common/nodemanager/nodemanager.domains also appears in *APPLICATIONS_BASE*/admin/nodemanager/*node_ name*/common/nodemanager.domains. Then add the domain path (for example, CRMDomain) to the nodemanager.domains file. Finally, start the Node Manager on the new computer using the installation in shared storage from the existing computers by passing the host name of the new computer as a parameter as follows:

```
SALESHOST2> WL_HOME/server/bin/startNodeManager new_computer_ip
```

11. Start and test the new Managed Server from the Oracle WebLogic Server Administration Console:

a. Ensure that the newly created Managed Server, SalesServer*n*, is running.

b. Access the application on the newly created Managed Server (http://SALESHOST2:8001/sales/faces/index). The application should be functional.

12. Configure Oracle HTTP Server to route to the Administration Server that contains the new Managed Servers, by setting the WebLogicCluster parameter to the list of computers in the cluster. Follow these steps for all the 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\SetupDomain_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,SALESHOST2:8001
</Location>
```

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 `SALESHOST2:8001/sales/faces/index` are accessible.

    e. Access the Sales server from Oracle HTTP Server using this URL:

        `WEBHOST1:7777/sales/faces/index`

    f. Login to the WebLogic Server Administration Console and navigate to **CRMDomain > Environments > Servers**.

    g. Shut down the `SALESHOST1:8001/sales/faces/index`.

    h. Access the Sales server from Oracle HTTP Server using this URL:

        `WEBHOST1:7777/sales/faces/index`

    If this works, it shows that Oracle HTTP Server routed the request to the new host (`SALESHOST2`) while `SALESHOST1` was down, and SALESHOST2 serviced the request.

You can also use these steps to scale out Managed Servers to additional hosts, for example, `SALESHOST3` and `SALESHOST4`.

15.5.2.1 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`).

If this works, it shows that Oracle HTTP Server routed the request to the new Managed Server (`SalesServer2`) while `SalesServer1` was down and `SalesServer2` serviced the request.

3. If you have not already done so, restart the Managed Server you shut down in the previous section.

15.5.3 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

15.5.3.1 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*.

15.5.3.2 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 15–3, " 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 15–4 but you can substitute the names and properties to scale any Oracle Fusion Applications cluster.

Table 15–4   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 datasource 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**.

16

Moving Components for Oracle Fusion Applications Across Environments

This chapter provides an overview of moving Oracle Fusion Applications components from one environment to another. It contains the following topic:

■ Introduction to Moving Oracle Fusion Applications

The procedures for performing this movement are not available at this time.

16.1 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 test environment. Once the system is stable and performs as desired, the production environment would then be created by moving all the components from the test environment, instead of redoing all the changes that were incorporated into the test 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 Universal Content Management 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 users, 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

Part V

Troubleshooting

This part provides information about how to troubleshoot the Oracle Fusion Applications environment. .

Part V contains the following chapters:

17

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.

17.1 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. Each automatically created incident 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.

All automatically created Oracle Fusion Applications incidents 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.

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.

17.1.1 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 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*.

17.1.2 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 "Diagnosing Problems" chapter in the *Oracle Fusion Middleware Administrator's Guide.*

17.1.3 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 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:

- Oracle Enterprise Manager 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 17.4.2, 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.

17.2 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

17.2.1 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 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:

   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.

17.2.2 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 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 incident information (older incidents are purged automatically when the allocated space limit is reached)

For more information, see the "Configuring the Diagnostic Framework" section in the *Oracle Fusion Middleware Administrator's Guide.*

17.2.3 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.

17.2.4 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 17–1. For information about additional QuickTrace properties, see Section 17.2.7.

Table 17–1   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 SITE but is not available for USER. AFLOG_INCIDENT_ENABLED is a new profile that has been introduced in Fusion to support application incident creation.	Y, N	Site	Y
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	Specifies the minimum level for tracing. This setting is similar to the AFLOG_LEVEL profile option for standard Oracle Fusion Applications logging to log files, but the two settings operate independently.	1000 (SEVERE), 900 (WARNING), 800(INFO), 700 (CONFIG), 500 (FINE), 400 (FINER), 300 (FINEST) and 0 (OFF)	Site	500 (FINE)

17.2.5 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 17.2.4.

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.*

17.2.6 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.

17.2.7 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 17.2.4.

In addition, Table 17–2 shows QuickTrace property settings that you can change using Fusion Applications Control.

Table 17–2   QuickTrace Properties

Property Name	Description	Server Restart Required	Default
Buffer Size (bufferSize)	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. By default, each WebLogic Server has a single COMMON QuickTrace buffer for storing multiple users' log records in memory. 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 Enable User Buffer to true, restarting the server, and setting the value of User Names for Reserve Buffer to a comma-separated list of those users.	Yes	5242880
Enable User Buffer (enableUserBuffer)	When this property is set to false, each QuickTrace handler uses a single COMMON buffer to cache log messages for all users. When this property is set to true, each QuickTrace handler maintains an individual buffer for each user who is listed in the User Names for Reserve Buffer (reserveBufferUserIDs) property. Log messages for users who are not listed in that property are cached in the COMMON buffer. 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 COMMON buffer.	Yes	false
User Names for Reserve Buffer (reserveBufferUserIDs)	Holds a list of application User Name or server User ID values separated by comma. When the Enable User Buffer (enableUserBuffer) 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 COMMON buffer. When the Enable User Buffer (enableUserBuffer) property is set to false, this property has no effect.	No	null

You can also change the properties in Table 17–2 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 17–2, 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.

17.2.8 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 setting the logging profile options, see Section 13.7.5.

17.2.9 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 17.5.2.

17.2.10 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.

17.2.11 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.

17.2.12 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_ALERTABLE` 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:

17.2.13 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 17.2.1.

17.2.14 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 13.7.5.

- 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 13.7.5.

- 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 13.7.5.

■ 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 17.2.9.

17.3 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 an incident is created, the Diagnostic Testing Framework automatically runs any Oracle Fusion Applications diagnostic tests that Oracle Fusion Applications developers have associated with the incident's particular error message. 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.

17.4 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

■ Viewing Problems and Incidents

■ Recovering from Incidents Generated During SOA Operations

■ Creating and Packaging Incidents Manually

■ Working with Automatically Created QuickTrace Dumps

■ Working with Other Diagnostic Dumps

17.4.1 Process for Investigating, Reporting, and Solving a Problem

When you become aware of a problem with a particular Oracle Fusion application, you can use the following general strategy to help diagnose and address the problem. While any problem may have unique features, the information that you gather in this way is likely to be useful during the troubleshooting process.

To help diagnose and address the problem, 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 17.4.4.

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 and how to work with QuickTrace dumps, see Section 17.4.2 and Section 17.4.5.

4. View recent information in the standard log files for all of the following:

   - Oracle Fusion Applications

   - Oracle Business Intelligence

   - Oracle Enterprise Scheduler Service

   - Service-Oriented Architecture (SOA)

   For more information about locating log files for Oracle Fusion Applications, see Section 13.6.3. 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 5.15, Section 13.7.4.3, or Section 17.5.3.2.

5. Open the Diagnostic Dashboard and search for diagnostic tests that are associated with the application in which the problem occurred. Run the diagnostic tests that seem most likely to be relevant and inspect the results.

   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 Oracle Enterprise Manager Fusion Applications Control. You can increase the amount of information to be gathered for a particular user or for a whole site.

   - 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.

17.4.2 Viewing Problems and 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, and you can use the Automatic Diagnostic Repository Command Interpreter (`adrci`) to work with incidents.

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 Oracle Enterprise Manager 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.**

To view incidents:

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).

   - If the incident notification referred to an Automatic Diagnostic Repository for Oracle Fusion Middleware, navigate to `MW_HOME/wlserver_versionNumber/server/adr` (for UNIX) or `MW_HOME\wlserver_versionNumber\server\adr` (for Windows).

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
      ```

17.4.3 Recovering from Incidents Generated During SOA Operations

In some cases in which an 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 17.4.2.

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 Oracle Enterprise Manager 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** screen appears.

7. If you have not already done so, use the information on the **Flow Trace** screen to determine how to correct the problem that generated the incident, then repeat steps as needed to return to the **Flow Trace** screen.

   For more information about troubleshooting SOA operations, see Chapter 21.

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 screen 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** screen and repeat Step 8 through Step 10.

17.4.4 Creating and Packaging Incidents Manually

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. For more information, see the "Creating an Incident Manually" section in the *Oracle Fusion Middleware Administrator's Guide.*

After creating the incident, you can manually 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.*

17.4.5 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)

17.4.5.1 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*.

17.4.5.2 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 respository 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 17.4.2 and the chapter about the ADR Command Interpreter in *Oracle Database Utilities*.

17.4.6 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 17.4.4.

   - 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 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 "Working with Diagnostic Dumps" section in the *Oracle Fusion Middleware Administrator's Guide*.

17.4.7 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 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.*

17.5 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 nonproduction environment available for use in testing.

17.5.1 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 17.2.1.

17.5.2 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 17.2.9. 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 17.2.9.

17.5.3 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](#)

17.5.3.1 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.

17.5.3.2 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 5.7.2.

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, then proceed to Step 8.

8. In the **Request Details** pane, select **Request Log** from the Action drop-down list.

   The **Log Messages** screen 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** screen.

   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.

17.5.3.3 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, fchoose **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='0000Hl9TwKUCslT6uBi8UH18lkWX000002')
   ```

17.5.3.4 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.

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='0000Hl9TwKUCslT6uBi8UH18lkWX000002')
     ```

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.*

17.5.4 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 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.*

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.

17.6 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.

17.6.1 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.

17.6.2 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.

17.6.3 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 adding files other than diagnostic test results 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.

17.6.4 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.

18

Troubleshooting Oracle Business Intelligence

This chapter describes common problems that you might encounter when using Oracle Business Intelligence and explains how to solve them.

This chapter includes the following topics:

- Introduction to Troubleshooting Oracle Business Intelligence

- Problems and Solutions for Oracle Business Intelligence Enterprise Edition and Oracle Business Intelligence Publisher

- Diagnosing Oracle Transactional Business Intelligence Query Problems

- Problems and Solutions for Oracle Enterprise Performance Management

- Using My Oracle Support for Additional Troubleshooting Information

In addition to this chapter, review the *Oracle Fusion Middleware Error Messages Reference* for information about the error messages you may encounter.

18.1 Introduction to Troubleshooting Oracle Business Intelligence

This section provides guidelines and a process for using the information in this chapter. Using the following guidelines and process will focus and minimize the time you spend resolving problems.

Guidelines

When using the information in this chapter, Oracle recommends:

- After performing any of the solution procedures in this chapter, immediately retrying the failed task that led you to this troubleshooting information. If the task still fails when you retry it, perform a different solution procedure in this chapter and then try the failed task again. Repeat this process until you resolve the problem.

- Making notes about the solution procedures you perform, symptoms you see, and data you collect while troubleshooting. Examples of information that you should collect include:

  - Screenshots of the error or behavior.

  - The user ID that you used when the error or issue occurred.

  - The steps that you used to perform the task.

> – From where you accessed Oracle BI. For example, embedded analytics, the reporting pane, reports and analytics work area, integrated search, BI Composer, Oracle BI Presentation Services, or Oracle BI Publisher.

If you cannot resolve the problem using the information in this chapter and you must log a service request, the notes you make will expedite the process of solving the problem.

Process

Follow the process outlined in Table 18–1 when using the information in this chapter. If the information in a particular section does not resolve your problem, proceed to the next step in this process.

Table 18–1   Process for Using the Information in this Chapter

Step	Section to Use	Purpose
1	Section 18.3	Perform problem-specific troubleshooting procedures for Oracle Business Intelligence Enterprise Edition and Oracle Business Intelligence Publisher. This section describes:
		■ Possible causes of the problems
		■ Solution procedures corresponding to each of the possible causes
2	Section 18.4	Use this section to diagnose Oracle Transactional Business Intelligence query issues.
3	Section 18.5	Use this section to diagnose Oracle Enterprise Management issues.
3	Section 18.6	Use My Oracle Support to get additional troubleshooting information about Oracle Fusion Applications or Oracle BI. My Oracle Support provides access to several useful troubleshooting resources, including Knowledge Base articles and Community Forums and Discussions.
4	Section 18.6	Log a service request if the information in this chapter and My Oracle Support does not resolve your problem. You can log a service request using My Oracle Support at `https://support.oracle.com`.

18.2 Getting Started with Troubleshooting

Use the following information to determine where to begin troubleshooting your Oracle BI issue.

Oracle BI Technologies

To properly troubleshoot an Oracle BI issue, you must identify in the Oracle Fusion application user interface which underlying Oracle BI technologies are being used. When identifying these technologies, consider the following:

■ The most common analyses are Online Transactional Business Intelligence analyses. Note that Oracle Transaction Business Intelligence can be referred to as real-time or transactional business intelligence. Oracle Transactional Business Intelligence objects analyze very recent data such as today's data or this week's data.

■ Oracle Fusion Financials also uses Financial Reporting technology.

■ Some analyses use Essbase, Real Time Decisions, or Oracle BI Publisher reports.

Oracle BI Query Log

A useful tool to determine the underlying technology for troubleshooting purposes is the Oracle BI query log: `nqquery.log`. From this log you can identify the datasource

and connectivity set-up (for example, transactional database, warehouse database, Financial Reporting datasource, or Essbase datasource) and the datasource objects such as table names and column names.

Warehouse Enablement Setting

When troubleshooting, you should also note whether during the Oracle Fusion Applications setup, the Warehouse Enablement was set to Yes, No, or NA. Use Table 18–2 to determine whether the warehouse source or Oracle Transactional Business Intelligence source services the subject area.

Table 18–2   Warehouse Enablement Settings

Setting	Means that ...
Yes	Real-Time Subject Area is serviced by OTBI source.
	Common Subject Area is services by Warehouse source.
No	Real-Time Subject Area is serviced by OTBI source.
	Common Subject Area is serviced by OTBI source.
	Some measures are not available.
	Some measured return zero.
NA	NA

18.3 Problems and Solutions for Oracle Business Intelligence Enterprise Edition and Oracle Business Intelligence Publisher

This section describes common problems and solutions. It contains the following topics:

- Oracle BI Presentation Services Fails to Start

- The BI Server is Not Running

- Cannot Log Into Oracle Business Intelligence as a User With the BISystem Role

- The "My Account" Link Does Not Display in Presentation Services

- Oracle BI Publisher Reports Are Missing from the Presentation Services Shared Folders

- Oracle BI Publisher Reports Are Missing from the Oracle BI Publisher Server Shared Folders

18.3.1 Oracle BI Presentation Services Fails to Start

In most cases, Presentation Services fails because the BI Server is not running.

Problem

The BI Server may not be running.

Solution

To determine if the BI Server is running:

1. Log into Fusion Applications Control.

2. Open the Business Intelligence node and select **coreapplication**.

3. Click the **Capacity Management** page, and then click the **Availability** tab.

4. Under **BI Server**, locate **coreapplication_obis1** and check its status.

5. If you discover that the BI Server is not running, see Section 18.3.2.

18.3.2 The BI Server is Not Running

For additional information about this issue, see the "Managing Oracle Business Intelligence" and "Diagnosing and Resolving Issues in Oracle Business Intelligence" chapters in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

Problem

The Cluster Controller might not be communicating with the BI Server and the BI Server log might contain information describing why it cannot communicate.

Solution

To check the Cluster Controller log files in Fusion Applications Control:

1. Log into Fusion Applications Control.

2. Open the **Business Intelligence** node and select **coreapplication**.

3. Click the **Diagnostics** tab, and then click the **Log Messages** subtab.

4. Go to the **View/Search Log Files** section and click **Cluster Controller Log.** The Log Messages page displays.

5. Search for an entry similar to the following entry:

```
[2011-02-28T12:54:51.000+00:00] [OracleBIServerComponent] [NOTIFICATION:1] []
[]
[ecid: 004bLSdkE549Xb9LJe_Aif0000Zg000000] [tid: b7f686c0] nqsserver:
Clustered Oracle BI Server started.  Version: 11.1.1.4.0.110109.0239.000.
[2011-02-28T12:54:52.000+00:00] [OracleBIServerComponent] [NOTIFICATION:1]
[] [] [ecid: 004bLSdkE549Xb9LJe_Aif0000Zg000000] [tid: 66f6ba0]  [43071]
A connection with Cluster Controller somemachine.company.com:7001 was
established.
```

Note that if this entry does not exist, then the Cluster Controller is *not* communicating correctly with the BI Server. If the Cluster Controller is not communicating with the BI Server, then go to the following "To check the BI Server log file in Fusion Applications Control" procedure.

To check the BI Server log file in Fusion Applications Control:

1. Log into Fusion Applications Control.

2. Open the **Business Intelligence** node and select **coreapplication**.

3. Click the **Diagnostics** tab, and then click the **Log Messages** subtab.

4. Go to the **View/Search Log Files** section and click **BI Server Log**. The Log Messages page displays.

5. Search for an entry similar to the following entry:

```
[2011-01-31T20:47:38.000-08:00] [OBIPS] [ERROR:10] []
[saw.security.odbcuserpopulationimpl.initialize] [ecid: ] [tid: ]
Odbc driver returned an error (SQLDriverConnectW).
State: HY000.  Code: 10058.
[NQODBC] [SQL_STATE: HY000] [nQSError: 10058]
A general error has occurred.
[nQSError: 73006] Cannot obtain Oracle BI Servers from either
```

```
the primary Cluster Controller (somemachine.company.com)
or the secondary Cluster Controller () specified for the
clustered DSN. (HY000)[[File:odbcuserpoploaderimpl.cpp Line:282
[2011-01-31T20:47:38.000-08:00] [OBIPS] [ERROR:16] []
[saw.security.odbcuserpopulationimpl.initialize] [ecid: ]
[tid: ] Unable to create a system user connection to BI Server during start up.
Trying again.
[[File:odbcuserpoploaderimpl.cpp
Line:283
```

18.3.3 Cannot Log Into Oracle Business Intelligence as a User With the BISystem Role

When logging into Oracle Business Intelligence as a user with the `BISystem` role, a message displays stating that you have entered an invalid user name or password.

Problem 1

A user ID with the `BISystem` role may be locked. The system can lock this user ID if an administrator makes multiple attempts to log in with the incorrect password.

Solution 1

Ask the Oracle Internet Directory administrator to go to the Oracle Internet Directory and check the user ID. If this user ID is locked, you must unlock it.

Problem 2

The password of the user ID with the `BISystem` role may not be synchronized with the credential store.

Solution 2

To confirm that the password is correct:

1. Log into Oracle WebLogic Server Administration Console.

2. Go to **Domain Structure** and click **Security Realms**, and select **myrealm**. The Settings for myrealm page displays.

3. Click the **Users and Groups** tab, and then the **Users** subtab.

4. In the **Users** table, locate and select the user that has the `BISystem` role assigned to it. The settings page displays.

5. Click the **Passwords** tab and reset the user password. Click **Save**.

6. Log into Fusion Applications Control.

7. Open the **WebLogic Domain** node and right-click **BIDomain**.

8. Select **Security**, and then select **Credentials**.

9. Select **oracle.bi.system** and the **select system.user**.

10. Edit the password to match the user password that you set in Oracle WebLogic Server.

11. Restart the Oracle BI Administration Server, Oracle Managed Server, and Oracle Process Manager and Notification Server. For detailed instructions on how to restart these servers, see the "Starting and Stopping Business Intelligence" chapter in the *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

Problem 3

The problem may be that the Oracle Internet Directory password is incorrect.

Solution 3

To check and correct the Oracle Internet Directory password:

1. Log into Fusion Applications Control.

2. Open the **Business Intelligence** node and select **coreapplication**.

3. Click the **Diagnostics** tab, and then click the **Log Messages** subtab.

4. Go to the **View/Search Log Files** section and click **Log Viewer** to search all log files. The Log Messages page displays.

5. Go to the **Search** section and expand the **Selected Targets** section.

6. In the **Selected Targets** table, locate **AdminServer** and click the corresponding **View** list of target files button.

7. Search for Oracle Internet Directory errors. If the error is due to the wrong password, you will find an error such as the following error:

```
<Feb 15, 2011 12:31:10 PM PST> <Error>  <Console> <BEA-240003> <Console
  encountered the following error weblogic.security.providers.authentication.
  LDAPAtnDelegateException: [Security:090294]could not get connection

    atweblogic.security.providers.authentication.
      LDAPAtnDelegate.getConnection(LDAPAtnDelegate.java:3483)

    at weblogic.security.providers.authentication.LDAPAtnDelegate.
      getConnection(LDAPAtnDelegate.java:3470)

    at weblogic.security.providers.authentication.LDAPAtnDelegate.
      listUsers(LDAPAtnDelegate.java:2258)

    at weblogic.security.providers.authentication.LDAPAuthenticatorImpl.
      listUsers(LDAPAuthenticatorImpl.java:178)

    at weblogic.security.providers.authentication.
      OracleInternetDirectoryAuthenticatorMBeanImpl.listUsers

    ... 110 more

Caused by: netscape.ldap.LDAPException: error result (49); Invalid
    credentials
```

8. To correct this issue, log into Oracle WebLogic Server Administration Console.

9. Go to **Domain Structure**, click **Security Realms**, and select **myrealm**.

10. Select the **Users and Groups** tab, and then select the **Users** subtab.

11. Select the user ID. The settings for the user ID display.

12. Click the **Passwords** subtab and reset the user's password to the correct password. If you are not sure of the password, see the Oracle Internet Directory administrator.

18.3.4 The "My Account" Link Does Not Display in Presentation Services

In the Global Header of Presentation Services, the My Account option is missing. Normally the user can select the Signed In As list and the My Account option displays.

Problem

The problem may be a global user ID (GUID) refresh issue that requires you to recover the catalog.

> **Note:** For more information, see "Regenerating User GUIDs" in *Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition*.

Solution

To recover the catalog:

1. Start a command prompt.

2. Enter the following commands:

   ```
   cd<Oracle_bi_Instance>/config/OracleBIPresentationServicesComponent/
       coreapplication_obips1/
   ```

3. Open the `instanceconfig.xml` file.

   > **Note:** Before you modify the `instanceconfig.xml`, it is critical that you understand how to properly modify this file. For more information, see the "Configuring the Oracle Business Intelligence System" chapter in the *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

4. Locate the `<catalog>` element and add the following:

   ```
   <UpdateAccountGUIDS>Recover</UpdateAccountGUIDs>
   ```

5. Save `instanceconfig.xml`.

6. Stop and restart Presentation Services.

7. Log into Fusion Applications Control.

8. Open the **Business Intelligence** node and select **coreapplication**.

9. Select the **Capacity Management** tab, and then select the **Availability** subtab.

10. Under **BI Presentation Services**, select **coreapplication_obips1**.

11. Click the **Stop Selected** button and confirm that **coreapplicataion_obips1** stopped.

12. Re-select **coreapplication_obips1**.

13. Click the **Start Selected** button and confirm that **coreapplication_obips1** started.

14. Open `instanceconfig.xml`.

15. Remove the following:

    ```
    <UpdateAccountGUIDS>Recover</UpdateAccountGUIDs>
    ```

    > **Note:** You must perform this step to ensure that your system is secure.

16. Save `instanceconfig.xml`.

17. Stop and restart the BI Server.

18.3.5 Oracle BI Publisher Reports Are Missing from the Presentation Services Shared Folders

When users log into Presentation Services and access the Oracle BI Presentation Catalog to browse the shared folders for their reports, the shared folders do not contain any reports.

Problem

The reports might not exist on the file system.

Solution

To confirm that the items exist on the file system:

1. Log into Fusion Applications Control.

2. Open the **Business Intelligence** node and select **coreapplication**.

3. Click the **Deployment** tab, and then click the **Repository** subtab.

4. Go to the **BI Presentation Catalog** section of the page and note the catalog path.

5. Use OS commands or a file explorer to navigate to the catalog path.

6. Browse the catalog for reports. Depending upon the outcome of this step, you may do one of the following:

- If the reports do not exist in the catalog, it is probably because of a GUID refresh issue. See the following "To refresh the GUIDs:" procedure.

- If the reports are not in the catalog, then it may be that you recently completed an unsuccessful catalog merge. Note that even if the last statement of the merge indicated success, the merge still might not have completed successfully. If you suspect that your catalog merge was not successful, then repeat the catalog merge.

To refresh the GUIDs:

Note: For more information, see "Regenerating User GUIDs" in *Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition*.

---

1. Start a command prompt.

2. Enter the following commands:

```
cd <Oracle_bi_Instance>/config/OracleBIServerComponent/coreapplication_obis1/

edit NQSconfig.ini .. set FMW_UPDATE_ROLE_AND_USER_REF_GUIDS=YES

./opmnctl stopproc ias-component=coreapplication_obis1

./opmnctl startproc ias-component=coreapplication_obis1
```

3. Repeat the previous step and set the `FMW_UPDATE_ROLE_AND_USER_REF_GUIDS` parameter in `NQSConfig.INI` back to NO.

> **Note:** You must perform this step to ensure that your system is secure.

4. Restart the BI Server. For detailed instructions on how to perform this task, see the "Starting and Stopping Business Intelligence" chapter in the *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

18.3.6 Oracle BI Publisher Reports Are Missing from the Oracle BI Publisher Server Shared Folders

When users log into Oracle BI Publisher and access the shared folders containing their reports, they see no reports. This problem can occur even though users can locate their reports in the Presentation Services Catalog's shared folders.

Problem 1

The user may not have the proper privileges to view the report.

Solution 1

To access Presentation Services and check the privileges required to access reports:

1. Log into Presentation Services as a user with administrator privileges.

2. Go to the **Global Header** and click the **Signed In As** list.

3. Click **My Account**. The My Account page displays.

4. Click the **Roles and Catalog Groups** tab and note the roles and groups in this page. Click **OK** to exit this page.

5. Go to the **Global Header** and click **Sign Out** to log out of Presentation Services.

6. Log back into BI Presentation Services using the log in credentials of the user who cannot view the reports in the Oracle BI Publisher Server shared folders.

7. Access the user's privileges using the previous steps.

8. If the user does not have the proper privileges, log into Presentation Services as a user with administrator privileges and navigate to the Manage Privileges page to assign the proper privileges to the user.

   For more information, see the "Setting Oracle BI Presentation Catalog Privileges for an Application Role" section in *Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition*.

Problem 2

An error may have occurred when Oracle BI Publisher connected to Presentation Services.

Solution 2

To override the `BISystemUser` account:

By default Oracle BI Publisher connects to Presentation Services using the `BISystemUser` account that is stored in the credential store. However, it is possible for users to override this.

1. Open the `xmlp-server-config.xml` file in the following directories:

```
(UNIX) BI_DOMAIN_HOME/config/bipublisher/repository/Admin/Configuration
(Windows) FA_MW_HOME\config\bipublisher\repository\Admin\Configuration
```

2. Confirm that the SAW_USERNAME and SAW_PASSWORD properties are not set. Note that if these properties have been set, it will cause an error. Before you change these properties, ask your administrator why they have been set.

3. To un-set these properties, set their values to "".

4. Restart the WebLogic domain. For more information, see Section 3.3.4.

Problem 3

The catalog may be corrupted.

Solution 3

To check if the catalog is corrupted:

1. Log into Presentation Services as a user with administrator privileges.

2. Go to the **Global Header** and click **Catalog**. The Catalog page displays.

3. In the **Search** pane, enter the following search criteria:

   ■ In the **Name** field, type *.

   ■ In the **Location** field, click the list and choose **All**.

   ■ In the **Type** field, click the list and choose **Data Model**.

4. Click **Search**. If the search results include the **Shared Folder**, then the catalog is corrupt. Note that if the Shared Folder is included in your search results, Oracle BI Enterprise Edition has classified it as a data model rather than a folder. Perform the following "To fix the corrupted catalog" procedure to fix the catalog.

To fix the corrupted catalog:

1. Open Catalog Manager and open the catalog in offline mode.

2. Navigate to **Shared Folders**.

3. Right-click the object in the Name column and select Properties.

4. Set the following properties to null:

   ■ **bip:DisplayName**

   ■ **Caption**

   ■ **compositeSignature**

   ■ **DESCRIPTION**

   ■ **objectName**

5. Save your changes.

18.4 Diagnosing Oracle Transactional Business Intelligence Query Problems

This section describes how you to use the NQQuery.log file and the AdminServer-diagnostic.log file to diagnose Oracle Transactional Business Intelligence query issues such as ODBC errors, ORA-errors, other database errors, and query performance issues.

The NQQuery.log file is useful for any analysis that uses the BI Server to query the datasource (for example, OTBI or OBIA). The NQQuery.log records the logical SQL sent by the Oracle BI Presentation layer to the BI Server layer and the physical or ANSI SQL sent by the BI Server to the datasource.

This section contains the following topics:

- Enabling the NQQuery.log File
- Presentation Services Request Query
- Query Trace
- Physical Query
- Using AdminServer-diagnostic.log File to Troubleshoot Oracle WebLogic Server
- Using View Data to Check the Datasource Connectivity

18.4.1 Enabling the NQQuery.log File

You can enable the NQQuery.log file to diagnose issues that you encounter in Oracle Transactional Business Intelligence queries. Use one of the following methods to enable the NQQuery.log file.

- You can enable query logging from the Oracle BI Administration Tool. For information and procedures about enabling the query log file, see "Managing the Query Log" in the *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

- You can enable the LogLevel system session variable for the repository by accessing the repository and setting the LogLevel system variable to 7. For information and procedures about setting this variable, see "About System Session Variables" section in the *Oracle Fusion Middleware Metadata Repository Builder's Guide for Oracle Business Intelligence Enterprise Edition (Oracle Fusion Applications Edition)*.

- If you have administrator permissions, you can log into Presentation Services, access the analysis that you want to troubleshoot, and temporarily change its logging level to 2. This will turn on logging for the analysis even when logging is turned off for all analyses. For more information, see "Setting the Query Logging Level" in *Oracle Fusion Middleware System Administrator's Guide for Oracle Business Intelligence Enterprise Edition*.

  If you use this method to enable the NQQuery.log file, look for the following:

  - An error like "no log found." An error like this indicates that there is a known error. Look in the NQQuery.log file, which contains additional logging.
  - The timestamp to help you locate the log entry.
  - The logical SQL as shown in the error message. This is the SQL sent by Presentation Services to the BI Server layer.
  - The physical SQL related to the above mentioned logical SQL. Search for "physical". The physical SQL is the SQL sent by BI Server through ODBC to the database.
  - For any more detailed errors.

Note that the NQQuery.log file is *not* searchable in the Fusion Applications Control Log Viewer.

After enabling query logging, if the BI Server caching is enabled, you must clear the cache before re-executing the analysis so that the NQQuery.log file will contain the correct information. Use the following procedure to perform this task.

To clear the BI Server cache:

1. Log into Presentation Services as an administrator.

2. In the **Global Header**, click **Administration**. The Administration page displays.

3. Click **Manage Sessions**. The Manage Sessions page displays.

4. Scroll to the Cursor Cache section of the page and click the **Close All Cursors** button.

18.4.2 Presentation Services Request Query

The request query executed from Presentation Services will be shown in the NQQuery.log file as follows:

```
RqList     0 as c1 GB,    Dim - Procurement Item.Category Name as c2 GB,     Dim
 - Procurement Item.Item Description as c3 GB,    # of PO Lines:[DAggr(Fact -
 Purchasing - Order.# of PO Lines by [ Dim - Procurement Item.Category Name, Dim
 - Procurement Item.Item Description] )] as c4 GBOrderBy: c1 asc, c2 asc, c3 asc
```

This query trace explains how a measure is calculated. It also shows the logical table source used to render the data. The query should use the logical table sources having priority 5 for the requests executed from the Oracle Transactional Business Intelligence Real Time Subject Areas.

18.4.3 Query Trace

The query trace shows the view object, its view criteria, and its view links that are used to execute the request. See the following:

```
<ADFQuery mode="SQLBypass" queryid="14604-3902"
locale="en"><Parameters></Parameters>
<Projection>
<Attribute><Name><![CDATA[CategoryName]]></Name>
    <ViewObject><![CDATA[FscmTopModelAM.PrcPoPublicViewAM.
    PurchasingItemP]]></ViewObject>
</Attribute>
<Attribute><Name><![CDATA[ItemDescription]]></Name>
    <ViewObject><![CDATA[FscmTopModelAM.PrcPoPublicViewAM.
    PurchasingItemP]]></ViewObject>
</Attribute>
<Attribute><Name><![CDATA[PurchasingDocumentHeaderTypeLookupCode]]></Name>
    <ViewObject><![CDATA[FscmTopModelAM.PrcPoPublicViewAM.StandardLinePVO]]>
    </ViewObject>
</Attribute>
<Attribute><Name><![CDATA[PoLineId]]></Name>
    <ViewObject><![CDATA[FscmTopModelAM.PrcPoPublicViewAM.StandardLinePVO]>]
    </ViewObject>
</Attribute>
<Attribute><Name><![CDATA[ItemNumber]]></Name>
    <ViewObject><![CDATA[FscmTopModelAM.PrcPoPublicViewAM.PurchasingItemP]]>
    </ViewObject>
</Attribute>
</Projection>
JoinSpec>
    <ViewObject>
```

```
        <Name><![CDATA[FscmTopModelAM.PrcPoPublicViewAM.PurchasingItemP]]></Name>
        <ViewLink><Name><![CDATA[oracle.apps.prc.po.publicView.analytics.link.
          PurchasingDocumentLinePVOToPurchasingItemPVO]]></Name></ViewLink>
    </ViewObject>
    <ViewObject>
        <Name><![CDATA[FscmTopModelAM.PrcPoPublicViewAM.StandardLinePVO]]></Name>
        <ViewLink><Name><![CDATA[oracle.apps.prc.po.publicView.analytics.link.
          PurchasingDocumentLinePVOToPurchasingItemPVO]]></Name></ViewLink>
    </ViewObject>
</JoinSpec>
<DetailFilter>
<ViewCriteria>
  <ViewCriteriaRow conjunction="VC_CONJ_AND">
  <Attribute><Name><![CDATA[PurchasingDocumentVersionCoSequence]]></Name>
    <ViewObject><![CDATA[FscmTopModelAM.PrcPoPublicViewAM.StandardLinePVO]]>
    </ViewObject>
  </Attribute>
      <Value><![CDATA[0]]></Value>
      <Predicate operator="OPER_EQ"/>
  </ViewCriteriaRow>
  <ViewCriteriaRow conjunction="VC_CONJ_AND" negated="true">
  <Attribute><Name><![CDATA[PurchasingDocumentVersionCoCanceledFlag]]></Name>
    <ViewObject><![CDATA[FscmTopModelAM.PrcPoPublicViewAM.StandardLinePVO]]>
    </ViewObject>
  </Attribute>
  <Value><![CDATA[Y]]></Value>
  <Predicate operator="OPER_EQ"/>
  </ViewCriteriaRow>
</ViewCriteria>
</DetailFilter>
</ADFQuery>
```

You or the administrator can review the query trace to check that the desired View Objects (VOs) and View Links are executed. The execution of the view objects and view must be based on the Logical model code to render the measure or attributes executed in the request.

18.4.4 Physical Query

The SQL Bypass database should be setup in the repository to send the physical query to the database directly, rather than through the ADF Server. The NQQuery.log file shows the physical query sent to the database if the SQL Bypass database is enabled. See the following excerpt from the log file:

```
WITH
SAWITH0 AS (select T744257.C501728333 as c1,
   T744257.C281594243 as c2,
   T744257.C469784899 as c3,
   T744257.C168071223 as c4
from
   (SELECT V110617254.CATEGORY_NAME AS C501728333,       V110617254.
    ITEM_DESCRIPTION
 AS C281594243,          V278099157.TYPE_LOOKUP_CODE263 AS C469784899,
V278099157.PO_LINE_ID AS C168071223,          V110617254.ITEM_NUMBER AS
C348883104,          V278099157.CO_SEQUENCE AS C78750419,          V278099157.CO_
CANCELED_FLAG AS C443371219,          V278099157.PO_HEADER_ID1 AS PKA_
PurchasingDocumentHeaderP0,          V278099157.VERSION_ID AS PKA_
PurchasingDocumentVersion0 FROM (SELECT PurchasingDocumentLine.PO_LINE_ID,
PurchasingDocumentHeader.PO_HEADER_ID AS PO_HEADER_ID1,
PurchasingDocumentHeader.TYPE_LOOKUP_CODE AS TYPE_LOOKUP_CODE263,
```

```
PurchasingDocumentVersion.CO_CANCELED_FLAG,
PurchasingDocumentVersion.CO_SEQUENCE,
PurchasingDocumentVersion.VERSION_ID,
(DECODE(PurchasingDocumentLine.ITEM_ID, NULL,
DECODE(PurchasingDocumentLine.VENDOR_PRODUCT_NUM, NULL,
(PurchasingDocumentLine.ITEM_DESCRIPTION || '[' || CategoryTranslation.
CATEGORY_NAME || ']'),(PurchasingDocumentLine.VENDOR_PRODUCT_NUM || '[' ||
PurchasingDocumentHeader.VENDOR_ID || ']')),TO_CHAR(PurchasingDocumentLine.
ITEM_ID))) AS ITEM_NAME FROM PO_LINES_ALL PurchasingDocumentLine,
PO_HEADERS_ALL PurchasingDocumentHeader, PO_VERSIONS PurchasingDocumentVersion
 WHERE (PurchasingDocumentLine.PO_HEADER_ID = PurchasingDocumentHeader.
PO_HEADER_ID AND PurchasingDocumentHeader.PO_HEADER_ID =
PurchasingDocumentVersion.PO_HEADER_ID) AND ( (
(UPPER(PurchasingDocumentHeader.TYPE_LOOKUP_CODE) = UPPER('STANDARD') ) ) ))
 V278099157, (SELECT DISTINCT ( DECODE(PL.ITEM_ID,NULL,DECODE(PL.
VENDOR_PRODUCT_NUM, NULL,(PL.ITEM_DESCRIPTION || '[' || TL1.CATEGORY
_NAME||']'),(PL.VENDOR_PRODUCT_NUM || '[' || PH.VENDOR_ID ||']')), TO
_CHAR(I.ITEM_NUMBER))) AS ITEM_NUMBER,  PL.ITEM_DESCRIPTION,
PL.CATEGORY_ID,TL1.CATEGORY_NAME, DECODE(PL.ITEM_ID, NULL,DECODE
(PL.VENDOR_PRODUCT_NUM, NULL, 'DESCRIPTION BASED ITEMS','SUPPLIER
 ITEMS'),'INVENTORY_ITEMS') AS ITEMTYPE  FROM PO_LINES_ALL PL,
EGP_CATEGORIES_TL TL1, PO_HEADERS_ALL PH, PO_SYSTEM_PARAMETERS_ALL SP,
EGP_SYSTEM_ITEMS_B I WHERE PL.CATEGORY_ID=TL1.CATEGORY_ID AND PL.CATEGORY_ID IS
NOT NULL AND PL.PO_HEADER_ID=PH.PO_HEADER_ID AND PL.PRC_BU_ID=SP.PRC_BU_ID AND
I.INVENTORY_ITEM_ID=PL.ITEM_ID) V110617254 WHERE V278099157.ITEM_NAME =
 V110617254.ITEM_NUMBER( + ) AND ( ( (V278099157.CO_SEQUENCE = 0 ) )  AND
( NOT ( (V278099157.CO_CANCELED_FLAG = 'Y' ) ) ) ) )) T744257),
SAWITH1 AS (select D1.c1 as c2,    D1.c2 as c3,
   case when D1.c3 = 'STANDARD' then D1.c4 end  as c4
from
   SAWITH0 D1), SAWITH2 AS (select D1.c2 as c2,    D1.c3 as c3,
   D1.c4 as c4,
   ROW_NUMBER() OVER (PARTITION BY D1.c2, D1.c3, D1.c4 ORDER BY D1.c2 DESC,
   D1.c3 DESC, D1.c4 DESC) as c5from      SAWITH1 D1),SAWITH3 AS (select
   count(distinct case D1.c5 when 1 then D1.c4 else NULL end )
   as c1,    D1.c2 as c2,    D1.c3 as c3from SAWITH2 D1
group by D1.c2, D1.c3)
select distinct 0 as c1,    D1.c2 as c2,    D1.c3 as c3,    D1.c1 as c4 from
   SAWITH3 D1
order by c2, c3
```

Use this physical SQL to diagnose which tables, columns joins, and filters are being used by the BI Server to gather data. Table names are aliased (for example, T744257). Search for the table alias in the RPD to find the actual table name. Run the SQL to check whether it works in SQL*Plus.

For performance issues, look for filter columns that are not indexed, and known database performance causes such as NOT IN clauses. Also use Oracle Database Control for performance advice.

The NQQuery.log file also indicates the user who executed the analysis, the timestamp, and the OBI Connection Object used for the analysis. From the Connection Object name, you can refer back to the Oracle BI repository to find the set up properties of the connection, such as whether it uses a native connection or ODBC.

The Oracle BI NQQuery.log file contains detailed BI Server errors, so it is recommended to check this log file. You can access this log file with the Fusion Application Control. The NQQuery.log file shows queries executed at the time of starting the BI Server service. Be sure to check this log for any query failures.

18.4.5 Using AdminServer-diagnostic.log File to Troubleshoot Oracle WebLogic Server

Oracle BI EE gets the view object physical query from the Oracle ADF Server. This query involves the view object query and the security predicate associated to it. You or the administrator need to research the WebLogic Server log file if you suspect an issue in the view object query. You will find this information in the `AdminServer-diagnostic.log` file. This file is located in the WebLogic Server domain associated to Oracle JDeveloper 11*g*.

The `AdminServer-diagnostic.log` file shows the user name, view object, attributes, the view link source and destination entities, and the view criteria for view objects (including security view criteria). In this file you will find the roles associated to the user and the security predicate associated to the roles and users in the roles. This file also shows the composite view object API called for the view object, along with view links and view criteria.

18.4.6 Using View Data to Check the Datasource Connectivity

The BI Server uses the properties in the connection pool object, which is located in the Oracle BI repository. Use the following procedure to confirm that your connectivity is set up correctly.

To confirm that your connectivity is set up correctly:

> **Note:** This procedure uses the ADF BC Datasource as an example.

1. In the Physical layer of the Administration Tool, expand the database object for the ADF Business Component data source.

2. Right-click a physical table and click **View Data**.

3. Check that the appropriate data displays. Note that if you have just imported, you may need to check in the new objects before you perform this test.

18.5 Problems and Solutions for Oracle Enterprise Performance Management

This section describes common problems and solutions for Oracle Enterprise Performance Management. It contains the following topics:

- Foundation Services
- Financial Reporting
- Hyperion Provider Services
- Essbase

18.5.1 Foundation Services

This section describes common problems and solutions for Foundation Services. It contains the following topics:

- Getting Started with Logging Basics for Foundation Services
- Web Traffic Snooping
- EPM Registry

- Hyperion Security Username/Password Authentication Fails during Enterprise Scheduler Services Essbase Cube Creation in the Domain

- EPM Workspace

- Allocation Manager

- Smart View

18.5.1.1 Getting Started with Logging Basics for Foundation Services

Log files are the best tools for analyzing what might be wrong with the system configuration. The logging configuration file, the defaults it ships with, and instructions for changing change those defaults will help with the analysis.

- Configuration: The `logging.xml` file in the following directories contains the loggers and their default levels:

  ```
  (UNIX) DOMAIN_HOME/servers/server_name/logs
  (Windows) DOMAIN_HOME\servers\server_name\logs
  ```

  The *DOMAIN_HOME* is named `BIDomain`.

- Output: The default locations of the loggers are rooted in the following directories:

  ```
  (UNIX) DOMAIN_HOME/servers/server_name/logs
  (Windows) DOMAIN_HOME\servers\server_name\logs
  ```

  These files include:

  - `/registry/registry.log`

  - `/css/css.log`

  - `/workspace/Framework.log`

  - `/workspace/workspace.log`

  - `/financialreporting/fr.log`

  - `CalcManager.log`

- Loggers: Increase or decrease the levels of these loggers:

  - **oracle.EPMCSS**

  - **oracle.bi.bifndnepm.epmreg**

  - **oracle.bi.bifndnepm.bpmui**

  - **oracle.bi.bifndnepm.workspace**

  - **oracle.EPMFR**

  - **oracle.EPMAnnotations**

  - **oracle.EPMJCR**

  - **oracle.EPMADM**

  - **oracle.bi.bifndnepm.calcmgr**

  To change logger levels, perform the following steps:

  a. Go to Fusion Applications Control.

  b. From the navigation pane, expand the farm and then **WebLogic Domain**.

  c. Right-click a Managed Server from within the domain (each server's log levels can be independently set).

    d. Choose **Logs** > **Log Configuration**.

    e. In the **Logger Name** column, expand the **oracle** runtime loggers to display loggers.

    f. Change the logging level as appropriate.

18.5.1.2 Web Traffic Snooping

Problem

Some components are not operating correctly, or appear to have some minor issues. This is a broad category, but snooping the web traffic can help development resolve some issues by looking for irregular status, content, and redirection, and just providing the flow of calls up to the problem area.

Solution

1. Connect to `http://www.fiddler2.com/fiddler2` and install version 2 (beta is stable). Launch Fiddler through one of these methods:

   - Start Menu: **Programs** > **Fiddler2**

   - Microsoft Internet Explorer: **Tools** > **Fiddler2**

   - Mozilla Firefox: **Tools** > **Monitor With Fiddler** > **Launch Fiddler Now**

     Ensure **Do not use Fiddler** is not selected.

2. Connect to the starting URL and continue normally until the problem occurs.

3. When the problem occurs, in Fiddler, choose **File** > **Save** > **All Sessions**.

4. Enter a name for the session archive (`.saz`) file.

> **Note:** Be aware that the file may contain sensitive information, such as user ID and password information. Therefore, manage the file carefully.

18.5.1.3 EPM Registry

This section describes common problems and solutions for EPM Registry. It contains the following topics:

- EPM Registry Initialization Fails

- Connection to Essbase Server Fails

- Host and Ports Do Not Match

18.5.1.3.1 EPM Registry Initialization Fails

Problem

The EPM Registry initialization fails in the `BIDomain` domain. Oracle Fusion General Ledger and Oracle Fusion Customer Relationship Management (CRM) code running within the `FinancialsDomain` and `CRMDomain` domains, respectively, has direct integration with the EPM Registry API and consumes EPM Security Component, which also tightly integrates with the EPM Registry API. When EPM Registry initialization fails, it greatly impacts communication with downstream EPM products in the BI domain.

Solution

To resolve this problem:

1. Login to Oracle WebLogic Server Console as the Oracle WebLogic Server administrative user.

2. From **Domain Structure** portlet, navigate to **Services** > **Data Sources** and select **EPMSystemRegistry Datasource** from the available list of datasources.

3. Click on the **Connection Pool** tab and confirm that the correct information is available for URL property.

4. Click on the **Targets** tab and confirm that the data source is targeted to all of the required Oracle WebLogic Servers and clusters running in the `Financials` domain.

5. Click on the **Monitoring** tab and confirm that the data source for all of the targeted Oracle WebLogic Servers and clusters is in **Running** state. Optionally, try clicking **Test Data Source** button to test the data source for a given server.

6. If, in Step 5, the data source was not in the **Running** state, click the **Control** tab and attempt to start the data source using the **Start** button.

7. If any errors were encountered during Step 5 or 6, consider restarting the Oracle WebLogic Server domain.

18.5.1.3.2   Connection to Essbase Server Fails

Problem

An error is returned when connecting to Essbase Server from the Fusion Applications domain, using the Essbase Cluster lookup URL that is built using the host and port information retrieved for `Essbase_FA_Cluster`, from EPM Registry using the Registry API.

Solution

To resolve this issue, perform the following validations with appropriate actions:

1. Verify that the host and port values for the `LOGICAL_WEB_APP` component in EPM System Registry, where `webAppType` is `PROVIDER_SERVICES_WEB_APP`, exactly matches the host and port that the APS Web application is running on in the BI domain. Follow these steps:

    a. Run the EPM Registry Editor utility using the following command:

    ```
    (UNIX) APPLICATIONS_CONFIG_
    HOME/BIInstance/config/foundation/11.1.2.0/epmsys_registry.sh view LOGICAL_
    WEB_APP | tee -a logical_webapp_report.txt
    (Windows) APPLICATIONS_CONFIG_
    HOME\BIInstance\config\foundation\11.1.2.0\epmsys_registry.bat view
    LOGICAL_WEB_APP > logical_webapp_report.txt
    ```

    b. In the `logical_webapp_report.txt` file that is generated as a result of Step 1, search for the following string:

    ```
    *"webAppType = PROVIDER_SERVICES_WEB_APP"
    ```

    c. For the `COMPONENT` containing the matching string, verify that `HOST/port/SSL_PORT` exactly match what the APS Web application is actually running on in the BI domain. If any of these values are different in the Essbase cluster lookup URL, they should be appropriately updated in EPM

Registry for the given Logical Webapp Component. This solution is discussed in Section 18.5.1.3.3.

d. Confirm that the APS Web application is in an Active state in the BI domain, by logging into the BI Admin Server console and reviewing the deployment profile of APS.

2. If verification of Step 1 successfully passes, verify the Oracle Access Manager protection policies and confirm that the Essbase cluster lookup URL is excluded from those policies

18.5.1.3.3 Host and Ports Do Not Match

Problem

If the host or port for Logical Web App components in the EPM Registry do not exactly match the actual host and port that EPM Web applications are running on in the `BIDomain` domain, there may be connection or launch issues for EPM Web applications from within the Fusion user interface. In this case, you must update the host and port for Logical Web App components in EPM Registry.

Solution

To resolve the issue:

1. Execute the following command so that respective component IDs are available in the file for later use.

```
(UNIX) APPLICATIONS_CONFIG_HOME/BIInstance/config/foundation/11.1.2.0/epmsys_
registry.sh view LOGICAL_WEB_APP | tee -a logical_webapp_report.txt
(Windows) APPLICATIONS_CONFIG_
HOME\BIInstance\config\foundation\11.1.2.0\epmsys_registry.bat view LOGICAL_
WEB_APP > logical_webapp_report.txt
```

The EPM Registry Dump (`.html` report) that is created by running `epmsys_registry.sh` and `epmsys_registry.bat` does not include the component IDs that are required to update the host and port properties for individual components.

After successfully executing this command, the `logical_webapp_report.txt` file now contains exported data for five Logical Web App components - something similar to:

```
COMPONENT - 1
 NAME -  Default
 ID -  a01b453873d1e7b5S7b70c01f12e34faed1bS7fdc
 TYPE -  LOGICAL_WEB_APP
 HOST -  hostname
 HYPERION HOME -  /scratch/aime1/work/mw8747/Oracle_BI1
 PROPERTIES -
         webAppType = CALC_WEBAPP
        context = calcmgr
```

The ID highlighted above in bold will be available for COMPONENT - [1-5]; this is what you will be required to use in subsequent steps.

2. For the individual ID of COMPONENT - [1-5], execute the following EPM Registry update command:

```
(UNIX) APPLICATIONS_CONFIG_HOME/BIInstance/config/foundation/11.1.2.0/epmsys_
registry.sh updateproperty \#ID/@port port
(Windows) APPLICATIONS_CONFIG_
```

```
HOME\BIInstance\config\foundation\11.1.2.0\epmsys_registry.bat updateproperty
#ID/@port port
```

For example:

```
./epmsys_registry.sh updateproperty
\#a01b453873d1e7b5S7b70c01f12e34faed1bS7fdc/@port 16050
```

3. Validate that you have successfully updated the host and port for Logical Web App components:

```
(UNIX)  APPLICATIONS_CONFIG_HOME/BIInstance/config/foundation/11.1.2.0/epmsys_
registry.sh view LOGICAL_WEB_APP | tee -a logical_webapp_update_report.txt
(Windows)  APPLICATIONS_CONFIG_
HOME\BIInstance\config\foundation\11.1.2.0\epmsys_registry.bat view LOGICAL_
WEB_APP > logical_webapp_report.tx
```

4. After these updates, restart the `BIDomain`, `CRMDomain` and `FinancialsDomain` domains, so that new changes take effect. See Section 3.3.4

18.5.1.4 Hyperion Security Username/Password Authentication Fails during Enterprise Scheduler Services Essbase Cube Creation in the Domain

Problem

The following error about Oracle Enterprise Scheduler Essbase Cube creation in the `BIDomain` domain is returned because of a Hyperion security username and password authentication failure:

```
EPMCSS-00301: Failed to authenticate user. Invalid credentials. Enter valid
credentials.
```

Solution

To resolve this issue, verify that the password stored in credential store for Application Identity (App ID) used by the Oracle Enterprise Scheduler job to connect to Essbase exactly matches the password for the given Application Identity in the underlying LDAP Identity Store:

1. Identify which App ID is used by the Oracle Enterprise Scheduler job in the given domain, to connect to an Essbase single cluster.

The following application identities are used for the respective domains:

- Oracle Fusion Customer Relationship Management: `FUSION_APPS_CRM_ADF_APPID`

- Oracle Fusion Projects: `FUSION_APPS_PRJ_ESS_APPID`

- Oracle Fusion General Ledger: `FUSION_APPS_GL_ESSBASE_APPID`

2. Run the WLST `listCred` script with the appropriate map and key. See the section "listCred" in the *Oracle Fusion Middleware Application Security Guide*.

3. The password obtained in Step 2 can then be used to compare the existing password for the given App ID in the underlying LDAP ID Store, using the following command line `ldapcompare`:

```
ldapcompare -h hostname -p port -D "bind_user_dn" -w bind_user_password -a
userpassword -v appid_password_from_credstore -b "appid_dn"
```

If Step 3 does not run with positive results, it implies that the credential store password and the LDAP store password failed the comparison test. In this case, passwords in credential store and LDAP ID Store should be synced.

18.5.1.5 EPM Workspace

This section describes common problems and solutions for EPM Workspace. It contains the following topics:

- Debugging

- Startup

- Microsoft Internet Explorer Truncated Menus

18.5.1.5.1 Debugging

Problem

Errors occur with certain functionality.

Solution

1. Check the status of Workspace by entering this URL:

   ```
   http://server_name:19000/workspace/status
   ```

   This URL displays a summary of any fatal errors that prevent Workspace from running, and a list of integrated products.

2. Enable the client debugging feature by selecting **Navigate** > **Administer** > **Workspace Server Settings** and selecting **Client Debug Enabled**.

   All users who log on after selecting this option will be affected. The URLs that can be accessed include the following:

   ```
   http://server_name:19000/workspace/debug/configInfo.jsp
   http://server_name:19000/workspace/debug/userInfo.jsp
   ```

   The first URL shows information about the system including metadata from integrated products (for example, menus, preferences panels, and so on); the second URL shows information about the current user, including groups, effective roles, and Workspace menu items as interpreted and filtered for the current user.

18.5.1.5.2 Startup

Problem

The `http://server_name:19000/workspace/` URL fails to respond or an error occurs.

Solution

Workspace initialization happens at the first request, not at server startup. If initialization fails, it is reattempted at every subsequent request. Any error during initialization is logged in the `workspace.log` file. The error message displays a page that replaces the normal Workspace Log on screen. Finally, the error is reported in the status report servlet, `/workspace/status`. Only the log files include Exception stack traces. The following tasks must successfully complete in order before Workspace allows logons, and can be matched up in the `workspace.log` file:

1. Parse and validate the file `/conf/WSProducts.xml` in the Workspace Web application.

2. Initialize the Shared Services Registry:

   a. Connect to the Shared Services Registry.

   b. Find the unique `WORKSPACE` component in the Shared Services Registry.

   c. Find the unique Workspace `LOGICAL_WEB_APP` component in the Shared Services Registry.

   d. View the Workspace configuration properties from `LOGICAL_WEB_APP`.

   e. Find at least one `WEB_SERVER` component that is a child of `WORKSPACE`.

3. Update `LOGICAL_WEB_APP` with host and port information from the Oracle WebLogic Server MBean.

4. Update `LOGICAL_WEB_APP` with any missing configuration properties.

5. Initialize Shared Services.

6. Scan the Shared Services Registry for integrated products, and match every discovered product with a configuration file from the following resources:

- `/conf/WSProducts.xml` (Workspace)

- `WorkspaceConfig` file attribute of `LOGICAL_WEB_APP` (all other products)

18.5.1.5.3 Microsoft Internet Explorer Truncated Menus

Problem

Menu lists appear truncated in either height or width.

Solution

To resolve this issue:

1. From Microsoft Internet Explorer, choose **Tools > Internet Options > Security > CustomLevel > Miscellaneous**.

2. For **Allow script-initiated windows without size or position constraints**, select **Enable**.

18.5.1.6 Allocation Manager

This section describes common problems and solutions for Allocation Manager. It contains the following topics:

- Essbase Applications Are Not Listed Under the Essbase Node

- Business Rules Cannot be Deployed to General Ledger

18.5.1.6.1 Essbase Applications Are Not Listed Under the Essbase Node

Problem

No Essbase applications are listed under the Essbase node.

Solution

This could be due to any of the following:

- You are not provisioned to work with Allocation Manager. You must have at least one of the following access privileges to work with Allocation Manager:

  – Create General Ledger Allocation Rules

- – Administer Allocation Rules

- – Generate General Ledger Allocation Rules.

- The Essbase cluster name is not `Essbase_FA_Cluster`. If the Essbase server is not under this cluster, it is ignored by Allocation Manager.

- The stripe ID for General Ledger is different than the stripe ID for Workspace, so different security may be applied when Allocation Manager communicates with Essbase. Check the Allocation Manager log file to ensure that the stripe ID (or the application ID) for Allocation Manager is the same as the stripe ID for General Ledger.

- The Essbase application contains an empty Comment field. Each Essbase application must contain at least one character in the Comment field, or it is ignored by Allocation Manager.

- The application is not an Essbase aggregate storage application; only Essbase aggregate storage applications may be used with General Ledger.

18.5.1.6.2 Business Rules Cannot be Deployed to General Ledger

Problem

Business rules cannot be deployed from Allocation Manager to General Ledger. You receive the error message:

```
no WSDLLOCATION set
```

Solution

In the EPM registry, ensure that the `WSDLLOCATION` property is defined for Allocation Manager. The `WSDLLOCATION` property contains the General Ledger web services URL.

18.5.1.7 Smart View

This section describes common problems and solutions for Smart View. It contains the following topics:

- Re-enabling Smart View in Microsoft Office Applications

- Enabling Advanced Logging

- Timeout Errors

- Shared Connections Panel Does Not Display Server Names

18.5.1.7.1 Re-enabling Smart View in Microsoft Office Applications

Problem

Smart View may become disabled in Microsoft Office applications.

Solution

Smart View is a COM add-in, which Microsoft Office applications can disable. To re-enable Smart View, follow instructions in Excel Help for re-enabling COM add-ins.

18.5.1.7.2 Enabling Advanced Logging

Problem

Smart View collects and records events, errors, and other information in a log file, typically `SmartViewLogs.log`. When you experience performance or other issues, you can enable the logging in this file of additional information about profiling and requests and responses between Smart View and the server. This additional information can help Oracle Support to troubleshoot your issues.

Solution

To enable additional troubleshooting information:

1. Close all Microsoft Office applications.

2. Select **Start**, and then **Run**.

3. Enter `regedit` and click **OK** to open the registry.

4. Navigate to `HKEY_CURRENT_USER\Software\Hyperion Solutions\HyperionSmartView\Options`.

5. Right click **Options**, select **New**, and then **String Value**.

6. Name the new string value **Profile**.

7. Double-click **Profile** to open **Edit String**.

8. From **Edit String**, under **Value Data**, enter one of the following values:

   - **0:** Logging is not enabled. Initialized value `bEnableProfiling` is set to `false`.

   - **1:** Profile entries are logged when an event completes. This setting provides little information about abrupt terminations, but performance is better than with 2.

   - **2:** Profile entries are logged immediately. Use this setting to determine the function in which a termination occurred. This setting provides the most detailed information but has the greatest negative impact on performance.

9. Close the registry.

10. Restart Excel.

11. From the Smart View ribbon, select **Options**.

12. Go to the **Advanced** page and ensure that **Route Messages to File** is selected. The log file, typically SmartViewLogs.log, will now begin recording profiling and request/response information between Smart View and the server in addition to the other information it records.

   > **Note:** To improve performance, when you no longer need to log profiling and request/response information, delete the **Profile** entry that you created in the registry.

18.5.1.7.3 Timeout Errors

Problem

If the server takes longer to process a Smart View operation than the timeout value set on the client computer, users may receive a connection timeout error, or zero values may be displayed for Smart View functions.

Solution

Increase the timeout limit for Smart View client computers. Smart View uses Win-Inet APIs to communicate with the provider. These are the same modules that Internet Explorer uses. To increase the timeout value for a Windows client computer:

1. In the Windows registry of the client computer, navigate to `HKEY_CURRENT_ USER\Software\Microsoft\Windows\CurrentVersion\InternetSettin gs`.

2. Set the following values:

   ```
   "ReceiveTimeout"=dword:00dbba00
   "KeepAliveTimeout"=dword:0002BF20
   "ServerInfoTimeout"=dword:0002BF20
   ```

18.5.1.7.4 Shared Connections Panel Does Not Display Server Names

Problem

The **Shared Connections** panel of the Smart View Panel displays no server names. This typically happens when your applications are not properly registered during installation and configuration.

Solution

Review the information in the *Oracle Hyperion Enterprise Performance Management System Installation and Configuration Guide* to ensure that applications and servers are properly configured.

18.5.2 Financial Reporting

This section describes common problems and solutions for Financial Reporting. It contains the following topics:

- Re-enabling Smart View in Microsoft Office Applications
- Enabling Advanced Logging
- Timeout Errors
- Shared Connections Panel Does Not Display Server Names

18.5.2.1 Setting Up Data Sources and Debugging Setup and Connectivity Issues

This section describes common problems and solutions for setting up data sources and debugging setup and connectivity issues:

- Unable to Connect to Essbase Cube
- Error While Creating Database Connection
- Cannot Set Up a Database Connection

18.5.2.1.1 Unable to Connect to Essbase Cube

Problem

You are unable to connect to an Essbase cube.

Solution

Verify that the Essbase Server name and port are correct. Verify the user credentials used to make the connection are correct and the User has at least read rights to the cube.

18.5.2.1.2 Error While Creating Database Connection

Problem

An error occurred while creating a database connection.

Solution

Check the Financial Reporting log file. If the log file has the following error, a Catalog GUID refresh is required.

```
"Caused by: javax.jcr.LoginException: access denied for user to path
/users/userid"
```

For more information, see "Regenerating User GUIDs" in *Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition*.

18.5.2.1.3 Cannot Set Up a Database Connection

Problem

You are unable to set up a database connection.

Solution

When you create a database connection, it is appended to the list in the Database Manager dialog box.

To add a database connection:

1. In Workspace, select **Tools**, then **Database Connection Manager**.

2. Click **New**.

3. In Database Connection Name, enter a name for the database connection.

4. Select a data source type of Essbase.

5. In **Server**, add the **Essbase Server Name** and **Port Number**.

6. Re-enter the **User ID** and **Password**.

7. To add application and database names, click the magnifying glass, and make your selections.

> **Note:** The **Application Lookup** button displays a tree view of the applications and corresponding databases; the applications are listed as parents and the databases are listed as children. You can search on an application or database. For data sources that are not associated with a database, only applications are listed.

8. Click **OK**.

    The database connection profile is appended to the list in Database Connection Manager dialog box.

18.5.2.2 Issues with Workspace (Access Privileges, Preferences, Change And Manage Database Connections)

This section describes common problems and solutions for setting up data sources and debugging setup and connectivity issues:

- Unable to Connect to Essbase Cube

- Error While Creating Database Connection

- Cannot Set Up a Database Connection

18.5.2.2.1 Error When Running a Report

Problem

When running a report, following error is displayed:

```
Error 1012:Report contains an invalid grid. The following dimensions could not be
found: <i>xxx</i>
```

Solution

This error may occur if the database was recently changed on the report. To resolve the issue, open and save the report that has mismatched dimensions. This causes the dimensions that existed in the old database connection but not in the new database connection to be removed. The dimension and its members that existed in the rows and columns are removed from the grid. If, as a result of the removal, no dimension exists in the row or column, you need to add a valid dimension to the cleared row or column in order for the report to run. Dimensions that exist in the new database connection but not in the old one, are added to the POV.

18.5.2.3 Common Administrative Tasks and How to Debug Issues

This section describes common problems and solutions for setting up data sources and debugging setup and connectivity issues:

- Log File Output Management and Diagnosis of Log Issues

- Error While Creating Database Connection

- Log File Output Management and Diagnosis of Log Issues

- Cannot Set Up a Database Connection

18.5.2.3.1 Log File Output Management and Diagnosis of Log Issues

Problem

Financial Reporting cannot be accessed from Workspace.

Solution

1. Check if the Financial Reporting web application is running. In Internet browser, enter protocol://server:port/hr/status.jsp. If the web application is running, the following message is displayed:

   ```
   The Oracle© Hyperion Financial Reporting, Fusion Edition Web application is
   available.
   ```

2. View the Financial Reporting log file fr.log from the following directories to see if there are any com.sun.xml.ws.wsdl.parser related errors:

```
(UNIX) DOMAIN_HOME/servers/server_name/logs/financialreporting
(Windows) DOMAIN_HOME\servers\server_name\logs\finsncialreporting
```

In this case, it is possible that analytics web application is not running. If `fr.log` contains multiple instances of `Caused by: javax.jcr.LoginException: access denied for user to path /users/userid`, then a Catalog GUID refresh is required.

For more information, see "Regenerating User GUIDs" in *Oracle Fusion Middleware Security Guide for Oracle Business Intelligence Enterprise Edition*.

18.5.2.3.2 Setting Up Financial Reporting TCP Ports for Firewall Environments and Debugging Issues Around It

Problem

Unable to connect to Financial Reporting through a firewall from the Financial Reporting Studio.

Solution

By default, Financial Reporting components communicate with each other through Remote Method Invocation (RMI) on dynamically assigned Transmission Control Protocol (TCP) ports. To communicate through a firewall, you must specify the port for each Financial Reporting component separated by the firewall in the `JConsole.exe` file, and then open the necessary ports in your firewall. In addition, you may need to open ports for the Reports Server RDBMs, for data sources that you report against, and for LDAP/NTMLM for external authentication. The Print Server defaults to 8297. This can be changed by modifying the `PrintServers` property in FR Mbeans and restarting. Once connected, all RMI Services can use anonymous ports by default for communication. Alternatively a range of ports can be configured for communication by setting `RMIPortRangeLower` and `RMIPortRangeUpper` within the Financial Reporting configuration. You can change the port assignments to use in a firewall environment for servers in the Financial Reporting Mbeans called `RMIPortRangeUpper` and `RMIPortRangeLower`.

To locate the Financial Reporting MBeans properties to modify in Fusion Applications Control:

1. From the navigation pane, expand the `BIDomain` farm, **Application Deployments**, and then **Internal Applications**.

2. Expand **FinancialReporting(11.1.1)(*bi_cluster_name*)**.

3. Right-click **FinancialReporting(11.1.1)(*bi_server_name*)** and choose **System MBean Browser**.

4. In the System MBean Browser page, expand **Application Defined MBeans**.

5. Expand **Application Defined MBeans**, **com.hyperion**, **Server:*bi_server_name***.
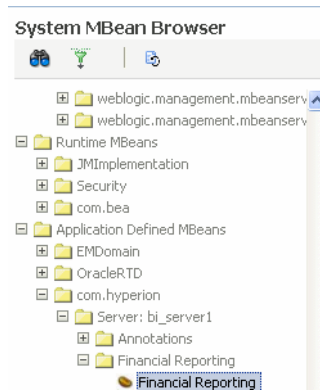
6. Expand **Financial Reporting**.

7. Click **Financial Reporting**.

8. In the Application Defined MBeans: Financial Reporting page, scroll down to **PrintServers**, **RMIPortRangeUpper**, and **RMIPortRangeLower.**

9. Click on each attribute and add a value in the **Value** field and click **Apply**.

18.5.2.3.3 Supporting PDF Print/Troubleshooting Financial Reporting Print Server Issues

Problem

Financial Reporting Print server is not working.

Solution

To resolve this issue:

1. Ensure that the Financial Reporting Print Server service has been created and started.

2. Verify 32-bit Ghostscript is installed by running. From the **Start** menu, choose **All Programs** > **Ghostscript**.

3. Examine the `FRPrintLogging.log` file in the *install_ directory*`\Oracle\FinancialReportingStudio\diagnostics\logs\FinancialReporting` directory.

4. Verify that the `PrintServers` property in Financial Reporting MBeans now shows the Financial Reporting Print Server computer and port number. The default port is 8297. This may show multiple financial Reporting Print Servers if more than one have been configured.

To locate the `PrintServers` property to modify in Fusion Applications Control:

a. From the navigation pane, expand the `BIDomain` farm, **Application Deployments**, and then **Internal Applications**.

b. Expand **FinancialReporting(11.1.1)(*bi_cluster_name*)**.

c. Right-click **FinancialReporting(11.1.1)(*bi_server_name*)** and choose **System MBean Browser**.

d. In the System MBean Browser page, expand **Application Defined MBeans**.

e. Expand **Application Defined MBeans**, **com.hyperion**, **Server:*bi_server_name*.**
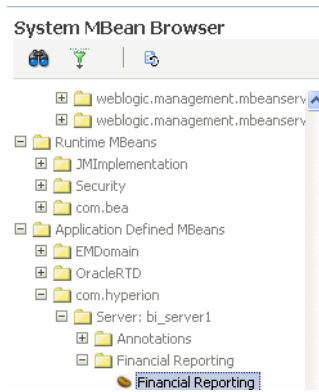
f. Expand **Financial Reporting**.

g. Click **Financial Reporting**.

h. In the Application Defined MBeans: Financial Reporting page, scroll down to **PrintServers**.

i. Click on the attribute and check the value in the **Value** field.

j. If necessary, adjust the value and click **Apply**.

5. Verify that the Financial Reporting Print Server server name and port can be reached from the Financial Reporting web application server machines. Use `telnet` to connect to the Financial Reporting Print Server.

18.5.2.3.4 Getting Error Application HReports is not Defined

Problem

When the Batch Scheduler is opened, the following error displays:

```
application HReports is not defined
```

This may happen if the `fressclient.ear` application is not started in the General Ledger's Oracle Enterprise Scheduler domain.

Solution

From the Oracle WebLogic Server Console for the Oracle Enterprise Scheduler domain, perform the following to steps:

1. Navigate to **Deployments** and select **fressclient**.

2. Select **Start**, then **Servicing all Requests**.

18.5.2.3.5 Permission Errors

Problem

When the Batch Scheduler is opened, or at the time of job submission, permission errors are shown. This may happen if the user is not granted the `oracle.as.scheduler.security.MetadataPermission` permission.

Solution

Check if the user has `oracle.as.scheduler.security.MetadataPermission` permission in the fscm stripe using Fusion Applications Control.

18.5.2.3.6 All the Jobs Fail

Problem

All the jobs fail.

Solution

This may happen if the Essbase server is not running.

To check if the Essbase server is running and restart (if needed):

1. Determine the current status:

   ```
   (UNIX) APPLICATIONS_CONFIG_HOME/BIInstance/bin/opmnctl status
   (Windows) APPLICATIONS_CONFIG_HOME\BIInstance\bin/opmnctl status
   ```

   OPMN generates a list of the running components and processes. The following message indicates that the Essbase Server (`essbaseserver1`) is already running:

   ```
   Processes in Instance: BIInstance
   -------------------------------+-------------------+--------+---------
   ias-component                  | process-type      |    pid | status
   -------------------------------+-------------------+--------+---------
   essbaseserver1                 | Essbase           |  27879 | Alive
   coreapplication_obiccs1        | OracleBIClusterCo~|  10828 | Alive
   coreapplication_obisch1        | OracleBIScheduler~|  18308 | Alive
   coreapplication_obijh1         | OracleBIJavaHostC~|  18337 | Alive
   coreapplication_obips1         | OracleBIPresentat~|  26455 | Alive
   coreapplication_obis1          | OracleBIServerCom~|  21716 | Alive
   ```

2. Restart the Essbase server:

   ```
   (UNIX) APPLICATIONS_CONFIG_HOME/BIInstance/bin/opmnctl restartproc
   ias-component=component_name
   (Windows) APPLICATIONS_CONFIG_HOME\BIInstance\bin\opmnctl restartproc
   ias-component=component_name
   ```

   For example

   ```
   APPLICATIONS_CONFIG_HOME/BIInstance/bin/opmnctl restartproc
   ias-component=essbaseserver1
   ```

18.5.3 Hyperion Provider Services

This section describes common problems and solutions for Oracle Hyperion Provider Services. It contains the following topics:

- Provider Services Server Not Running

- Provider Services Version Information

- Maximum Numbers of Rows Exceeded in Smart View

- Monitoring Provider Sessions

- Monitoring Active-Active Essbase Clusters

- Enabling and Disabling Active-Active Essbase Cluster Components

18.5.3.1 Provider Services Server Not Running

Problem

Provider Services appears not to be functional. You may need to determine whether Provider Services is running.

Solution

To determine whether Provider Services is running:

1. Launch a web browser.

2. Enter the connection URL for the corresponding products:

   - Smart View Provider: `https://Provider_Services_server:Provider_Services_port/aps/SmartView`

   - Java API: `https://Provider_Services_server:Provider_Services_port/aps/JAPI`

   If an HTTP error is returned, the Provider Services server is not running.

3. You may need to find out if you are connecting to the right port.

   If you are running Provider Services with Oracle Access Manager, then the default Provider Services port number is 9704. If you are not running Provider Services with Oracle Access Manager, then the default Provider Services port number is 4443.

4. To start the Provider Services server, use the following script from the `fusionapps` Middleware directory:

   ```
   (UNIX) FA_MW_HOME/user_projects/domains/bi_foundation_domain_
   name/bin/startWebLogic.sh
   (Windows) FA_MW_HOME\user_projects\domains\bi_foundation_dmain_
   name\bin\startWebLogic.cmd
   ```

18.5.3.2 Provider Services Version Information

Problem

Provider Services version information is not provided in a user interface.

Solution

You can find the Provider Services version information in the following locations:

- Operating system console window that is displayed when the Provider Services server is running

- Provider Services log files

18.5.3.3 Maximum Numbers of Rows Exceeded in Smart View

Problem

In Excel, a `Maximum numbers of rows [5000] Exceeded` error is returned if there are more than 5,000 rows on the Smart View grid. By default, the maximum number of rows is set to 5000 and the maximum number of columns to 256. In Excel 2003, you cannot exceed these limits. However, in Excel 2007 and 2010, there are no limits on the number of rows or columns, and you can change the default settings.

Solution

To change maximum row and column settings in Excel 2007 or 2010:

1. Launch Excel.

2. From the Smart View ribbon, select **Open**, and then **Smart View Panel**.

3. In the Smart View Panel, click **Edit Provider Services** to open the **Provider Services** server preferences dialog box.

4. Change the settings for **Number of Rows and Number of Columns** as needed. Enter 0 to remove any limits to the number of rows or columns.

18.5.3.4 Monitoring Provider Sessions

Problem

You may need to monitor the number of sessions on a Provider Services server.

Solution

To monitor Provider Services server sessions, from Administration Services Console:

1. From **Enterprise View** or a custom view, under the **Provider Services** node, select a provider.

2. Under the provider node, select **Action**, and then **Sessions**.

18.5.3.5 Monitoring Active-Active Essbase Clusters

Problem

You may need to monitor active-active Essbase clusters.

Solution

To monitor Provider Services server sessions, from Administration Services Console:

1. From Enterprise View or a custom view, under the **Provider Services** node, select a provider.

2. Under the provider node, expand **Essbase Clusters** and select the active-active Essbase cluster node you want to monitor.

3. Choose **Action**, and then edit the active-active Essbase cluster.

18.5.3.6 Enabling and Disabling Active-Active Essbase Cluster Components

Problem

You may need to enable or disable individual active-active Essbase cluster components.

Solution

To enable or disable an active-active Essbase cluster component, follow the steps in the *Oracle Hyperion Enterprise Performance Management System High Availability and Disaster Recovery Guide.*

18.5.4 Essbase

This section describes common problems and solutions. It contains the following topics:

- Getting Started with Logging Basics for Essbase

- Essbase Agent Startup Fails with Error

- Essbase Agent Startup Fails Due to serviceinstanceref ref="audit" Entry in jps-config-jse.xml

- [Essbase Agent Startup Fails with an Error While Loading Shared Libraries](#)

- [opmnctl Commands Fail to Execute](#)

- [An Application Stops Responding](#)

- [An Essbase Application Will Not Start](#)

- [ASO Database Corruption Error](#)

- [Essbase Login Credentials Are Unknown When Essbase is Included In the Oracle Business Intelligence Installation](#)

- [Cannot Stop an Application Process](#)

- [Changing the Essbase Ports (High-Availability Mode)](#)

- [Changing the Essbase Ports (Non-High-Availability Mode)](#)

- [Data Load Fails with the "Load Buffer Does Not Exist" Error](#)

- [Data Load Fails With Resource Usage Error](#)

- [Essbase Fails to Start in Cluster Mode](#)

- [Essbase Login Failed Due to Invalid Credentials](#)

- [Failed to Open a File on UNIX](#)

- [GL Writeback Fails with "Accounting Date Conversion" Error](#)

- [GL Writeback Fails With "Group ID Node" Error](#)

- [GL Writeback Fails With "Not a Valid GL Application" Error](#)

- [GL Writeback Fails With "SQL Database Connection" Error](#)

- [Network Timeout](#)

- [OPMN fails to start Essbase in High-Availability Mode](#)

- [Restructure Failure](#)

- [Security File is Corrupt](#)

- [Status of Essbase Agent Connection](#)

- [Unable to Write File During Data Load or Building Aggregate Views](#)

18.5.4.1 Getting Started with Logging Basics for Essbase

Log files are the best tools for analyzing what might be wrong with the system configuration. The logging configuration file, the defaults it ships with, and instructions for changing change those defaults will help with the analysis.

- Essbase log files

  ```
  (UNIX) APPLICATIONS_CONFIG_
  HOME/BIInstance/diagnostics/logs/Essbase/essbaseserver_name/essbase/Essbase.log
  (Windows) APPLICATIONS_CONFIG_
  HOME\BIInstance\diagnostics\logs\Essbase\essbaseserver_name\essbase\Essbase.log
  ```

- Oracle Diagnostic Logging (ODL) logs

  ```
  (UNIX) APPLICATIONS_CONFIG_
  HOME/BIInstance/diagnostics/logs/Essbase/essbaseserver_name/essbase/ESSBASE_
  ODL.log
  (Windows) APPLICATIONS_CONFIG_
  HOME\BIInstance\diagnostics\logs\Essbase\essbaseserver_name\essbase\ESSBASE_
  ODL.log
  ```

- Lease Manager logs

  (UNIX) *APPLICATIONS_CONFIG_
  HOME*/BIInstance/diagnostics/logs/Essbase/*essbaseserver_
  name*/essbase/leasemanager_essbase_*machine_name*.log
  (Windows) *APPLICATIONS_CONFIG_
  HOME*\BIInstance\diagnostics\logs\Essbase*essbaseserver_name*\leasemanager_
  essbase_*machine_name*.log

- Shared Services logs

  (UNIX) *APPLICATIONS_CONFIG_
  HOME*/BIInstance/diagnostics/logs/Essbase/*essbaseserver_
  name*/essbase/SharedServices_Security_Client.log
  (Windows) *APPLICATIONS_CONFIG_
  HOME*\BIInstance\diagnostics\logs\Essbase*essbaseserver_name*\SharedServices_
  Security_Client.log

- OPMN Logs

  (UNIX) *APPLICATIONS_CONFIG_HOME*/BIInstance/diagnostics/logs/OPMN/opmn/opmn.log
  (Windows) *APPLICATIONS_CONFIG_
  HOME*\BIInstance\diagnostics\logs\OPMN\opmn\opmn.log

- Esssbase ping log

  (UNIX) *APPLICATIONS_CONFIG_
  HOME*/BIInstance/diagnostics/logs/OPMN/opmn/EssbasePing.log
  (Windows) *APPLICATIONS_CONFIG_
  HOME*\BIInstance\diagnostics\logs\OPMN\opmn\EssbasePing.log

- OPMN debug log

  (UNIX) *APPLICATIONS_CONFIG_HOME*/BIInstance/diagnostics/logs/OPMN/opmn/debug.log
  (Windows) *APPLICATIONS_CONFIG_
  HOME*\BIInstance\diagnostics\logs\OPMN\opmn\debug.log

  > **Note:** The system will generate the OPMN debug log only if DEBUG
  > mode is turned on in `opmn.xml`.

- OPMN console log

  (UNIX) *APPLICATIONS_CONFIG_
  HOME*/BIInstance/diagnostics/logs/OPMN/opmn/console*.log
  (Windows) *APPLICATIONS_CONFIG_
  HOME*\BIInstance\diagnostics\logs\OPMN\opmn\console*.log

  > **Note:** The system will generate the OPMN console log only if `DEBUG`
  > mode is turned on in `opmn.xml`.

18.5.4.2 Essbase Agent Startup Fails with Error

Problem

The Essbase Agent startup fails with the following error:

```
Fatal Error: FUSIONAPPID not specified in the cfg file
```

Solution

To modify `essbase.cfg` and restart the Essbase server:

1. Locate the `essbase.cfg` file in the following directories:

   ```
   (UNIX) APPLICATIONS_CONFIG_HOME/BIInstance/Essbase/essbaseserver_name/bin
   (Windows) APPLICATIONS_CONFIG_HOME\BIInstance\Essbase\essbaseserver_name\bin
   ```

2. Make sure that the following entry is present in `essbase.cfg`:

   ```
   FUSIONAPPID appidname
   ```

3. Restart the Essbase server.

   To restart the Essbase server using `opmnctl`:

   a. Determine the current status:

   ```
   (UNIX) APPLICATIONS_CONFIG_HOME/BIInstance/bin/opmnctl status
   (Windows) APPLICATIONS_CONFIG_HOME\BIInstance\bin/opmnctl status
   ```

   OPMN generates a list of the running components and processes. The following message indicates that the Essbase Server (`essbaseserver1`) is already running:

   ```
   Processes in Instance: BIInstance
   --------------------------------+--------------------+---------+---------
   ias-component                   | process-type       |     pid | status
   --------------------------------+--------------------+---------+---------
   essbaseserver1                  | Essbase            |   27879 | Alive
   coreapplication_obiccs1         | OracleBIClusterCo~ |   10828 | Alive
   coreapplication_obisch1         | OracleBIScheduler~ |   18308 | Alive
   coreapplication_obijh1          | OracleBIJavaHostC~ |   18337 | Alive
   coreapplication_obips1          | OracleBIPresentat~ |   26455 | Alive
   coreapplication_obis1           | OracleBIServerCom~ |   21716 | Alive
   ```

   b. Restart the Essbase server:

   ```
   (UNIX) APPLICATIONS_CONFIG_HOME/BIInstance/bin/opmnctl restartproc
   ias-component=component_name
   (Windows) APPLICATIONS_CONFIG_HOME\BIInstance\bin\opmnctl restartproc
   ias-component=component_name
   ```

   For example

   ```
   APPLICATIONS_CONFIG_HOME/BIInstance/bin/opmnctl restartproc
   ias-component=essbaseserver1
   ```

   To start Essbase server using Fusion Applications Control:

   a. From the navigation pane, expand the farm and then **WebLogic Domain** for the `BIDomain` domain.

   b. Expand **Essbase Servers**, and then select the Essbase server.

   c. From the **Essbase Server** menu, choose **Administration**, then **Ports Configuration**.

   d. Select the **Listen** port, and then click **Edit**.

   e. Change the port number, and then click **OK**.

   f. From the **Essbase Server** menu, choose **Control**, then **Restart**.

18.5.4.3 Essbase Agent Startup Fails Due to serviceinstanceref ref="audit" Entry in jps-config-jse.xml

Problem

The Essbase Agent startup fails with the following error:

```
1051223 - Single Sign On function call [css_init] failed with error [CSS Error:
CSS method invocation error: getInstance: Failed to get CSSSystem instance,
please check SharedServices_Security_Client.log for more information
```

Solution

To modify `essbase.cfg` and restart the Essbase server:

1. Locate the `essbase.cfg` file in the following directories:

    ```
    (UNIX) APPLICATIONS_CONFIG_HOME/BIInstance/Essbase/essbaseserver_name/bin
    (Windows) APPLICATIONS_CONFIG_HOME\BIInstance\Essbase\essbaseserver_name\bin
    ```

2. Make sure that the following entry is commented out using XML style comments in the `essbase.cfg` file. For example, `<!-- -->`:

    ```
    serviceinstanceref ref="audit"
    ```

3. Restart the Essbase server.

    To restart the Essbase server using `opmnctl`:

    a. Determine the current status:

    ```
    (UNIX) APPLICATIONS_CONFIG_HOME/BIInstance/bin/opmnctl status
    (Windows) APPLICATIONS_CONFIG_HOME\BIInstance\bin/opmnctl status
    ```

    OPMN generates a list of the running components and processes. The following message indicates that the Essbase Server (`essbaseserver1`) is already running:

    ```
    Processes in Instance: BIInstance
    --------------------------------+--------------------+---------+---------
    ias-component                   | process-type       |     pid | status
    --------------------------------+--------------------+---------+---------
    essbaseserver1                  | Essbase            |   27879 | Alive
    coreapplication_obiccs1         | OracleBIClusterCo~ |   10828 | Alive
    coreapplication_obisch1         | OracleBIScheduler~ |   18308 | Alive
    coreapplication_obijh1          | OracleBIJavaHostC~ |   18337 | Alive
    coreapplication_obips1          | OracleBIPresentat~ |   26455 | Alive
    coreapplication_obis1           | OracleBIServerCom~ |   21716 | Alive
    ```

    b. Restart the Essbase server:

    ```
    (UNIX) APPLICATIONS_CONFIG_HOME/BIInstance/bin/opmnctl restartproc
    ias-component=component_name
    (Windows) APPLICATIONS_CONFIG_HOME\BIInstance\bin\opmnctl restartproc
    ias-component=component_name
    ```

    For example

    ```
    APPLICATIONS_CONFIG_HOME/BIInstance/bin/opmnctl restartproc
    ias-component=essbaseserver1
    ```

    To start the Essbase server using Fusion Applications Control:

    a. From the navigation pane, expand the farm and then **WebLogic Domain** for the `BIDomain` domain.

  b. Expand **Essbase Servers**, and then select the Essbase server.

  c. From the **Essbase Server** menu, choose **Administration**, then **Ports Configuration**.

  d. Select the **Listen** port, and then click **Edit**.

  e. Change the port number, and then click **OK**.

  f. From the **Essbase Server** menu, choose **Control**, then **Restart**.

18.5.4.4 Essbase Agent Startup Fails with an Error While Loading Shared Libraries

Problem

The Essbase Agent startup fails with the following error:

```
Error while loading shared libraries: libARicu24.so: cannot open shared object
file: No such file or directory
```

Solution

To modify opmn.xml file and restart the Essbase server:

1. Locate the opmn.xml file in the following directories:

```
(UNIX) APPLICATIONS_CONFIG_HOME/config/OPMN/opmn
(Windows) APPLICATIONS_CONFIG_HOME\BIInstance\config\OPMN\opmn
```

2. Update the following ODBC-related entries in opmn.xml:

```
<variable append="true" id="LD_LIBRARY_PATH" value="$ORACLE_
HOME/common/ODBC/Merant/6.0/lib$:$ORACLE_
HOME/jdk/jre/lib/i386/server$:$ESSBASEPATH/bin"/>

<variable id="ODBCINI" value="$ORACLE_HOME/common/ODBC/Merant/6.0/odbc.ini"/>

<variable id="ODBCINST" value="$ORACLE_
HOME/common/ODBC/Merant/6.0/odbcinst.ini"/>
```

3. Restart the Essbase server:

  a. Determine the current status:

```
(UNIX) APPLICATIONS_CONFIG_HOME/BIInstance/bin/opmnctl status
(Windows) APPLICATIONS_CONFIG_HOME\BIInstance\bin/opmnctl status
```

  OPMN generates a list of the running components and processes. The
  following message indicates that the Essbase Server (essbaseserver1) is
  currently running:

```
Processes in Instance: BIInstance
-------------------------------+--------------------+---------+---------
ias-component                  | process-type       |     pid | status
-------------------------------+--------------------+---------+---------
essbaseserver1                 | Essbase            |   27879 | Alive
coreapplication_obiccs1        | OracleBIClusterCo~ |   10828 | Alive
coreapplication_obisch1        | OracleBIScheduler~ |   18308 | Alive
coreapplication_obijh1         | OracleBIJavaHostC~ |   18337 | Alive
coreapplication_obips1         | OracleBIPresentat~ |   26455 | Alive
coreapplication_obis1          | OracleBIServerCom~ |   21716 | Alive
```

  b. Restart the Essbase server:

```
(UNIX) APPLICATIONS_CONFIG_HOME/BIInstance/bin/opmnctl restartproc
```

```
ias-component=component_name
(Windows) APPLICATIONS_CONFIG_HOME\BIInstance\bin\opmnctl restartproc
ias-component=component_name
```

For example

```
APPLICATIONS_CONFIG_HOME/BIInstance/bin/opmnctl restartproc
ias-component=essbaseserver1
```

18.5.4.5 opmnctl Commands Fail to Execute

Problem

A communication error occurs when attempting to issue commands using the opmnctl command line for OPMN.

For example:

```
qtfhp3:/vol1/nnguyen/rc11/Oracle/Middleware/user_
projects/epmsystem1/bin]$ opmnctl status RCV: No such file or
directory Communication error with the OPMN server local port.
Check the OPMN log files opmnctl status: opmn is not running.
```

Solution

To modify opmn.xml:

1. Edit opmn.xml to assign a different local and a remote port to OPMN, or just a remote port to OPMN. The currently assigned ports may already be used by another process. See the following example:

   ```
   <notification-server interface="any"> <ipaddr
   remote="<hostname>"/> <port local="6711" remote="6712"/>
   ```

2. Restart OPMN and try the opmnctl command again.

18.5.4.6 An Application Stops Responding

Problem

An Essbase application stops responding or shuts down abnormally.

Solution

To determine why the application is not responding:

1. Check the Essbase.log file for the following error message:

   ```
   1002089 RECEIVED ABNORMAL SHUTDOWN COMMAND - APPLICATION
   TERMINATING
   ```

2. Check the Essbase.log file for the following message:

   ```
   Exception error log [log00001.xcp] is being created...
   ```

3. If either of the above messages are found, contact Technical Support.

   Note that if you find the log00001.xcp file, save it. Oracle Support will need this file to troubleshoot the issue.

   ```
   (UNIX) APPLICATIONS_CONFIG_
   HOME/BIInstance/diagnostics/logs/Essbase/essbaseserver_
   name/essbase/app/appname/log000001.xcp
   (Windows) APPLICATIONS_CONFIG_
   HOME\BIInstance\diagnostics\logs\Essbase\essbaseserver_
   ```

name\essbase\app*appname*\log000001.xcp

18.5.4.7 An Essbase Application Will Not Start

Problem

An application does not respond after being started, or an error message like the following is displayed in the Essbase log file:

```
Network Error 10048 : Unable to Bind Host Server Socket On Port
4213
```

Solution

A port conflict may be occurring. Correct the port-related entries in `opmn.xml`.

For example, update the Essbase port range to match the port that Essbase is actually using:

```
<port id="essbase-port-range" range="32768-33768"/>
```

18.5.4.8 ASO Database Corruption Error

Problem

During an operation on an Essbase aggregate storage database, the following error is displayed:

```
Persistent data does not match the outline. There is no member
in dimension [Abc] for member number [12345]. Data is corrupted.
```

Solution

To clear the database and reload it from the original sources or from a saved exported backup:

1. Save the list of current aggregate views.

2. Clear all data in the database.

3. Reload the data from original sources or from a backup.

4. Rebuild the aggregate views using the list specified.

18.5.4.9 Essbase Login Credentials Are Unknown When Essbase is Included In the Oracle Business Intelligence Installation

Problem

You do not know the Essbase login credentials when Essbase is included in the Oracle Business Intelligence installation.

Solution

Use the Weblogic default user name and password to log in to Essbase. The default user name and password are `weblogic` and `welcome1`.

18.5.4.10 Cannot Stop an Application Process

Problem

An Essbase application process cannot be shut down.

Solution

To check the application log for active processing:

1. Access the application log. Note the following locations:

   ```
   (UNIX) APPLICATIONS_CONFIG_
   HOME/BIInstance/diagnostics/logs/Essbase/essbaseserver_
   name/essbase/app/appname/appname.log
   (Windows) APPLICATIONS_CONFIG_
   HOME\BIInstance\diagnostics\logs\Essbase\essbaseserver_
   name\essbase\app\appname\appname.log
   ```

2. Check the application log and look for the following message:

   ```
   RECEIVED SHUTDOWN COMMAND – SERVER TERMINATING.
   ```

3. If you see the proceeding message, wait a few minutes to allow the application to self-terminate. The application may be doing cleanup tasks associated with the shutdown process.

4. If you see the following message, the application is still processing user requests. Wait for user requests or any other operation that is still in progress to terminate, and then try shutting down the application.

   ```
   Cannot unload database <dbname> while user <username> is
   performing database operation. Wait for the user to complete
   the operation, or ask the user to abort it. Log out all users
   and then unload the database. Cannot unload database <dbname>
   when it is still in use
   ```

5. If you need to terminate the application, perform the following steps:

   a. If you need to terminate the application even when there are user requests in progress on the database, you can run the following sequence MaxL statements to forcefully terminate user requests:

   ```
   #Display active sessions to see current requests display
   session on application <appname>
   #Disallow new connections to the application alter
   application <appname> disable connects;
   #Force logout of all, and terminate requests alter system
   logout session on database <appname>.<dbname> force;
   ```

   b. After forcefully logging out all users, wait for a few minutes, and then use the following MaxL statement to check if any user requests are still running:

   ```
   #Display active sessions display session on application
   <appname>;
   ```

   c. If there are no requests running, attempt to stop the application process.

18.5.4.11 Changing the Essbase Ports (High-Availability Mode)

Problem

The Essbase port numbers need to be changed when Essbase is installed in high-availability mode.

Solution

To run the Essbase failover automation script to change the Essbase port:

1. Go to the following location:

```
(UNIX) APPLICATIONS_CONFIG_HOME/BIInstance/bin/essbase_ha
(Windows) APPLICATIONS_CONFIG_HOME\BIInstance\bin\essbase_ha
```

2. Look for the `%SHARED_FOLDER%` location specified in the file
   `essfoenv.properties`.

3. Edit the `%SHARED_FOLDER%\EssFoConfig.properties` file or `%SHARED_
   FOLDER%/EssFoConfig.properties` file.

4. Look for `SYSTEM_AGENT_PORTNUMBER(1|2)`, and change the port number.
   There are two Essbase instances, so make sure to change the one that you need.

5. Execute `essfoconfig.sh` (UNIX) or `essfoconfig.bat` (Windows) without
   any parameter. The Help prints so you can see a list of options for parameters.

6. Execute the following command:

   ```
   essfoconfig.sh update prpName
   ```

   where *prpName* is the port number property in Step 4.

   For example, the following command updates the agent port of the first Essbase
   instance:

   ```
   essfoconfig.sh update SYSTEM_AGENT_PORTNUMBER1
   ```

18.5.4.12 Changing the Essbase Ports (Non-High-Availability Mode)

Problem

The Essbase port numbers need to be changed when Essbase is installed in
non-high-availability mode.

Solution

Go to Fusion Applications Control and change the Essbase port.

1. From the navigation pane, expand the farm and then **WebLogic Domain** for the
   `BIDomain` domain.

2. Expand **Essbase Servers**, and then select the Essbase server.

3. From the **Essbase Server** menu, choose **Administration**, then **Ports
   Configuration**.

4. Select the **Listen** port, and then click **Edit**.

5. Change the port number, and then click **OK**.

6. Restart Essbase server. From the **Essbase Server** menu, choose **Control**, then
   Restart.

18.5.4.13 Data Load Fails with the "Load Buffer Does Not Exist" Error

Problem

Following a data load failure to a specified load buffer, the following error displays for
subsequent loads to the same load buffer:

```
Data load buffer [123] does not exist.
```

Solution

This error may occur when data loads initialize a load buffer and load multiple data
files to it before committing it to the cube. If an error occurs that causes one of the

steps to fail, then the load buffer is automatically destroyed, causing all subsequent steps to fail with the previous error.

Use one of the following methods to correct this error:

- Determine the cause of the original data load error and try to resolve it.

- Set the Essbase data load options to indicate that the data load should not be aborted on error; instead, have data load errors written or appended to the log file

18.5.4.14 Data Load Fails With Resource Usage Error

Problem

Upon attempting to load data into an Essbase aggregate storage database, the following error displays:

```
Specified load buffer resource usage [100] is above currently
available value [0].
```

Solution

Other ongoing data load operations have reserved part of the cache for their load buffers. Use one of the following methods to correct this error:

- Reduce the resource usage for this data load, and try again. If you are using MaxL, you must explicitly create the load buffer, using the optional `resource_usage` argument. For example:

  ```
  alter database AsoSamp.Sample initialize load_buffer with
  buffer_id 1 resource_usage .5;
  ```

- Wait for other operations to finish.

- To see what reservations have been made to the cache resources, run the following MaxL statement:

  ```
  query database "app"."db" list load_buffers;
  ```

  For more information about the `query database`, `alter database`, and other MaxL statements, see the *Oracle Essbase Technical Reference*.

18.5.4.15 Essbase Fails to Start in Cluster Mode

Problem

Essbase does not start when it is in cluster mode and Oracle Business Intelligence domain is installed with a non-domain-qualified host name.

Solution

The error occurs when the Essbase host specified in the cluster configuration properties file, `EssFOConfig.properties`, is domain-qualified, but is not in the EPM Registry, leading to a mismatch.

For more information about viewing the EPM Registry, see the "Viewing the Components in the Shared Services Registry," section in the *Oracle Hyperion Enterprise Performance Management System Installation and Configuration Guide*.

To check if there is an Essbase host name qualification mismatch:

1. Use the following command to return all the Essbase clusters that were configured in the registry:

```
(UNIX) APPLICATIONS_CONFIG_HOME/BIInstance/config/foundation/11.1.2/epmsys_
registry.sh view CLUSTER
(Windows) APPLICATIONS_CONFIG_HOME\BIInstance\config\foundation\11.1.2\epmsys_
registry.bat view CLUSTER
```

2. Drill into the specific Essbase Server instance underneath this cluster, and view that node to see the host name:

```
(UNIX) APPLICATIONS_CONFIG_HOME/BIInstance/config/foundation/11.1.2/epmsys_
registry.sh view #Essbase_Server_GUID
(Windows) APPLICATIONS_CONFIG_HOME\BIInstance\config\foundation\11.1.2\epmsys_
registry.bat view #Essbase_Server_GUID
```

Note that GUID is the global unique ID of an Essbase Server instance as returned from the full viewing of the EPM Registry.

3. Access the EPM registry and confirm that the host in the cluster configuration properties file displays exactly how Essbase is configured in the EPM Registry.

4. Restart Essbase.

18.5.4.16 Essbase Login Failed Due to Invalid Credentials

Problem

When the user attempts to log in to Essbase, the login attempt fails. The problem may be an authentication failure. When the Essbase login fails due to an authentication problem, it reports the following errors:

```
ERROR – 103 - Unexpected Essbase error 1051440.
```

```
ERROR – 1051440 - Essbase user [bi-001] Authentication Fails
against the Shared Services Server with Error [EPMCSS-1009004:
Failed to read data from the policy store.].
```

Solution

To determine login failure problems by gathering detailed error messages:

1. Locate the essbase.cfg file in the following directories:

```
(UNIX) APPLICATIONS_CONFIG_HOME/BIInstance/Essbase/essbaseserver_name/bin
(Windows) APPLICATIONS_CONFIG_HOME\BIInstance\Essbase\essbaseserver_name\bin
```

2. Using a text editor, add the following text on its own line:

```
LOGINFAILUREMESSAGEDETAILED
```

For more information about this and other essbase.cfg configuration settings, see the *Oracle Essbase Technical Reference*.

3. Stop and restart the Essbase server. For detailed procedures, see "Starting and Stopping Essbase Server" in the *Oracle Essbase Database Administrator's Guide*.

4. Use the following paths to access the Essbase.log file, which contains any error messages:

```
(UNIX) APPLICATIONS_CONFIG_
HOME/BIInstance/diagnostics/logs/Essbase/essbaseserver_name/essbase/Essbase.log
(Windows) APPLICATIONS_CONFIG_
HOME\BIInstance\diagnostics\logs\Essbase\essbaseserver_
name\essbase\Eessbase.log
```

5. In the same path where `Essbase.log` is found, see the `SharedServices_Security_Client.log` to locate security-related error messages.

To gather debug-level error messages:

> **Note:** Keep the debug statistics available for diagnostic purposes, in case you need to contact Technical Support.

1. Locate the `essbase.cfg` file in the following directories:

   ```
   (UNIX) APPLICATIONS_CONFIG_HOME/BIInstance/Essbase/essbaseserver_name/bin
   (Windows) APPLICATIONS_CONFIG_HOME\BIInstance\Essbase\essbaseserver_name\bin
   ```

2. Using a text editor, add the following text on its own line:

   ```
   AGENTLOGMESSAGELEVEL DEBUG
   ```

   For more information about this and other `essbase.cfg` configuration settings, see *Oracle Essbase Technical Reference*.

3. Stop and restart the Essbase server. For detailed procedures, see the "Starting and Stopping Essbase Server" section in *Oracle Essbase Database Administrator's Guide*.

4. Use the following paths to access the `Essbase.log` file, which contains any error messages:

   ```
   (UNIX) APPLICATIONS_CONFIG_
   HOME/BIInstance/diagnostics/logs/Essbase/essbaseserver_name/essbase/Essbase.log
   (Windows) APPLICATIONS_CONFIG_
   HOME\BIInstance\diagnostics\logs\Essbase\essbaseserver_
   name\essbase\Eessbase.log
   ```

To check for and correct BI domain login errors:

See Table 18–3 to identify if you have a debug login failure problem. Go to each file listed in the table and search for the corresponding error message. If one or both of your files contain errors, use the following procedure to correct them.

Table 18–3   File Names and Error Messages

File Name	Error Message
`AdminServer.out`	`<oracle.epm.oem.essbase.mbeans.EssbaseServerRegistrationMBeanImpl>`
	`<BEA-000000> <Child MBean registration failed due to some reason Not able to get Credential Store Factory from EM java.lang.Exception: Not able to get Credential Store Factory from EM`
`BIDomain.log`	`Exception while retrieving data from Essbase Cannot connect to olap service.`
	`Cannot connect to Essbase Server. Error:EssbaseError(1051293): Login fails due to invalid login credentials com.essbase.api.base.EssException: Cannot connect to olap service.`
	`Cannot connect to Essbase Server. Error:Essbase`

To add a valid Essbase user name to the credential store from Fusion Applications Control, perform the following:

1. From the navigation pane, expand the farm and then **WebLogic Domain** for the `BIDomain` domain.

2. Right-click **BIDomain** and click **Security**.

3. Click **Credentials**. The Credentials page displays.

4. In the Credential table, expand **essbaseserver**.

5. Select the Essbase server and click the **Edit** the selected credential key button. The Edit Key page displays.

6. Verify that the user name and password are correct.

18.5.4.17 Failed to Open a File on UNIX

Problem

The following error occurs on a UNIX platform:

```
Failed to open file [filename]: a system file error occurred.
Please see application log for details.
```

Solution

To open a file on UNIX:

1. Confirm that the specified file exists.

2. If the file exists, increase the maximum number of open file descriptors. To do this, consult the UNIX operating system's documentation for the `ulimit` command.

18.5.4.18 GL Writeback Fails with "Accounting Date Conversion" Error

Problem

GL writeback fails with the following error:

```
The accounting date conversion failed for %s to %s in the
ACCOUNTING_DATE column. Unable to proceed with GL export.
```

Problem

This error may occur when GL writeback cannot format the date string in the way that GL tables expect.

To fix this issue, make sure that the `<DATE_FORMAT>` tag in the `cubeMap.xml` file has date format represented in one of the following acceptable formats, and that it matches with the date value stored in the accounting date alias table of the outline. The following are the acceptable date formats:

```
mon dd yyyy
Month dd yyyy
mm/dd/yy
mm/dd/yyyy
yy.mm.dd
dd/mm/yy
dd.mm.yy
dd-mm-yy
dd Month yy
dd mon yy
Month dd, yy
mon dd, yy
mm-dd-yy
```

```
yy/mm/dd
yymmdd
dd Month yyyy
dd mon yyyy
dd/mon/yy
yyyy-mm-dd
yyyy/mm/dd
Day, Month dd, yyyy
```

18.5.4.19 GL Writeback Fails With "Group ID Node" Error

Problem

GL writeback fails with the following error:

```
The group id node has active GL operations in state [<state_name>]. Use the
appropriate API call in sequence. Unable to proceed with [api name] call.
```

Solution

This error may occur if more than one user tries to use the GL writeback-related API calls with the same group ID. To fix this issue, ensure that all users using GL writeback operations are using different group IDs.

18.5.4.20 GL Writeback Fails With "Not a Valid GL Application" Error

Problem

GL writeback fails with the following error:

```
This application is not a valid GL application for Essbase.
```

Solution

The error may occur if the `cubeMap.xml` file is either not present in the *app/ db* directory, or is not parsed properly. For the GL writeback to work properly, the `cubeMap.xml` file must be present and contain the required information to be parsed successfully.

To confirm that `cubeMap.xml` was parsed correctly:

1. Ensure that `cubeMap.xml` file exists in the *app/ db* directory and contains the correct information.

2. To ensure that `cubeMap.xml` was found and parsed successfully, check the Essbase application log for the following entry:

   ```
   Parsing of cubeMap.xml file succeeded
   ```

18.5.4.21 GL Writeback Fails With "SQL Database Connection" Error

Problem

GL writeback fails with the following error:

```
Failed to Establish Connection With SQL Database Server. See log
for more information.
```

Solution

The error may occur if the SQL drivers are not set up correctly.

To confirm that the SQL drivers are set up correctly:

1. Ensure that the ODBC Merant drivers path in `odbcinst.ini` is correct.

2. Ensure that the **ODBCINST** variable in `opmn.xml` is pointing to the correct `odbcinst.ini` file.

3. Ensure that an Oracle GL target database is running. For the GL writeback to work, a GL system is required, with `cubeMap.xml` correctly set up with the GL system information as follows:

   ```
   <HOST>somemachine.company.com</HOST> <PORT>port</PORT> <SID>SID</SID>
   ```

4. Ensure that `Essbase.cfg` contains driver descriptors information in the following format, so that Essbase can connect to drivers in the `odbcinst.ini`:

   ```
   BPM_Oracle_DriverDescriptor "DataDirect 6.0 Oracle Wire
   Protocol"
   ```

18.5.4.22 Network Timeout

Problem

When performing an operation against an Essbase cube, one of the following errors occurs:

- `Network error [12345]: Cannot Send Data`

- `Network error [12345]: Cannot Receive Data`

Solution

- The error may indicate that Essbase has terminated abnormally. If an `xcp` file is found in the following location, save it and contact Oracle Support:

  ```
  (UNIX) APPLICATIONS_CONFIG_
  HOME/BIInstance/diagnostics/logs/Essbase/essbaseserver_
  name/essbase/app/appname/log000001.xcp
  (Windows) APPLICATIONS_CONFIG_
  HOME\BIInstance\diagnostics\logs\Essbase\essbaseserver_
  name\essbase\app\appname\log000001.xcp
  ```

- In `essbase.cfg`, increase the network timeout parameters using the settings `NETDELAY` and `NETRETRYCOUNT`. For more information, see "Essbase.cfg Configuration Settings" in *Oracle Essbase Technical Reference*.

18.5.4.23 OPMN fails to start Essbase in High-Availability Mode

Problem

OPMN fails to start Essbase in high-availability mode.

Solution

For Essbase clustering to work, all Essbase failover cluster data must be on shared storage with the `ARBORPATH` variables set correctly.

To confirm that Essbase can start in high-availability mode:

1. Ensure that OPMN is started by a network or domain user.

2. Ensure that `ARBORPATH`s are specified as mapped drives, and not as UNC paths.

18.5.4.24 Restructure Failure

Problem 1

During a dimension build or outline restructure, error code `1130203` occurs.

Solution 1

The error indicates that not enough memory was available to perform the operation. Use one of the following methods to correct this error:

- Increase the amount of virtual memory available to the operating system.

- If Essbase is running on a 32-bit platform, keep in mind that the maximum memory available to Essbase is between 2 GB and 4 GB, regardless of how much RAM is installed on the machine. Because Essbase loads both the old and new outlines in memory during the restructure, there is a restriction on the largest outline that can be modified on such a machine. For example, if the platform only allows 2 GB of memory to be used by a process, and an Essbase application process is already using 200 MB memory, then the maximum-size outline that can be modified is 900 MB. Switching to a 64-bit installation of Essbase will lift this limitation.

Problem 2

During dimension build or outline restructure, outline validation fails with one of the following errors:

- `There were errors validating the outline. Please check the error file.`

- `Outline has errors`

Solution 2

Review the application log or the error file (if available) to see what specific errors are causing the validation error.

Problem 3

During dimension build or outline restructure, one of the following errors occurs:

- `Cannot write the new outline file during the restructuring of [%s]`

- `Error writing outline change log file for database [%s]`

Solution 3

Check the application log for other errors. If none are found, check to see if the file system is full. Remember that making a change to an Essbase outline requires at least as much free space as the existing outline.

18.5.4.25 Security File is Corrupt

Problem

The `Console.log` displays the following error:

`Fatal Error: Invalid item index in security file`

Solution

The Essbase security file, `essbase.sec`, is invalid. Use one of the following methods to correct this error:

- Restart Essbase using the latest backup security file by copying `essbase_timestamp.bak` to `essbase.sec`. Both files are located in the following directories

  ```
  (UNIX) APPLICATIONS_CONFIG_HOME/BIInstance/Essbase/essbaseserver_name/bin
  (Windows) APPLICATIONS_CONFIG_HOME\BIInstance\Essbase\essbaseserver_name\bin
  ```

  For more information, see "Managing the Essbase Security File" in *Oracle Essbase Database Administrator's Guide*.

- Open the `essbase.cfg` file and set the `ENABLESWITCHTOBACKUPFILE` setting to `TRUE`. This setting allows Essbase to automatically use a backup security file. The `essbase.cfg` file is located in the following directories:

  ```
  (UNIX) APPLICATIONS_CONFIG_HOME/BIInstance/Essbase/essbaseserver_name/bin
  (Windows) APPLICATIONS_CONFIG_HOME\BIInstance\Essbase\essbaseserver_name\bin
  ```

  If you edit `Essbase.cfg`, restart Essbase to enable the change.

  For more information, see "Essbase.cfg Configuration Settings" in *Oracle Essbase Technical Reference*.

18.5.4.26 Status of Essbase Agent Connection

Problem

The Administrator does not know if Essbase is running.

Solution

Use the following OPMN command to determine if the Essbase agent is running:

```
(UNIX) APPLICATIONS_CONFIG_HOME/BIInstance/bin/opmnctl status
(Windows) APPLICATIONS_CONFIG_HOME\BIInstance\bin/opmnctl status
```

To start Essbase server using Fusion Applications Control:

1. From the navigation pane, expand the farm and then **WebLogic Domain** for the `BIDomain` domain.

2. Expand **Essbase Servers**, and then select the Essbase server to see if it is running.

3. If it is not running, From the **Essbase Server** menu, choose **Control**, then **Restart**.

18.5.4.27 Unable to Write File During Data Load or Building Aggregate Views

Problem

When attempting to load data into an Essbase aggregate storage database or while building aggregate views, the following error is displayed:

```
Failed to extend file [<path>/ess0001.dat]: a system file error
occurred. Please see application log for details.
```

Solution

This error indicates that the file system is out of space. Essbase can require significant temporary disk storage while performing a data load or aggregate view build. The "default" tablespace is the location where cube data is stored. The "temp" tablespace is

where Essbase writes temporary data while building the cube. By default, both tablespaces are on the disk drive where Essbase is installed.

During the operation, the space required by the temp tablespace is at least as big as the resulting change in the database size, so the total free space required is at least twice as big as the resulting change to the database size. For example, loading a database that is 1GB will require at least 2GB of free space. Building 10GB worth of aggregate views will require at least 20GB of free space. Note that after the operation is complete, the files created in the temp tablespace are deleted.

Use one of the following methods to correct this error:

- To correct this error, use MaxL statements to view or change the tablespace settings.

- If the tablespace locations do not have enough free space, consider moving one or both tablespaces to different disk drives, or limit the size of the existing file locations and add new locations on other disk drives. Note that you cannot remove a tablespace location that already contains data.

18.6 Using My Oracle Support for Additional Troubleshooting Information

You can use My Oracle Support (formerly MetaLink) to help resolve Oracle Fusion Applications problems. My Oracle Support contains several useful troubleshooting resources, such as:

- Knowledge base articles

- Community forums and discussions

- Patches and upgrades

- Certification information

You can access My Oracle Support at `https://support.oracle.com`.

19

Troubleshooting Oracle Enterprise Crawl and Search Framework

This chapter describes common problems that you might encounter when using Oracle Enterprise Crawl and Search Framework (ECSF) and explains how to solve them.

This chapter includes the following topics:

- Problems and Solutions
- Diagnosing Enterprise Crawl and Search Problems
- Using My Oracle Support for Additional Troubleshooting Information

In addition to this chapter, review the *Oracle Fusion Middleware Error Messages Reference* for information about the error messages you may encounter.

19.1 Problems and Solutions

This section describes common problems and solutions. It contains the following topics:

- Missing Parameters for Search Engine Instance
- Cannot Create Search Engine Instance
- Searchable Object Not Listed in Fusion Applications Control for ECSF
- Searchable Object Deployment Fails
- Searchable Object Undeployment Fails
- Search Category Deployment Fails
- Cannot Start, Undeploy, or Delete the Index Schedule
- Crawl Fails
- Invalid Channel Feed Type Error
- ECSF Query Error
- FND-6603 Error

19.1.1 Missing Parameters for Search Engine Instance

You do not see any parameters for the search engine instance you created.

Problem

Seed data is not set up on your Oracle Fusion Applications database.

Solution

Execute the `ECSF_SEED_DATA.sql` script against your Oracle Fusion Applications database.

19.1.2 Cannot Create Search Engine Instance

You are unable to create a search engine instance record.

Problem

An existing search engine instance record is present.

Solution

1. Execute the following SQL statement:

   ```
   select * from ECSF_ENGINE_INSTANCE
   ```

2. Delete the row that contains null values for all fields except the ID and Name fields, and commit the deletion.

3. Recreate the search engine instance.

19.1.3 Searchable Object Not Listed in Fusion Applications Control for ECSF

Your searchable object is not listed in the table of the Add a Searchable Object to Selected Engine Instance dialog.

Following are two possible causes and solutions for this issue:

Problem

The searchable object is not in the Oracle Fusion Applications database. You can check by executing the following SQL statement against your Oracle Fusion Applications database:

```
select * from ECSF_SEARCH_INDEX_OBJECT
```

Solution

Register the searchable object in the Oracle Fusion Applications database. For information, see Task 2, "Register the Searchable Objects".

Problem

The searchable object has already been associated with an engine instance. This is evident when a value exists in the `SEI_ID` column.

Solution

Disassociate the searchable object from the current engine instance. For information, see Task 3, "Disassociate the Searchable Objects from Search Engine Instances".

You can identify the engine instance with which the searchable object is associated by executing the following SQL statement:

```
select * from ECSF_ENGINE_INSTANCE where ID = sei_id_value_of_
your_searchable_object_record
```

Alternatively, you can browse through the engine instances in Fusion Applications Control for ECSF.

19.1.4 Searchable Object Deployment Fails

Your attempt to deploy a searchable object using Fusion Applications Control for ECSF fails.

Following are two possible causes and solutions for this issue:

Problem
The ECSF application, which is deployed to Oracle WebLogic Server, does not contain the search metadata file.

Solution
Make sure the ECSF application contains the search metadata file, and redeploy it to Oracle WebLogic Server. For information, see the "Deploying and Crawling Searchable Objects" chapter in *Oracle Fusion Applications Developer's Guide*.

Problem
The ECSF_RUNTIME_URL search engine instance parameter is not set, or is set incorrectly.

Solution
Make sure the ECSF_RUNTIME_URL parameter is set correctly.

19.1.5 Searchable Object Undeployment Fails

Your attempt to undeploy a searchable object from an engine instance fails.

Problem
The searchable object you are trying to undeploy does not exist in the Oracle SES data sources.

Solution
Make sure the searchable object you are trying to undeploy is deployed to an engine instance, as indicated when the DEPLOYED flag is set to TRUE, or the Deployed column is checked in Fusion Applications Control for ECSF.

19.1.6 Search Category Deployment Fails

Your attempt to deploy the search category to the engine instance fails.

Problem
There are no deployed searchable objects associated with the search category.

Solution
Make sure you associate searchable objects with the search category before you attempt to deploy the search category. For information, see Task 3, "Associate Searchable Objects with the Search Categories".

19.1.7 Cannot Start, Undeploy, or Delete the Index Schedule

You cannot start, undeploy, or delete the index schedule.

Problem

The index schedule is in `ERROR_MANUAL_RECOVERY` state, which results from restarting the database.

Solution

You must recover the stuck ESS request. For information, see the "Stuck Scheduling Requests" section in *Oracle Secure Enterprise Search Administrator's Guide*.

19.1.8 Crawl Fails

You can view the data in feeds, but you receive an error when you attempt to crawl the data into the search engine.

Problem

You did not specify a search plug-in for your searchable object, so the default security plug-in is being used. The default security plug-in requires you to select a secure attribute.

Solution

Select a secure attribute when you define searchable attributes. For information, see the "Creating Searchable Objects" chapter in *Oracle Fusion Applications Developer's Guide*.

19.1.9 Invalid Channel Feed Type Error

You receive the following error during a scheduled crawl:

```
Error: Invalid Channel Feed type "error".
```

Problem

When the crawl was scheduled to start, the crawl state of the searchable object was `CRAWLING`. ECSF allows a searchable object to be crawled by only one process at a time. If you are certain that the searchable object is not being crawled, then the last crawl was abnormally terminated, and the ECSF metadata was not properly updated to indicate that the crawl had concluded.

Solution

Manually start the crawl to reset the state of the crawl and enable the searchable object to be crawled. For information, see Task 5, "Start the Index Schedules".

19.1.10 ECSF Query Error

You receive the following error when you click on a facet value:

```
unexpected ECSF Query Error: oracle.ecsf.SearchException: Field
'facet_attribute' cannot be used for this query because it has not been
indexed in on the search engine
```

Problem

You deployed the searchable object to the Oracle SES instance, so it became available for search and the facet tree was rendered with its facet values. However, the crawl

had either failed or no documents were indexed, so Oracle SES does not contain the searchable attribute.

Solution

Determine why the crawl failed, undeploy the searchable object, and then redeploy it to the Oracle SES instance to generate a new index.

19.1.11 FND-6603 Error

You receive the FND-6603 error that halts search functionality.

Problem

The search categories could not be fetched, or results could not be obtained from all the selected search categories. There may be other causes for the error.

Solution

You can change the server startup parameters to increase the web service timeout value to 5 minutes (500,000 milliseconds). For information, see the "How to Set the System Parameter for Web Service" section in *Oracle Fusion Applications Developer's Guide*.

You can also undeploy the categories that are failing to load. For information, see Task 1, "Undeploy the Search Categories".

19.2 Diagnosing Enterprise Crawl and Search Problems

This section describes general approaches for diagnosing ECSF problems. It contains the following topics:

- Diagnosing ECSF Command Line Administration Utility Issues
- Diagnosing Fusion Applications Control for ECSF Issues
- Diagnosing Failures During Deploy/Undeploy Operations for Search Categories

19.2.1 Diagnosing ECSF Command Line Administration Utility Issues

To diagnose issues related to ECSF Command Line Administration Utility, you can check the log file (`ecsfCmdLineAdminLog.txt`) located in the log subdirectory where `runCmdLineAdmin.sh` or `runCmdLineAdmin.bat` file is executed.

19.2.2 Diagnosing Fusion Applications Control for ECSF Issues

To diagnose issues related to Fusion Applications Control for ECSF, you can also check for error messages in the log files located in *DOMAIN_HOME*/sysman/log.

If the parameters for the selected search engine instance are not configured correctly, a window with error messages appears. The error message contains three entries: a summary of the error, the specific error message from the runtime server, and the detailed stack trace information. The messages in the first and the second entries are internationalized (that is, they are translated to the language of the selected locale).

19.2.3 Diagnosing Failures During Deploy/Undeploy Operations for Search Categories

If you experience failures while deploying or undeploying search categories, check for the following:

- The database connection is successful.

- The searchable object is available in the Oracle Enterprise Manager repository.

- The associated searchable objects are deployed.

- The deployed object is available in Oracle SES.

19.3 Using My Oracle Support for Additional Troubleshooting Information

You can use My Oracle Support (formerly MetaLink) to help resolve Oracle Fusion Applications problems. My Oracle Support contains several useful troubleshooting resources, such as:

- Knowledge base articles

- Community forums and discussions

- Patches and upgrades

- Certification information

You can access My Oracle Support at `https://support.oracle.com`.

20

Troubleshooting Oracle Enterprise Scheduler

This chapter describes common problems that you might encounter when using Oracle Enterprise Scheduler and explains how to solve them.

This chapter contains the following topics:

- Introduction to Troubleshooting Oracle Enterprise Scheduler
- Getting Started with Troubleshooting Oracle Enterprise Scheduler Jobs
- Getting Started with Troubleshooting an Oracle Enterprise Scheduler Cluster
- Problems and Solutions
- Using My Oracle Support for Additional Troubleshooting Information

To gain insight into the log details generated by Oracle Fusion applications, see Chapter 17. Also, review the *Oracle Fusion Middleware Error Messages Reference* for information about the error messages you may encounter.

20.1 Introduction to Troubleshooting Oracle Enterprise Scheduler

This section provides guidelines and a process for using the information in this chapter. Using the following guidelines and process will focus and minimize the time you spend resolving problems.

Guidelines

When using the information in this chapter, Oracle recommends:

- After performing any of the solution procedures in this chapter, immediately retrying the failed task that led you to this troubleshooting information. If the task still fails when you retry it, perform a different solution procedure in this chapter and then try the failed task again. Repeat this process until you resolve the problem.

- Making notes about the solution procedures you perform, symptoms you see, and data you collect while troubleshooting. If you cannot resolve the problem using the information in this chapter and you must log a service request, the notes you make will expedite the process of solving the problem.

Process

Follow the process outlined in Table 20–1 when using the information in this chapter. If the information in a particular section does not resolve your problem, proceed to the next step in this process.

Table 20–1   Process for Using the Information in this Chapter

Step	Section to Use	Purpose
1	Section 20.2	Get started troubleshooting Oracle Enterprise Scheduler jobs.
2	Section 20.3	Get started troubleshooting an Oracle Enterprise Scheduler clustered environment.
3	Section 20.4	Perform problem-specific troubleshooting procedures. This section describes: ■ Possible causes of the problems ■ Solution procedures corresponding to each of the possible causes
4	Section 20.5	Use My Oracle Support to get additional troubleshooting information about Oracle Fusion Applications or Oracle Enterprise Scheduler. My Oracle Support provides access to several useful troubleshooting resources, including Knowledge Base articles and Community Forums and Discussions.
5	Section 20.5	Log a service request if the information in this chapter and My Oracle Support does not resolve your problem. You can log a service request using My Oracle Support at `https://support.oracle.com`.

20.2 Getting Started with Troubleshooting Oracle Enterprise Scheduler Jobs

You may want to troubleshoot the following typical issues that can arise when running Oracle Enterprise Scheduler jobs.

■ Asynchronous jobs remain in running state indefinitely.

■ An asynchronous job hangs or crashes.

■ Oracle Enterprise Scheduler is down when the remote scheduled job completes or there are network problems such that Oracle Enterprise Scheduler does not receive the completion status from the remote job.

■ A scheduled job is ready to execute, but does not execute.

■ A scheduled job is placed in manual error recovery state where troubleshooting is needed.

■ Oracle Enterprise Scheduler is throwing errors.

■ A scheduled job ends in error.

For troubleshooting Oracle Enterprise Scheduler, use the standard Oracle WebLogic Server system log. For information about viewing job request logs, see Section 5.15

This section contains the following topics:

■ Troubleshooting Asynchronous Scheduled Jobs

■ Troubleshooting Process or Spawned Jobs

■ Steps for Manual Recovery

■ Job Diagnostics

20.2.1 Troubleshooting Asynchronous Scheduled Jobs

Asynchronous jobs run on separate JVMS, including asynchronous Oracle BI Reporting and Publishing, PL/SQL and Java jobs that invoke asynchronous SOA or ADF Business Components services. When handling asynchronous scheduled jobs,

Oracle Enterprise Scheduler depends on the remote job sending a completion status defining the job outcome after running. However, the completion status may not be generated, or it may get lost for any of the following reasons:

- The scheduled job has crashed.

- The job is stuck in a hanging state.

- Oracle Enterprise Scheduler was down when the job completed.

- Network problems.

In any of these cases, an asynchronous job stays in running state indefinitely. As a result, subsequent steps in a job set may not execute, or an incompatible job may be blocked indefinitely.

Oracle Enterprise Scheduler displays information in Fusion Applications Control that enables you to locate the job on the remote system, including an external identifier.

You can take any of the following actions to troubleshoot an asynchronous job that is stuck in running state:

- Use the remote system to troubleshoot the job and determine the outcome of its execution. For more information, see Section 20.2.1.1, Section 20.2.1.2 and Section 20.2.1.3.

- Cancel the job. For more information, see Section 5.7.3.

- Configure a timeout for asynchronous jobs. For more information, see Section 5.7.1.

When configuring timeouts for jobs, you can use Fusion Applications Control to display all jobs that have timed out. However, a job that has timed out is still in running state. You must manually change the state of jobs that have timed out. Status callbacks are still accepted for timed out jobs and the job transitions to completion.

You can change the status of asynchronous jobs that have not timed out. This might happen if a timeout has not been configured, the completion status was lost and you notice that the job has been running for a long time.

20.2.1.1 Troubleshooting Asynchronous Java Jobs

In the case of asynchronous Java jobs (including jobs that invoke remote Oracle SOA Suite or Oracle ADF Business Components services), the log records are tagged with the ECID. You can view logs across the domain by ECID to troubleshoot the job execution. Oracle ADF Business Components, SOA composites, the web services stack and Oracle Fusion application log records with the ECID.

For asynchronous SOA jobs, the audit trail for the instance in Fusion Applications Control can be used to troubleshoot composite execution, as described in Section 20.4.4.

Example 20–1 shows a sample log message file the JRF web services stack in `server-name`-diagnostic.log, with inline comments for each log. This log is stored in the following directories:

```
(UNIX) DOMAIN_HOME/servers/server_name/logs
(Windows) DOMAIN_HOME\servers\server_name\logs
```

Example 20–1   Asynchronous job logging messages in the ADF Business Components log

Sending message to JMS queue "oracle.j2ee.ws.server.async.DefaultRequestQueue" for asynchronous processing of service b99d80e5-42aa-423a-9e98-f6f88b8b79dfRequest.

```
[2010-08-30T18:28:13.519-07:00] AdminServer NOTIFICATION []
oracle.j2ee.ws.common.jaxws.JAXWSMessages [tid: ACTIVE.ExecuteThread: '0' for queue:
'weblogic.kernel.Default (self-tuning)'] [userId: <anonymous>]
[ecid: 0000If5kZGS76EWLHyo2yf1CV5eh000000,0:1] [WEBSERVICE\_PORT.name:
 EmployeeModuleServiceSoapHttpPort] [APP: ADFBCAsync] [J2EE\_MODULE.name: ADFBCAsync\-ejb]
[WEBSERVICE.name: EmployeeModuleService] [J2EE\_APP.name: ADFBCAsync] [MessageID:
urn:uuid:01994234-4442-4dee-82a6-e1e04407af56]
```

An asynchronous request message is received and successfully recorded for the service
EmployeeModuleService with the replyTo address
http://adc2180314:7001/ADFBCAsyncCallback/EmployeeModuleServiceCallbackResponseImplService.

```
[2010-08-30T18:28:13.724-07:00] AdminServer NOTIFICATION []
oracle.j2ee.ws.common.jaxws.JAXWSMessages [tid: ACTIVE.ExecuteThread: '0' for queue:
'weblogic.kernel.Default (self-tuning)'] [userId: <anonymous>]
[ecid: 0000If5kZGS76EWLHyo2yf1CV5eh000000,0:1] [WEBSERVICE\_PORT.name:
EmployeeModuleServiceSoapHttpPort] [APP: ADFBCAsync] [J2EE\_MODULE.name: ADFBCAsync\-ejb]
[WEBSERVICE.name: EmployeeModuleService] [J2EE\_APP.name: ADFBCAsync] [MessageID:
urn:uuid:01994234-4442-4dee-82a6-e1e04407af56]
Unknown macro: {/service/common/}
```

Started asynchronous request processing for the service EmployeeModuleService with the message
selector "b99d80e5-42aa-423a-9e98-f6f88b8b79dfRequest". Transaction enabled: "true".

```
[2010-08-30T18:28:13.811-07:00] AdminServer NOTIFICATION []
oracle.j2ee.ws.common.jaxws.JAXWSMessages [tid: ACTIVE.ExecuteThread: '2' for queue:
'weblogic.kernel.Default (self-tuning)'] [userId: OracleSystemUser]
[ecid: 0000If5kZGS76EWLHyo2yf1CV5eh000000,0] [APP: ADFBCAsync] [MessageID:
urn:uuid:01994234-4442-4dee-82a6-e1e04407af56]
```

Completed asynchronous request processing. A response is sent to the client.

```
[2010-08-30T18:28:17.307-07:00] AdminServer NOTIFICATION []
oracle.j2ee.ws.common.jaxws.JAXWSMessages [tid: ACTIVE.ExecuteThread: '2' for queue:
'weblogic.kernel.Default (self-tuning)'] [userId: OracleSystemUser]
[ecid: 0000If5kZGS76EWLHyo2yf1CV5eh000000,0] [APP: ADFBCAsync] [MessageID:
urn:uuid:01994234-4442-4dee-82a6-e1e04407af56]
```

Started asynchronous response processing for the service EmployeeModuleService with the message
selector "b99d80e5-42aa-423a-9e98-f6f88b8b79dfResponse".

```
[2010-08-30T18:28:17.330-07:00] AdminServer NOTIFICATION []
oracle.j2ee.ws.common.jaxws.JAXWSMessages [tid: ACTIVE.ExecuteThread: '0' for queue:
'weblogic.kernel.Default (self-tuning)'] [userId: OracleSystemUser]
[ecid: 0000If5kZGS76EWLHyo2yf1CV5eh000000,0] [APP: ADFBCAsync] [MessageID:
urn:uuid:01994234-4442-4dee-82a6-e1e04407af56]
Unknown macro: {/service/common/}
```

Completed asynchronous response processing successfully. The client should have received the
response at this point.

```
[2010-08-30T18:28:17.825-07:00] AdminServer NOTIFICATION []
oracle.j2ee.ws.common.jaxws.JAXWSMessages [tid: ACTIVE.ExecuteThread: '0' for queue:
'weblogic.kernel.Default (self-tuning)'] [userId: OracleSystemUser]
[ecid: 0000If5kZGS76EWLHyo2yf1CV5eh000000,0] [APP: ADFBCAsync] [MessageID:
urn:uuid:01994234-4442-4dee-82a6-e1e04407af56]
```

Completed asynchronous response processing with exceptions. The client does not receive any
response.

```
[2010-08-31T09:55:33.939-07:00] AdminServer ERROR [] oracle.j2ee.ws.common.jaxws.JAXWSMessages
[tid: ACTIVE.ExecuteThread: '5' for queue: 'weblogic.kernel.Default (self-tuning)'] [userId:
OracleSystemUser] [ecid: 0000If94mKW76EWLHyo2yf1CVJG1000000,0] [APP: j2wpojoasync] [MessageID:
urn:uuid:5b9a5134-1416-4bda-95fd-da6e624466d7]
```

*The response will not be sent again as the callback service has replied with a SOAP fault. The HTTP
response code is 500.*

```
[2010-08-31T10:17:35.852-07:00] AdminServer ERROR [] oracle.j2ee.ws.common.jaxws.JAXWSMessages
[tid: ACTIVE.ExecuteThread: '0' for queue: 'weblogic.kernel.Default (self-tuning)'] [userId:
OracleSystemUser] [ecid: 0000If99plA76EWLHyo2yf1CVJ^i000000,0] [APP:
ADFBCAsyncInvalidCallbackCreds] [MessageID: urn:uuid:7e355428-f68f-4c3f-a368-baf529048323]
```

For more information about configuring parameters for a job, see Section 5.7.1. For
more information about viewing the *server-name*-diagnostic.log file, see
"Managing Log Files and Diagnostic Data" chapter in the *Oracle Fusion Middleware
Administrator's Guide*.

20.2.1.2 Troubleshooting Asynchronous PL/SQL Jobs

PL/SQL jobs can be identified in the Oracle Database by their job names. You can find
the database job names associated with the request on the Job Details page in Fusion
Applications Control. For more information about viewing job request details, see
Section 5.7.3.

20.2.1.3 Troubleshooting Asynchronous Oracle Business Intelligence Publisher Jobs

When viewed in Fusion Applications Control, a given Oracle Business Intelligence
Publisher job includes a direct URL pointing to the Oracle Business Intelligence
Publisher server on the Job Details page. Once the Oracle Business Intelligence
Publisher job starts running on the Oracle Business Intelligence Publisher server, the
scheduled job request attains the property value bip.status_url. This property
value holds the URL of the Oracle Business Intelligence Publisher server which is used
to diagnose Oracle Business Intelligence Publisher report execution. For more
information about viewing job request details, see Section 5.7.3.

20.2.2 Troubleshooting Process or Spawned Jobs

Spawned job processes can be identified by a combination of host name, process ID
and process group ID. The Job Details page in Fusion Applications Control shows this
and additional information. For more information about viewing job request details,
see Section 5.7.3.

20.2.3 Steps for Manual Recovery

Job requests may require manual error recovery for a number of reasons. For example,
in the case of an asynchronous job request, the job implementation cannot know
whether the job request launched successfully and an error manual recovery exception
is thrown. When a job request is placed in manual error recovery, the job request must
be manually completed. If the job request actually was successfully launched and
returns a completion status on completion, this status is ignored.

A job request may be placed in manual error recovery for the following reasons:

- An asynchronous Java job request throws an
  ExecutionManualRecoveryException, which indicates to Oracle Enterprise

Scheduler that manual recovery is necessary. The job request is placed in `ERROR_MANUAL_RECOVERY` state. The cause is set to `Cause.PROCESS_MANUAL_RECOVER_ERROR (209)`.

- An asynchronous Java job request throws a runtime exception or an error. If a thrown exception is not handled by the job implementation, Oracle Enterprise Scheduler cannot know whether the remote job was invoked, such that manually recovery is required. The cause is set to `Cause.PROCESS_MANUAL_RECOVER_ERROR (209)`.

- If Oracle Enterprise Scheduler crashes after executing an asynchronous job request, but before the job returns and is marked in the database as complete, Oracle Enterprise Scheduler cannot know whether the remote job has been invoked. The job request transitions to `ERROR_MANUAL_RECOVERY`, which holds onto incompatibility locks. The cause is set to `Cause.PROCESS_RECOVER (210)`.

- A spawned job runs in a clustered environment, the job request runs on the first instance of Oracle Enterprise Scheduler, which goes down, along with the associated Perl agent. If the first instance of Oracle Enterprise Scheduler will not be back up and running for a while, Oracle Enterprise Scheduler does not know whether the spawned process is actually still running. Manual detection and recovery are required. The cause is set to `Cause.PROCESS_RECOVER(210)`.

20.2.3.1 Handling Synchronous Java Jobs Requiring Manual Recovery

If a Java job times out or runs for too long, you can either let it run to completion as usual, or attempt to terminate it.

To recover synchronous Java jobs:

1. If the job is not in the `ERROR_MANUAL_RECOVERY` state, cancel the job. See Section 5.7.5. If the job remains in `CANCELLING` state for an unreasonable amount of time, continue to the next step.

2. Take one of the following actions:

   - If, after some time, the request transitions to a terminal state, no other intervention is required.

   - If the request remains in `CANCELLING` state, then determine the Oracle Enterprise Scheduler server on which the request is running by finding the server and cluster names.

     Determine the cluster name from the Oracle Enterprise Scheduler process group. Process group information is shown in search results only when the scope for the job request search is **All Scheduling Services sharing the ESS repository**. See Section 5.7.2 for information on performing searches.

3. If canceling the job is not effective, restart the relevant Oracle Enterprise Scheduler server:

   a. From the navigation pane, expand the farm and then **WebLogic Domain**.

   b. Expand the Oracle Enterprise Scheduler cluster and select the Oracle Enterprise Scheduler server.

   c. In the WebLogic Server home page, from the **WebLogic Server** menu, choose **Control** > **Shut Down**.

   d. After the server is shut down, from the WebLogic Server home page, from the **WebLogic Server** menu, choose **Control** > **Start Up**.

4. In the Request Details page, from the **Action** menu, select **Recover Stuck Request**.

20.2.3.2 Handling Stuck Asynchronous Jobs Requiring Manual Recovery

When an asynchronous job request requires manual recovery, follow these basic steps. Additional steps will depend on the job type.

If a request is stuck (marked for manual recovery or taking too long or timed out) and it is an asynchronous remote job, you first check to see if the remote job is still running.

To handle asynchronous jobs:

1. Check to see if the remote job is still running:

   a. Identify the remote job by navigating to the Request Details page. See Section 5.7.3 for information about how to view the Request Details page.

      - For PL/SQL jobs, navigate to the Request Details page and click the **Execution Type** icon. The database session information displays, correlating the scheduled job request with the database scheduler job. This indicates whether the job is still executing.

      - For spawned jobs, navigate to the Request Details page, click the **Execution Type** icon to display the spawned process information. This correlates the scheduled job request with the non-Oracle Enterprise Scheduler job, in this case the operating system process. This correlation indicates whether the job is still executing.

   b. Verify that the remote job is no longer running.

      - If the remote job was not successfully created on the remote system, set the status of the job request to `ERROR`.

      - If the remote job was created and has finished executing, determine its status and set the status of the job request accordingly.

      - If the remote job instance has not finished executing, wait until it completes and set the job request status accordingly.

2. Once the remote job is no longer running, terminate the job request in Oracle Enterprise Scheduler, so that Oracle Enterprise Scheduler is no longer is keeping track of the job.

   a. Navigate to the Search request page by clicking the **Scheduling Service** menu and selecting **Search Job Request**.

   b. From the **Quick Search** list, for asynchronous jobs marked for manual recovery (`ERROR_MANUAL_RECOVERY`), select **Asynchronous requests that need manual recovery**. Requests will be in `RUNNING` state. For asynchronous jobs not in the `ERROR_MANUAL_RECOVERY` state, rather than search with the **Asynchronous requests that need manual recovery** option, search for the known requests that need to be updated.

   c. In the search results, click the request ID to display the Request Details page.

   d. In the Request Details page, from the **Action** menu, choose **Recover Stuck Request**.

   e. In the Recover Stuck Request dialog box, set the state accordingly for the job request. Optionally, add a description for the status of the job request.

      If you set the status to `ERROR`, the description you add displays in the Request Details page.

20.2.4 Job Diagnostics

A scheduled job may not execute for a number of reasons, or it may fail. Either way, Fusion Applications Control provides built-in diagnostics in the Job Details page. For jobs that fail with an error, the Job Details page displays the reason and provides access to the job request log.

Fusion Applications Control provides the following:

- Access logs for the job request.

- Database session information for PL/SQL job requests, shown in the Job Details page for a PL/SQL job request.

- Spawned process information for process job requests, shown in the Job Details page for process job requests.

- A message displays in the Job Details page for job requests in wait, ready or blocked state specifying the reason the job request is in that state.

- Error and warning messages display in the Job Details page. Additional details also display, such as stack traces and so on.

- For retried job requests, the Job Details page displays the number of times the job request was retried, the time of the next retry and the number of additional times the job request is to be retried in the event of an error.

- For job requests that require manual recovery or have timed out, the Job Details page displays a message regarding a need for manual recovery.

For more information about viewing job request details, see Section 5.7.3.

Table 20–2 shows the associated diagnostic codes for each state along with a description and additional information that is provided. If a request is in a state that does not appear in the table, its diagnosis will contain only the request state.

Table 20–2   Job Request States and Associated Diagnostic Codes

State	Diagnostic Code	Description	Related Documentation
BLOCKED	BLOCKED	Blocked due to incompatible job request or requests. Includes the request ID of the blocking request.	For information about cancelling a job request, see Section 5.7.5.
COMPLETED	POSTPROCESS_DELAY	The job request is delayed by the post-processor. Includes the time at which the delay ends.	
PAUSED	PAUSED	The job request is the parent of one or more subrequests and has been paused. Includes the request ID of a subrequest in a non-terminal state, if there is one.	For information about resuming a paused job request, see Section 5.7.4.
READY	NO_ACTIVE_SERVER	No server is active in the process group. Includes the name of the process and isolation groups.	For information about activating an Oracle Enterprise Scheduler, see Section 5.6.
READY	REQUESTED_ PROCESSOR_NOT_ ACTIVE	The server specified by the job request SYS_requestedProcessor property is not available. Includes the name of the requested processor.	
READY	NO_APPLICATION	The application is either not deployed or not active. Includes the name of the application, process group and isolation group.	

Table 20–2   (Cont.) Job Request States and Associated Diagnostic Codes

State	Diagnostic Code	Description	Related Documentation
READY	PROCESSOR_STOPPED	The request cannot be processed because there is no server with the application deployed with an active processor.	For information about starting a request processor, see Section 5.6.2.
READY	PROCESSOR_FAILED	The request cannot be processed because the processor has failed on all servers to which the application is deployed.	
READY	PROCESSOR_WAIT	The job request is waiting for an available processor thread. Includes the name of a work assignment that could process the job request.	
READY	INACTIVE_WORK_ ASSIGNMENT_WAIT	Waiting for a work assignment to become active. Includes the name of an inactive work assignment that could process the job request if it were active.	
READY	NO_LOADED_WORK_ ASSIGNMENT	There is a bound work assignment that could process the request, but the binding is not loaded. The server may be down. Includes the name of a work assignment that could process the request but is not loaded.	For information about activating an Oracle Enterprise Scheduler, see Section 5.6.
READY	NO_BOUND_WORK_ ASSIGNMENT	You must bind a specialized work assignment to the job request.	For information about binding a work assignment, see Section 5.4.2.
READY	THROTTLED	The job request is asynchronous and the number of active asynchronous jobs of the same type is at the allowed limit. Includes the work assignment, workshift, the asynchronous job type and the asynchronous limit.	Change the number of threads allocated to jobs, or the asynchronous job limits for the work shift. For more information about editing a work shift, see Section 5.11.1.1.
READY	DISABLED_WORK_ ASSIGNMENT	There is a bound work assignment that could process the request, but the work assignment is disabled. Includes the name of the disabled work assignment.	Enable the work assignment. For more information about editing a work assignment, see Section 5.11.1.1.
RUNNING	PREPROCESS_DELAY	The request is delayed by the pre-processor. Includes the time at which the delay ends.	
RUNNING	TIMED_OUT	The job request has timed out. Typically, this code displays for timed out asynchronous Java job requests.	
RUNNING	NO_ACTIVE_SERVER	The job request processing cannot continue because no server is active in the process group. Includes the name of the process and isolation groups.	Verify that the Oracle Enterprise Scheduler server is running (see Section 5.8). If necessary, restart one of the Oracle Enterprise Scheduler components, as described in Section 5.6.

Table 20–2 (Cont.) Job Request States and Associated Diagnostic Codes

State	Diagnostic Code	Description	Related Documentation
RUNNING	REQUESTED_ PROCESSOR_NOT_ ACTIVE	Job request processing cannot continue because the server specified by the job request SYS_requestedProcessor property is not available. Includes the name of the requested processor.	Verify that the Oracle Enterprise Scheduler server is running (see Section 5.8). If necessary, restart one of the Oracle Enterprise Scheduler components, as described in Section 5.6.
RUNNING	NO_APPLICATION	Job request processing cannot continue because the application is either not deployed or not active. Includes the name of the application, process group and isolation group.	
RUNNING	PROCESSOR_STOPPED	Job request processing cannot continue because there is no server with the application deployed with an active processor.	Verify that the Oracle Enterprise Scheduler server is running (see Section 5.8). Start a request processor, as described in Section 5.6.2.
RUNNING	PROCESSOR_FAILED	Job request processing cannot continue because the processor has failed on all servers to which the application is deployed.	Verify that the Oracle Enterprise Scheduler server is running (see Section 5.8). Restart the request processor, as described in Section 5.6.2.
RUNNING	PROCESSOR_WAIT	Waiting for an available processor thread to handle the update event for an asynchronous Java job request. Includes the name of a work assignment that could process the request; indicates whether the job request has disabled update events.	
RUNNING	INACTIVE_WORK_ ASSIGNMENT_WAIT	Waiting for a work assignment to become active to handle the update event for an asynchronous Java request. Includes the name of an inactive work assignment that could process the request if it were active; indicates whether the job request has disabled update events.	Activate the inactive work assignment so that the job request can be processed. For more information, see Section 5.11.1.1.
RUNNING	NO_LOADED_WORK_ ASSIGNMENT	Processing of an asynchronous Java job request may be delayed. There is a bound work assignment that could process the update event, but the binding is not loaded. The server may be down. Includes the name of a work assignment that could process the job request but is not loaded; indicates whether the job request has disabled update events.	Verify that the Oracle Enterprise Scheduler server is running (see Section 5.8). If necessary, restart one of the Oracle Enterprise Scheduler components, as described in Section 5.6.
RUNNING	NO_BOUND_WORK_ ASSIGNMENT	Processing of an asynchronous Java job request may be delayed. There is no bound work assignment that can handle the update event for the asynchronous Java job request; indicates whether the request has disabled update events.	Bind a work assignment to the request processor. For more information, see Section 5.4.2.

Table 20–2   (Cont.) Job Request States and Associated Diagnostic Codes

State	Diagnostic Code	Description	Related Documentation
RUNNING	DISABLED_WORK_ ASSIGNMENT	Processing of an asynchronous Java job request may be delayed. There is a bound work assignment that could process the update event for the job request, but the work assignment is disabled. Includes the name of the disabled work assignment; indicates whether the job request has disabled update events.	Enable the work assignment. For more information about editing a work assignment, see Section 5.11.1.1.
WAIT	FUTURE_START	The job request has a scheduled time in the future. Includes the scheduled time for the job request.	
WAIT	RETRY_DELAY	An error occurred in the job request, and the job request has been delayed before being retried. Includes the scheduled retry time.	
WAIT	NO_APPLICATION	The job request scheduled time has been reached, but the job request cannot be dispatched because the application is not available (is either not deployed or not active). Includes the name of the application, process group and isolation group.	
WAIT	PARENT_NOT_PAUSED	The job request is a subrequest whose parent has not paused. The subrequest remains in WAIT state until its parent pauses. Includes the parent request ID.	For information about pausing a paused job request, see Section 5.7.4.
WAIT	DEFERRED	The job request is an instance of a recurrence for which the previous recurrence instance is still running. Oracle Enterprise Scheduler prevents concurrent execution of recurrent instances, and the next recurrence remains in WAIT state while the current recurrence is active.	
WAIT	DISPATCHER_ STOPPED	The job request cannot be dispatched because there is no server with the application deployed with a running dispatcher.	Restart the request dispatcher. For more information, see Section 5.6.2
WAIT	DISPATCHER_FAILED	The job request cannot be dispatched because the dispatcher has failed on all servers to which the application is deployed.	Restart the request dispatcher or other Oracle Enterprise Scheduler components. For more information, see Section 5.6.2 and Section 5.6.1.
terminal state	TERMINAL	The job request is in a terminal state (SUCCEEDED, WARNING, ERROR, CANCELLED, EXPIRED, VALIDATION_ FAILED, FINISHED). Includes the application, process group, isolation group, work assignment and workshift.	

> **Note:** For more information about job states, see *Oracle Fusion Applications Java API Reference for Oracle Enterprise Scheduler Service.*

20.3 Getting Started with Troubleshooting an Oracle Enterprise Scheduler Cluster

Troubleshooting an Oracle Enterprise Scheduler clustered environment involves the following.

- Finding Performance and Scalability Issues
- Using a Shared Database
- Tuning Oracle Enterprise Scheduler System Performance

20.3.1 Finding Performance and Scalability Issues

It is possible to detect any issues with performance and scalability problems by viewing performance metrics in Fusion Applications Control. Metrics include Oracle WebLogic Server, JVM-level metrics and plots, as well as Oracle Enterprise Scheduler level metrics.

The system-level tools specific to the operating system running on the server can be used as an additional diagnostic tool. System level tools indicate how machine resources are utilized at various times, such as network, memory, CPU, I/O, and so on. Such system tools are especially useful for tuning job implementations. Database tools enable identifying problems in the database. For remote jobs, such as ADF Business Components or SOA Java jobs running on a remote system, you can use the corresponding Fusion Applications Control and system level tools on those servers.

Fusion Applications Control provides the following types of Oracle Enterprise Scheduler metrics to help identify problems:

- Completed job request statistics by job name. Shows run count, run time, success rate and last run job request statistics for completed requests by job name.
- Completed job request statistics by user. Shows completed request count and run time for completed requests by user.
- Completed job request statistics by work assignment. Shows wait time, processing time, completed and failed count for completed job requests by work assignment.
- Completed job request count by status. Displays completed job requests in a variety of terminal states.

For more information about metrics, see Section 9.7.

20.3.2 Using a Shared Database

A common database can be used across multiple Fusion Applications domains. In this case, the database may be loaded from multiple sources. To help with this, Oracle Enterprise Manager allows you to see running and waiting jobs, as well as metrics across the database from multiple Oracle Enterprise Scheduler systems.

20.3.3 Tuning Oracle Enterprise Scheduler System Performance

The following potential performance and scalability issues may occur in the context of job requests or Oracle Enterprise Scheduler runtime.

- Jobs are saturating the CPU of Oracle Enterprise Scheduler servers.

- Jobs are overloading the remote systems where asynchronous jobs are running.

- Ready jobs are filling the queue, despite the availability of spare CPU power, such that job output is delayed.

- Multiple domains are sharing a database. A great number of concurrently running database-intensive jobs across domains are slowing down the database.

- Performance and scalability are affected by the running of large jobs at the end of a financial quarter or month.

- Performance and scalability are affected by the concurrent running of two or more jobs that interact very intensively with the database.

Tuning involves changing job implementations or changing schedules, Oracle Enterprise Scheduler cluster size, processor bindings for work assignments, throttling and thread limits.

20.3.3.1 Tuning Clusters

Clusters are the basic mechanism for enabling scalability and high availability. When a job runs, it is equally likely to run on any processor on which it is eligible to run at that time. By carefully controlling the size of the Oracle Enterprise Scheduler cluster, it is possible to better distribute job executions across the cluster so that servers do not become overloaded. In the case of remote jobs, the jobs actually execute on a remote system and consume very few resources in the cluster (except for a blocked thread for synchronous jobs). If jobs running locally are overloading the system, the first step is to revisit the cluster size configuration.

20.3.3.2 Processor Bindings

Some jobs can physically execute only on a given server; these jobs have been bound to run only on that server. If too many jobs are bound to a particular processor, the benefits of a clustered environment are effectively moot. For the purposes of high availability, avoid tying a job to just one server while enabling the job to run on at least two servers. Otherwise, the job will not run if the bound processor is down.

Rather than relying on clusters to randomly distribute work, you may have a set of long, resource intensive jobs to be run locally within a given scheduling window. In this case, you can bind jobs to specific processors and explicitly control the distribution of these jobs.

A processor is an Oracle Enterprise Scheduler instance. One Oracle Enterprise Scheduler instance runs on one cluster node. As one cluster node typically runs on a single computer, a processor normally equates to a computer.

At times, a clustered environment runs well until the scheduled periods during which a number of resource intensive jobs run. In order to maintain performance, you can configure the cluster with extra idle nodes that are activated during busy periods so as to handle the extra job load. You use standard Oracle WebLogic Server cluster methodologies to enable this clustering. For more information about Oracle WebLogic Server clustering, see the "High Availability for WebLogic Server" chapter in the *Oracle Fusion Middleware High Availability Guide*.

Job performance can vary depending on whether executing jobs are synchronous or asynchronous. Synchronous jobs consume a single thread throughout their execution, and are normally short lived. (An exception might be a process or spawned job that loads a database.) Asynchronous jobs consume a thread at the beginning and end of execution for a very short time, but they otherwise run independently. Asynchronous

jobs are typically long running and continue execution across server restarts. Typical examples of asynchronous jobs are PL/SQL jobs, Oracle BI Publisher jobs that generate reports, Java jobs that invoke remote asynchronous ADF Business Components services, and Java jobs that invoke remote asynchronous SOA services.

Throttling limits the maximum number of jobs that may execute concurrently. This is important to avoid flooding the system with too many concurrently running jobs. For synchronous jobs, this limit is imposed by limiting the number of threads available for execution. For PL/SQL jobs and other asynchronous jobs, this limit is imposed by defining a maximum concurrency limit for PL/SQL and asynchronous jobs respectively. Asynchronous throttling limits are set on the work assignment to which the jobs are assigned. For more information about setting asynchronous job limits, see Section 5.11.2.1

It is possible that all threads configured are used up for synchronous jobs thereby blocking asynchronous jobs from starting. This can be avoided if asynchronous and synchronous jobs are not combined in a single work assignment.

20.3.3.3 Using Job Incompatibility to Manage Performance

You can configure a job incompatibility not only to prevent two incompatible jobs from running, but in order to prevent both intensive jobs from heavily loading the same resource. In order to maintain good performance, you can define an incompatibility for such jobs so that they never run at the same time. For more information about defining a job incompatibility, see Section 5.10.3.2.

20.3.3.4 Tuning Oracle Enterprise Scheduler for Optimal Performance

You can tune the following Oracle Enterprise Scheduler components:

- Request dispatcher

- Request processor

- Connection pool size

- RDBMS Scheduler

Tuning the Dispatcher

The dispatcher tuning parameters apply to the Oracle Enterprise Scheduler request dispatcher. The request dispatcher manages requests that are awaiting their scheduled execution. The request processor handles the job requests once they have executed.

Parameters are as follows.

- **Dispatcher Enabled**: Indicates whether the request dispatcher is enabled on the Oracle Enterprise Scheduler server. When disabled, that Oracle Enterprise Scheduler server will not dispatch job requests whose scheduled execution time has arrived. By default, this parameter is enabled.

- **Maximum Poll Interval**: Specifies the maximum frequency, in seconds, at which the request dispatcher checks for job requests that are ready to be dispatched. The default value is 15 seconds.

Tuning the Processor

The processor tuning parameters apply to the Oracle Enterprise Scheduler request processor. The request processor manages job requests whose scheduled execution time has arrived, and are ready to execute.

Parameters are as follows.

- **Processor Enabled:** Indicates whether the request processor is enabled on the Oracle Enterprise Scheduler server. If disabled, the Oracle Enterprise Scheduler server will not process requests that are ready to be executed. By default, this parameter is enabled.

- **Maximum Processor Threads:** Specifies the maximum number of threads used to process job requests. This represents the total number of worker threads that might run concurrently for all active work assignments for the Oracle Enterprise Scheduler server. By default, this parameter is set to 25.

- **Starvation Threshold:** Indicates the wait time, in minutes, before a job request that is ready to be executed will be considered starved and eligible to be processed by a starvation worker thread. The starvation worker processes only those job requests that have been ready longer than the starvation time. A starvation worker is not created if the threshold value is equal to zero.

  If enabled (meaning the parameter value is greater than zero), a starvation worker thread is created for each active work assignment for the Oracle Enterprise Scheduler server. The Maximum Processor Threads parameter does not apply to starvation workers. By default, the value of this parameter is set to zero, such that no starvation worker is created.

Tuning the Connection Pool Size for the Oracle Enterprise Scheduler Internal Data Source

The connection pool size for the Oracle Enterprise Scheduler internal JDBC data source should be based on the request processor tuning values configured for the Maximum Processor Threads and Starvation Threshold parameters.

The recommended pool size if the Starvation Threshold parameter is disabled (its value is equal to zero) is the number of maximum processor threads plus twenty.

The recommended pool size if the Starvation Threshold parameter is enabled (its value is greater than 0) is the number of maximum processor threads, along with the number of bound work assignments plus twenty.

Tuning the RDBMS Scheduler

The RDBMS scheduler is capable of auto-turning. To enable auto-tuning, set `job_queue_processes` to 0. Leave `JOB_QUEUE_PROCESSES` to the default value at 1000. For more information about the `JOB_QUEUE_PROCESSES` parameter, see the *Oracle Database Reference*.

20.3.3.5 Tuning Dead Database Connections

Oracle Enterprise Scheduler spawned jobs connect to the database using SQL*Net. If the spawned jobs are canceled, Oracle Enterprise Scheduler kills these processes at the operating system level. It is possible, however, that the database connections used by these processes still exist in the database.

To reduce dead connections in the database, use the `SQLNET.EXPIRE_TIME` configuration option by setting this value to the desired value. For more information about the `SQLNET.EXPIRE_TIME` parameter, see the "Parameters for the sqlnet.ora File" chapter in *Oracle Database Net Services Reference*.

20.4 Problems and Solutions

This section describes common problems and solutions for Oracle Enterprise Scheduler. It contains the following topics:

- Job Remains in WAIT State

- Synchronous Job Continues in RUNNING State for Too Long

- Asynchronous Jobs Remain in RUNNING State and Do Not Complete

- Asynchronous Java SOA Job Remains In RUNNING State

- Asynchronous Java Oracle ADF Business Components Job Remains In RUNNING State

- Asynchronous PL/SQL Job Remains in RUNNING State

- BI Publisher Job Remains In RUNNING State

- BI Publisher Job States Mismatch with Oracle Enterprise Scheduler States

- Job Does Not Execute at Scheduled Time

- Job Never Execute and Go into ERROR State Immediately

- Asynchronous Java Job Requires Manual-Error Recovery

- Spawned (Process Type) Job Requires Manual Error Recovery

- Job Remains in CANCELLING State

- Job Completes, but Goes into a WARNING state

- Newly Added Server Is Not Being Utilized or Running Inappropriate Jobs

- Oracle Enterprise Scheduler Runtime System Is Throwing Errors

- Metadata Access Denied Error When Accessing Job Metadata

- Oracle Enterprise Scheduler Is Running Out Of Database Connections

- Insufficient Privilege Error on Request Error

- Empty Process List in Schedule New Process Dialog

- Slow Display of Job Definition List of Values in Schedule New Process Dialog

- Job Queue Full Due to a Hanging Job

In addition to the recommended solutions, consider reviewing Section 20.3.3 for tuning tips.

20.4.1 Job Remains in WAIT State

Problem

When a user submits a job, the job can remain in the WAIT state for too long without progressing to the RUNNING state.

Solution

To resolve this problem, verify the current status of Oracle Enterprise Scheduler from Fusion Applications Control:

1. Verify the request processor and request dispatcher are running:

   a. From the navigation pane, expand the farm, and then **Scheduling Services**.

   b. Select the **ESSAPP** application for the appropriate Managed Server.

   c. In the Scheduling Service home page, in the **Scheduler Components** section, ensure the Request Processor has a status of **Started**.

If it is not running, start it. See Section 5.6.2.

    d. In the Scheduling Service home page, in the **Scheduler Components** section, ensure the Request Dispatcher has a status of **Started**.

    If it is not running, start it Section 5.6.2.

2. Verify the `ESSAPP` application is running:

    a. From the navigation pane, expand the farm, and then **Scheduling Services**.

    b. Select the **ESSAPP** application for the appropriate Managed Server.

    c. In the Scheduling Service home page, in the **Scheduler Components** section, ensure the Request Dispatcher has a status of **Started**.

    d. In the WebLogic Server home page, in the **Deployments** section, ensure the `ESSAPP` applications is running.

    If it is not running, start it. See Section 5.6.1.

3. Check if concurrency or threads are configured is too small by looking at processor and work assignment configuration.

    ■ From the **Scheduling Service** menu, choose **Request Processor** > **Configure** to review the **Thread Count** field in the Configure Request Processor page. See Section 5.4.2.

    ■ From the **Scheduling Service** menu, choose **Work Allocation** > **Work Assignments** to review the configuration in the Work Assignments page. See Section 5.11.1.

20.4.2 Synchronous Job Continues in RUNNING State for Too Long

Problem

When the user submits a job, it remains in the `RUNNING` state for too long.

Solution

The job may be in `RUNNING` state because the Oracle Enterprise Scheduler server has crashed and recovery has not taken place.

To resolve this problem, determine the current status of Oracle Enterprise Scheduler from Fusion Applications Control:

1. Verify the request processor is running:

    a. From the navigation pane, expand the farm and then **Scheduling Services**.

    b. Select the **ESSAPP** application for the appropriate Managed Server.

    c. In the Scheduling Service home page, in the **Scheduler Components** section, ensure the **Request Processor** is enabled and is started.

    If it is not running, start it. See Section 5.6.2.

2. Verify the Oracle Enterprise Scheduler server is running and start if it is not running:

    a. From the navigation pane, expand the farm and then **WebLogic Domain**.

    b. Select the Oracle Enterprise Scheduler cluster.

    c. From the WebLogic Cluster page, in the **Servers** section, view the **Status** column to determine if the Oracle Enterprise Scheduler server is running.

    d. If it shows a status of down (red down arrow) for a server, click the server name in the **Name** column.

    e. In the WebLogic Server home page, from the **WebLogic Server** menu, choose **Control** > **Start Up**.

3. Look at the job output to see if the job is making progress. See Section 5.7.3.

When an Oracle Enterprise Scheduler server is restarted, synchronous jobs running on that server are transitioned to the ERROR state.

20.4.3 Asynchronous Jobs Remain in RUNNING State and Do Not Complete

Problem

Oracle BI Reporting and Publishing, PLSQL, and Java jobs that invoke asynchronous SOA services and Java jobs that invoke asynchronous Oracle Application Development Framework (Oracle ADF) Business Component services run on separate Java Virtual Machines (JVMs) or machines. In these cases, Oracle Enterprise Scheduler depends on the remote job sending a completion status at end of processing that defines the job outcome. However this message may never get generated or lost, resulting in the job staying in the RUNNING state. In addition, subsequent steps in a job set may not execute, or an incompatible job may be blocked indefinitely.

Solution

To resolve this issue, follow the actions described in Section 20.2.1.

20.4.4 Asynchronous Java SOA Job Remains In RUNNING State

Problem

Jobs that invoke asynchronous SOA services run on separate Java Virtual Machines (JVMs) or machines. In these cases, Oracle Enterprise Scheduler depends on the remote job sending a completion status at end of processing that defines the job outcome. However this message may never get generated due to various reasons.

Solution

To resolve this issue, you must troubleshoot the native job.

To resolve this problem for asynchronous SOA jobs:

1. Search for the request, as described in Section 5.7.2.

2. On the Request Details page, from the **Action** menu, select **Request Log** to view the log message. See Section 5.7.3 for further information about the Request Details page.

The Log Message page displays. By default when user navigates to view the logs for a request, only messages that are logged in Oracle Enterprise Scheduler cluster scope are shown. (If the ESSAPP application is not deployed to a cluster the messages that are logged in the Managed Server scope are shown). However, Oracle Enterprise Scheduler propagates the ECID associated with the request across subsystems, such as Oracle SOA Suite and Oracle ADF.

3. Make note of the value in the **ECID** field.

4. From the **Broaden Target Scope** list, select the */farm_name/domain_name* **(Oracle WebLogic Domain)** to view messages across the domain.

5. In the Log Messages page for the Oracle WebLogic Server domain, in the **Selected Targets** section, ensure the search includes the **ECID** field with the value from the Request Details page.

6. Search and view log records for Oracle SOA Suite and the ECID and note any issues. Section 5.15.4.

7. View the audit trail for an SOA composite application instance using the ECID:

    a. From the navigation pane, expand the farm, **SOA**, and then **soa-infra**.

    b. From the SOA Infrastructure page, click the **Instances** tab.

    c. In the **Search** section, enter the ECID in the **ECID** field.

    d. Click **Search** to find the instance with the ECID.

    e. Select the instance by clicking the ID in the **Instance ID** field from the Instances table.

       The Flow Trace page displays.

    f. View the audit trail for the instance and observe if the composite completed successfully or completed with error. See the "Viewing the Audit Trail and Process Flow of a BPEL Process Service Component" section in the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*. for more information about the Flow Trace window.

8. If the SOA composite is complete and the job is still running in Oracle Enterprise Scheduler, manually complete the job in the Request Details page. For related information, see Section 20.2.1.

20.4.5 Asynchronous Java Oracle ADF Business Components Job Remains In RUNNING State

Problem

Jobs that invoke asynchronous Oracle ADF Business Components services run on separate JVMs or computers. In these cases, Oracle Enterprise Scheduler depends on the remote job sending a completion status at end of processing that defines the job outcome. However this message may never get generated due to various reasons.

Solution

To resolve this issue, you must troubleshoot the native job.

To resolve this problem for synchronous Oracle ADF Business Components jobs:

1. Search for the request, as described in Section 5.7.2.

2. On the Request Details page, from the **Action** menu, select **Request Log** to view the log message. See Section 5.7.3 for further information about the Request Details page.

   The Log Message page displays. By default when user navigates to view the logs for a request, only messages that are logged in Oracle Enterprise Scheduler cluster scope are shown. (If the ESSAPP application is not deployed to a cluster the messages that are logged in the Managed Server scope are shown). However, Oracle Enterprise Scheduler propagates the ECID associated with the request across subsystems, such as Oracle SOA Suite and Oracle ADF.

3. Make note of the value in the **ECID** field.

4. From the **Broaden Target Scope** list, select the */farm_name/domain_name* **(Oracle WebLogic Domain)** to view messages across the domain.

5. In the Log Messages page for the Oracle WebLogic Server domain, in the **Selected Targets** section, ensure the search includes the **ECID** field with the value from the Request Details page and the **Component Name** field.

6. Search and view log records for Oracle ADF Business Components and web services stack for the ECID and note any issues. See the "Viewing and Searching Log Files" section in the *Oracle Fusion Middleware Administrator's Guide.*

7. Observe if the Oracle ADF Business Components completed successfully or completed with error.

8. If the service is complete and the job is still running in Oracle Enterprise Scheduler, manually complete the job in the Request Details page. For related information, see Section 20.2.1.

20.4.6 Asynchronous PL/SQL Job Remains in RUNNING State

Problem 1

PL/SQL jobs run on separate machines. In these cases, Oracle Enterprise Scheduler depends on the remote job sending a completion status at end of processing that defines the job outcome. However this message may never get generated due to various reasons.

Solution 1

PL/SQL jobs can be identified in the Oracle Enterprise Scheduler by their job names. Job definition names are available from the Request Details page in the Fusion Applications Control associated with the request. See Section 5.7.3.

To resolve this issue, troubleshoot the native job. For more information, see Section 20.2.1.2.

Problem 2

Oracle Enterprise Scheduler internally uses the Database Management System (DBMS) scheduler to schedule PL/SQL jobs. In some cases, the DBMS scheduler has not scheduled the job request, even though Oracle Enterprise Scheduler has submitted the job to the DBMS scheduler and set its state to RUNNING.

Solution 2

To resolve this issue:

1. Verify the DBMS Scheduler resource usage. See the "Administering Oracle Scheduler" chapter in the *Oracle Database Administrator's Guide*.

2. Change the PL/SQL job throttle limit by configuring PL/SQL job limits. See Section 5.11.2.

20.4.7 BI Publisher Job Remains In RUNNING State

Problem

Oracle Business Intelligence Publisher jobs run on separate computers. In these cases, Oracle Enterprise Scheduler depends on the remote job sending a completion status at end of processing that defines the job outcome. However this message may never get

generated due to various reasons. It is required to troubleshoot the native job implementation to resolve.

Solution

To resolve this issue, you must troubleshoot the native job. See Section 20.4.8 and Section 20.2.1.3.

20.4.8 BI Publisher Job States Mismatch with Oracle Enterprise Scheduler States

Problem

Under normal circumstances, the BI Publisher server state should match the Oracle Enterprise Scheduler job state. However, in case of any BI Publisher job errors, server issues or other network issues, the BI Publisher server state may not match the Oracle Enterprise Scheduler job state.

Solution

It is important to understand that the Oracle Enterprise Scheduler runtime invokes the Oracle BI Publisher web service asynchronously and then needs to wait for it to successfully complete to update the Oracle Enterprise Scheduler job status. Therefore, between the initial web service invocation, there could be various things that could go wrong. The table below lists the typical job life cycle, the potential failure points, and the Oracle Enterprise Scheduler and BI Publisher state values seen at each of those states. Use the table below to debug any state mismatch between the Oracle Enterprise Scheduler job and the BI Publisher server:

Processing Step/Error	Server	Expected ESS Request State	Expected BI Publisher State	Description
1. Oracle Enterprise Scheduler request is in the `RUNNING` state, calls BI Publisher job	Oracle Enterprise Scheduler	`RUNNING`	None	
2. BI Publisher Oracle Enterprise Scheduler job begins to execute.	Oracle Enterprise Scheduler	`RUNNING`	None	
3. BI Publisher Oracle Enterprise Scheduler job invokes BI Publisher web service.		`RUNNING`	None	
3a. Error calling the web service, for example, the BI Publisherserver is down or crashes while processing the web service.	Oracle Enterprise Scheduler	`ERROR_ MANUAL_ RECOVERY`	Could be `None` or `RUNNING`, depending on where the crash occurs	An error message is seen in the Oracle Enterprise Scheduler server log and with Oracle Enterprise Scheduler request in Fusion Applications Control. BI Publisher job state depends on why the web service failed. This requires manual recovery.
4. BI Publisher web service creates `"BIP Job ID"` with `RUNNING` status, enqueues to JMS, and commits.	BI Publisher	`RUNNING`	`RUNNING`	
5. BI Publisher Oracle Enterprise Scheduler job finishes after web service processing.	Oracle Enterprise Scheduler	`RUNNING`	`RUNNING`	

Processing Step/Error	Server	Expected ESS Request State	Expected BI Publisher State	Description
5a. Error in the BI Publisher Oracle Enterprise Scheduler job before returning.		`ERROR_ MANUAL_ RECOVERY`	`RUNNING`	Error message seen in the Oracle Enterprise Scheduler server log and with Oracle Enterprise Scheduler request in Fusion Applications Control. BI Publisher job is running, but Oracle Enterprise Scheduler job that invoked it has failed. This needs manual recovery because the Oracle Enterprise Scheduler job cannot be in a terminal `SUCCESS` or `ERROR` state while the BI Publisher job is `RUNNING`.
5b. Oracle Enterprise Scheduler server crash anywhere during BI Publisher Oracle Enterprise Scheduler job execution (crash during BI Publisher Oracle Enterprise Scheduler job executable call between Steps 2-5).		`ERROR_ MANUAL_ RECOVERY`	Could be `None` or `RUNNING`, depending on where the crash occurs	Error message seen in the Oracle Enterprise Scheduler server log and with Oracle Enterprise Scheduler request in Fusion Applications Control. It is not known if the web service has been invoked or not, so manual intervention is required.
6. BI Publisher server begins processing the BI Publisher job	BI Publisher	`RUNNING`	`RUNNING`	
6a. BI Publisher server crashes during processing and restarts		`RUNNING` to `ERROR` (See steps 11 through 13)	`RUNNING` (before restart) to `ERROR` (after restart)	
6b. Job is taking a long time to run		`RUNNING`	`RUNNING`	BI Publisher "request job history" UI should show which stage of processing the job is in. When BI Publisher server restarts, it will mark the BI Publisher state as `ERROR` and then notify Oracle Enterprise Scheduler that the job had an error.
BI Publisher report is successful				
7. BI Publisher job completes successfully, BI Publisher sets the BI Publisher job to `SUCCESS` and commits	BI Publisher	`RUNNING`	`SUCCESS`	
8. BI Publisher invokes Oracle Enterprise Scheduler web service to notify Oracle Enterprise Schedulerof job completion.	BI Publisher	`RUNNING`	`SUCCESS`	
8a. Oracle Enterprise Scheduler server is down while BI Publisher makes web service call to notify Oracle Enterprise Scheduler of job completion, other Oracle Enterprise Scheduler servers in the cluster are running.	BI Publisher	See Steps 9 through 10.	`SUCCESS`	Fusion Applications Control should show that the Oracle Enterprise Scheduler server is down. Service invocation should be routed successfully to another Oracle Enterprise Scheduler server in the cluster.
8b. All Oracle Enterprise Scheduler servers are down while BI Publisher makes web service call to notify Oracle Enterprise Scheduler of job completion.	BI Publisher	`RUNNING`	`SUCCESS`	Can detect this case after the Oracle Enterprise Scheduler servers are restarted and the BI Publisher request is `SUCCESS` while the Oracle Enterprise Scheduler request is `RUNNING` (or timed out). Job will remain `RUNNING` and administrator has to cancel and recover request in Fusion Applications Control.

Processing Step/Error	Server	Expected ESS Request State	Expected BI Publisher State	Description
9. Oracle Enterprise Scheduler calls job post-processor.	Oracle Enterprise Scheduler	COMPLETED	SUCCESS	
9a. Oracle Enterprise Scheduler job post-processor has an error.	Oracle Enterprise Scheduler	WARNING	SUCCESS	Error message seen in the Oracle Enterprise Scheduler server log and with ESS request in Fusion Applications Control.
10. Oracle Enterprise Scheduler completes the request after post-processing finishes successfully	Oracle Enterprise Scheduler	SUCCESS	SUCCESS	
BI Publisher report has error.				Error message seen in the Oracle Enterprise Scheduler server log and with Oracle Enterprise Scheduler request in Fusion Applications Control as well as in BI Publisher console for the job.
11. BI Publisher job completes with error, BI Publisher sets the BI Publisher job to ERROR and commits	BI Publisher	RUNNING	ERROR	
12. BI Publisher invokes Oracle Enterprise Scheduler web service to notify Oracle Enterprise Scheduler of job completion	BI Publisher	RUNNING	ERROR	
12a. Oracle Enterprise Scheduler server is down while BI Publisher makes web service call to notify Oracle Enterprise Scheduler of job completion, other Oracle Enterprise Scheduler servers in the cluster are running.	BI Publisher	RUNNING	ERROR	See Steps 8a and 8b. BI Publisher retries delivery of the status update in all cases if Oracle Enterprise Scheduler web service invocation fails
13. Oracle Enterprise Scheduler updates request status in the Oracle Database to ERROR	Oracle Enterprise Scheduler	ERROR	ERROR	
Request is cancelled.				**Note:** The timing refers to when the AsyncCancellable method is invoked, not when cancel is initiated. The states in this section assume that if the BI Publisher runs, it runs successfully unless canceled. For more details on Oracle Enterprise Scheduler cancellation, see Section 5.7.5.
14. Request in RUNNING state canceled prior to BI Publisher Oracle Enterprise Scheduler job execution.		CANCELLING to CANCELLED	None	Process phase prior to ExecuteInitiate. Should automatically transition from CANCELLING to CANCELLED.
15. Request canceled during BI Publisher Oracle Enterprise Scheduler job execution before BI Publisher job begins.		CANCELLING to CANCELLED	None	BI Publisher Oracle Enterprise Scheduler job should prevent BI Publisher job from being initiated by having checkpoints in the execute method. On return, request should transition to CANCELLED.
16. Request canceled during BI Publisher Oracle Enterprise Scheduler job execution after BI Publisher job initiated.		CANCELLING to CANCELLED	RUNNING to CANCELLED	BI Publisher Oracle Enterprise Scheduler job execute method can ignore cancel. Cancel() method should attempt to cancel BI Publisher job. BI Publisher job returns status (CANCELLED or other). Request ends up in CANCELLED state.

Processing Step/Error	Server	Expected ESS Request State	Expected BI Publisher State	Description
17. Request canceled while BI Publisher job is running.		CANCELLING to CANCELLED	RUNNING to CANCELLED	Cancel() method should attempt to cancel BI Publisher job. BI Publisher job returns status (CANCELLED or other). Request ends up in CANCELLED state.
18. Request canceled after BI Publisher job, before callback received (when is BI Publisher state set to SUCCESS.)		CANCELLING to CANCELLED	SUCCESS	Cancel() method should attempt to cancel BI Publisher job. BI Publisher job returns status (CANCELLED or other). Request ends up in CANCELLED state.
19. Request canceled after callback received, before post-processor		CANCELLING to CANCELLED	SUCCESS	
20. Request canceled while post-processor is running		SUCCESS	SUCCESS	Once post-processing is begun, the cancel operation has no effect.
21. Request canceled after post-processing		SUCCESS	SUCCESS	

20.4.9 Job Does Not Execute at Scheduled Time

Problem

When a job's scheduled time arrives, it does not execute.

Solution

To resolve this problem, view the Request Details page in the Fusion Applications Control. This page provides built in diagnostics to show what the issue is. For jobs that fails with an error, the Request Details page shows the reason and provides access to the job request log from the **Action** menu. See Section 5.7.3.

For more information job diagnostics, see Section 20.2.4.

20.4.10 Job Never Execute and Go into ERROR State Immediately

Problem

When the user submits a spawn job, it goes into ERROR state immediately.

Solution

To resolve this problem, verify that the RequestFileDirectory directory is created and set up on the server. This directory can be located anywhere. The Oracle WebLogic Server writes and reads from it, so the Oracle Fusion Middleware administration account server can access this directory. If this directory does not exist, all jobs will move to ERROR state. The directory is a shared file system shared across the Oracle Enterprise Scheduler cluster.

To find the RequestFileDirectory value in the Fusion Applications Control:

1. Search the logs to diagnose any specific issues found. See Section 20.4.4 and Section 20.4.5.

2. Verify that RequestFileDirectory directory is created:

   a. From the navigation pane, expand the farm, **WebLogic Domain**, and select the Oracle Enterprise Scheduler server.

    b. From the **WebLogic Server** menu, choose **System MBean Browser**.

    c. In the System MBean Browser page, expand **Application Defined MBeans**.

    d. Expand **oracle.adf.share.connections**, **Server:** *ess_server_name*, **Application: ESSAPP**, **ADFConnections**, **ADFConnections**, **EssConnection**.
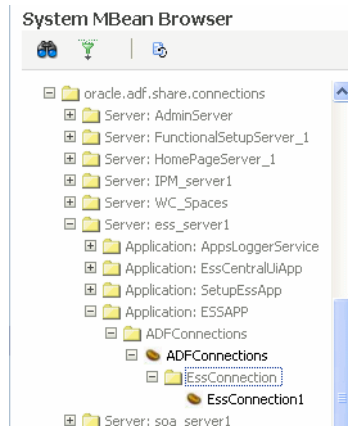
    e. Click **EssConnection1**.

    f. In the Application Defined MBeans: EssConnection:EssConnection1 page, view the attribute value for **RequestFileDirectory**.

20.4.11 Asynchronous Java Job Requires Manual-Error Recovery

Problem

A job gets placed in the ERROR_MANUAL_RECOVERY state. There are a number of reasons why a job may end up in error manual recovery. For example for an asynchronous job, the job implementation may not know if the job was successfully launched due to an error and throws the error manual recovery exception. For more reasons on why a job can end up in error manual recovery, see Section 20.2.3.

Solution

To resolve this issue, manually update the job status to complete it. See Section 20.2.3.1 and Section 20.2.3.2.

20.4.12 Spawned (Process Type) Job Requires Manual Error Recovery

Problem

A spawned process type job gets placed in the ERROR_MANUAL_RECOVERY state.

Solution

To solve this issue, transition the request to a terminal state:

1. Identify the spawned host and process ID in the Request Details page. See Section 5.7.3 for more information about the Request Details page.

2. If the process is still running on the host, wait for it to complete or terminate it.

3. When the process is no longer running, recover the request to begin the transition to an error state. It will be subject to auto-retries if configured. From Fusion Applications Control, perform the following steps:

    a. Search for the request, as described in Section 5.7.2. When selecting a value for **Status**, select **Cancelled** and **Error Manual Recovery**.

    b. On the Request Details page, from the **Action** menu, select **Recover Stuck Request**. See Section 5.7.3 for more information about the Request Details page.

    A request in the `CANCELLED` state is put to `CANCELLED` state and a request in `ERROR_MANUAL_RECOVER_STATE` is put to `ERROR` state, specifying an appropriate error message. The error message specified by the user will be shown on Request Details page.

    c. Manually update the job status to complete it. See Section 20.2.3.1 and Section 20.2.3.2.

For more information on manually recovering a job, see Section 20.2.3

20.4.13 Job Remains in CANCELLING State

Problem

Sometimes when a job is cancelled, it stays in `CANCELLING` state and does not get cancelled. Results of a cancellation request depends on the stage of processing for the request when the cancel happens and the results of that stage. Many jobs in Oracle Fusion Applications are implemented as asynchronous ADF service invocations. The Oracle Fusion Applications infrastructure does not support cancellation of in-flight service requests, and as a result, the job does not cancel as expected. In addition, there are additional reasons why a job may get stuck in `CANCELLING` state.

- For PL/SQL jobs, Oracle Enterprise Scheduler will attempt to kill the RDBMS scheduler job. For spawned process, Oracle Enterprise Scheduler will try to kill the running process. If the job is successfully killed, the request will transition to `CANCELLED` state. If the job completes before it can be killed, the state to which the request will transition depends on the result of the job execution. For these type of jobs, this issue should not occur.

- Asynchronous Java job: The request was canceled, but the remote job never contacted Oracle Enterprise Scheduler with its terminal status. This could happen if the job is still executing because either `AsyncCancellable` interface was not implemented or the remote cancel operation did not succeed. It could also happen if the remote system is unable to respond.

- Synchronous Java job: The request was canceled and the job is still executing. This could happen if either `Cancellable()` interface was not implemented or the job's `Executable.execute()` method still did not return after `Cancellable.cancel` was invoked. For more information about the `Executable` interface, see the chapter "Use Case Oracle Enterprise Scheduler Sample Application" in *Oracle Fusion Applications Developer's Guide for Oracle Enterprise Scheduler*. For more information about the Cancellable interface, see the section "How to Cancel a Scheduled Java Job" in the chapter "Working with Extensions to Oracle Enterprise Scheduler" in Oracle Fusion Applications Developer's Guide for Oracle Enterprise Scheduler.

- Job sets (Parent-Child Requests): In cases of job sets, the cancellation operation propagates to all eligible child requests. Until all child requests are completed, the parent request will remain in `CANCELLING` state.

- SOA Java job: In cases where an job ends in an error, look for the ECID in the Fusion Applications Control by locating the composite instance and looking at the audit trail and logs tagged with the ECID to see what happened. For more

information about finding the ECID, see the "Viewing the Audit Trail and Process Flow of a BPEL Process Service Component" section in the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

- Oracle ADF Business Component Java Job: In cases where requests end in an error, look for the message ID in the log file for the message ID in the `server-name`-`diagnostic.log` file in the following directories to see what happened:

```
(UNIX) DOMAIN_HOME/servers/server_name/logs
(Windows) DOMAIN_HOME\servers\server_name\logs
```

For more information about viewing log files, see the "Viewing and Searching Log Filess" section in the *Oracle Fusion Middleware Administrator's Guide.*

Solution

Manually intervene in the Fusion Applications Control to complete the job stuck in `CANCELLING` state.

To recover asynchronous requests that are stuck, determine whether the remote job is still executing. If it is, nothing should be done on the Oracle Enterprise Scheduler side. If the remote job is no longer executing, then perform the following from Fusion Applications Control to complete the request:

1. Search for the request, as described in Section 5.7.2.

2. On the Request Details page, from the **Action** menu, select **Recover Stuck Request**. See Section 5.7.3 for more information about the Request Details page.

For information about debugging any state mismatch between the Oracle Enterprise Scheduler job and the BI Publisher server, see Section 20.4.8.

For synchronous Java jobs, wait for the job to complete. If the job is irrevocably hung, then the server on which it is executing must be restarted.

20.4.14 Job Completes, but Goes into a WARNING state

Problem

Oracle Enterprise Scheduler tries to upload log/output files to Oracle Universal Content Management. If upload fails, the request will be marked to `WARNING` state.

Solution

To resolve this problem, use Fusion Applications Control:

1. Check that the Oracle Universal Content Management Content Server is up and running.

   Oracle Universal Content Management is located in the `CommonDomain` domain in the Oracle Fusion Setup product family.

   a. From the navigation pane, expand the farm, **Content Management**, and then **Content Server**.

   b. Select the **Oracle Universal Content Management - Content Server** application for the appropriate Managed Server.

   c. In the UCM home page, in the **Scheduler Components** section, ensure the Request Processor has a status of **Started**.

If it is not running, start it. See Section 5.6.2.

    d. In the Scheduling Service home page, in the **General** section, ensure the Content Server has a state of Active.

    e. If the state is not **Active**, from the **UCM** menu, choose **Control** > **Start**.

2. Search and view log records for issues related to Oracle Universal Content Management or attachments. See the "Viewing and Searching Log Files" section in the *Oracle Fusion Middleware Administrator's Guide.*

3. Check the attachments configuration for the job hosting application is correct:

    a. From the navigation pane, expand the farm, **Application Deployments**.

    b. Expand *domain_name***EssApp**, and then select *domain_name***EssApp**.

    The Application Deployment page displays.

    c. From the **Application Deployment** menu, choose **WebCenter** > **Service Configuration**.

    The WebCenter Service Configuration page displays.

    d. Click **Content Repository**.

    e. In the Manage Content Repository Connection sections, click **Edit** to view and modify the entry for the **FusionAppsContentRepository** connection.

20.4.15 Performance and Scalability Goes Down When Two Very Database Intensive Jobs Run at the Same Time

Problem

Performance and scalability goes down when two very database intensive jobs run at the same time.

To determine job performance issues:

1. Use operating-system level or Database Management System (DBMS) tools to check performance metrics across the Oracle Database to ensure the performance bottleneck is isolated to these two jobs running simultaneously and that tuning practices have been attempted.

2. Determine the execution time of the job:

    a. From the navigation pane, expand the farm, and then **Scheduling Services**.

    b. Select the **ESSAPP** application for the appropriate Managed Server.

    c. In the Scheduling Service home page, click the **Top 10 Long Running Job Requests** tab to view the jobs.

Solution

To solve this problem, perform one of two solutions:

■ Mark the two very database intensive jobs as incompatible. Sometimes we want to ensure two jobs never run at the same time, not because they are incompatible and would corrupt the system if run simultaneously, but because they heavily load the same resource. In this case, the jobs can be defined to be incompatible so they never run at the same time. See Section 5.10.3.

■ Schedule the jobs at different times. See Section 5.7.1.

20.4.16 Newly Added Server Is Not Being Utilized or Running Inappropriate Jobs

Problem

A newly added Managed Server does not run jobs as expected or is not being utilized at all.

You will see in the job history over time that no jobs are running on this server or you notice in the job history that a particular job is running on this server that should not run.

To view the job history:

1. From the navigation pane, expand the farm, and then **Scheduling Services**.

2. Select the **ESSAPP** application for the appropriate Managed Server.

3. You will see in the job history over time that no jobs are running on this server.

4. Search for the request, as described in Section 5.7.2, to view the list of jobs that already have executed.

5. From the **Request ID** column in the Request Search page, click a job to go to the Request Details page for the job.

6. In the **Execution Trail** section of the Request Details page, view the **Dispatcher**, **Processor**, **Work Assignment**, and **Workshift** the job ran on.

Solution

After a new server is added, Oracle Enterprise Scheduler determines if the default work assignment can be used based on how other processors are bound. If it cannot use the default work assignment, it configures the new server to only run the health check service (internal work assignment `ESSInternalWA`). Revisit this default configuration to configure the work assignment binding to this server as desired and removing the internal work assignment after the health check is complete. For more details, see Section 5.4.1.

20.4.17 Oracle Enterprise Scheduler Runtime System Is Throwing Errors

Problem

Oracle Enterprise Scheduler runtime system is not behaving properly and is throwing errors or encountering problems when processing a job.

Solution

To identify and solve this problem, review the Oracle Enterprise Scheduler system logs to troubleshoot this issue. From Fusion Applications Control, perform the following:

1. Search for the request, as described in Section 5.7.2.

2. On the Request Details page, from the **Action** menu, select **Request Log** to view the log message. See Section 5.7.3 for further information about the Request Details page.

3. To adjust the log levels:

    a. From the navigation pane, expand the farm, **WebLogic Domain**, Oracle Enterprise Scheduler cluster, and select the Oracle Enterprise Scheduler server (for example, **ess_server1**).

    The WebLogic Server home page displays.

> **b.** From the **WebLogic Server** menu, select **Logs** > **Log Configuration** to display the Log Configuration page.

You can configure the Oracle Enterprise Scheduler server loggers for an Oracle WebLogic Server by modifying the `logging.xml` file of that Oracle WebLogic Server. By default, there is no explicit logger entry for the Oracle Enterprise Scheduler and it will inherit 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 server logger can be found in the Oracle WebLogic Server diagnostic log file for that Oracle WebLogic Server. For more information on logging and log levels, see Section 5.15.4.

> **Note:** The logger 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.

20.4.18 Metadata Access Denied Error When Accessing Job Metadata

Problem

When accessing job metadata, users receive a metadata access denied error in the Standard Report Submission, when there is an attempt to submit a job.

Solution

To resolve this problem,

1. Determine the application role that is supposed to have metadata permissions on the job. See the "Mapping Application Roles to an External Role" section in the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)*.

2. From Fusion Applications Control, perform the following steps to verify the permissions:

   a. In the navigation pane, expand the farm and then **WebLogic Server Domain**.

   b. Select the domain.

   c. From the **WebLogic Domain** menu, choose **Security** > **Application Policies**.

   The Application Policies page displays.

   d. In the **Search** section, choose the application or application stripe to search, enter the data to match (a principal name or a permission name or both), and click the blue **Search application security grants** icon. In the results table at the bottom of the page, search the grants for the application role and see if the permissions are granted. If

   e. Add the permissions, as described in the "Managing Application Policies" section of the *Oracle Fusion Middleware Application Security Guide*.

   f. Determine the enterprise role that is supposed to map to the application role. 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)*.

3. From the Oracle WebLogic Server Administration Console, perform the following to steps to determine the enterprise role of the user.

   a. From the left pane, from the **Domain Structure**, select **Security Realms**.

   b. On the Summary of Security Realms page select the name of the realm.

   c. On the Settings for Real Name page, click the **Users and Groups** tab to check the user's group.

      If you do not find the user, add the user to the group, as described in Step 4.

4. Add the user to the group:

   a. In the **Users** table, select the user you want to add to a group.

   b. On the Settings for User Name page select **Groups**.

   c. Select a group or groups from the Available list box:

      - To locate a group in a large list, type the first few characters of the name.

      - To select multiple groups, Ctrl-click each group.

      - To add a user to a group, click the right arrow to move the selection to the **Chosen** list box.

      - To remove a user from a group, select the group in the **Chosen** list box and click the left arrow.

   d. Click **Save**.

20.4.19 Oracle Enterprise Scheduler Is Running Out Of Database Connections

Problem

If Oracle Enterprise Scheduler is running out of database connections, there could be a problem with connection leaks in Oracle Enterprise Scheduler.

Description

See the "Running out of Data Source Connections" section in the *Oracle Fusion Middleware Administrator's Guide*.

20.4.20 Insufficient Privilege Error on Request Error

Problem

The following error is reported:

```
User name does not have sufficient privilege to do name operation on request
number
```

Solution

To resolve this problem,

1. Check who submitted the request. A user should be able to operate on the requests submitted. If the login user is not the request submitter, go to Step 2.

2. Use Oracle Authorization Policy Manager to check if the submitter is assigned privileges on the request. See the "Locating Policies Associated with a Database Resource" chapter in the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)*. When using Oracle Authorization Policy Manager, in the **Search** area of the **Manage Database**

Resources and Policies tab, search for the ESS_REQUEST_HISTORY database resource.

3. Search the policies for the application role that is expected to receive the grants. See "Searching Database Resources" chapter in the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)*.

4. If the policy or condition is missing, use Oracle Authorization Policy Manager to add them. See the "Managing Database Resource Conditions" chapter in the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)*.

5. Use Oracle Authorization Policy Manager to check if the submitter is assigned privileges on the request. See the "Searching Database Resources" chapter in the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)*.

20.4.21 Empty Process List in Schedule New Process Dialog

Problem

When an application user logs into an Oracle Fusion application to schedule a new Oracle Enterprise Scheduler job by using the following procedure and finding no value for the process:

1. Click the **Navigator** link.

2. Choose **Tools** > **Scheduled Processes**.

3. In the Scheduled Processes page, in the **Search Results** section, click **Schedule New Process**.

4. In the Schedule New Process dialog, click the arrow button next to the **Process Name** list and find no values.

This issue is usually the result of a permissions problem.

Solution

To resolve this problem,

1. Check who submitted the request. A user should be able to operate on the requests submitted. If the login user is not the request submitter, go to Step 2.

2. Use Oracle Authorization Policy Manager to check if the submitter is assigned privileges on the request. See the "Locating Policies Associated with a Database Resource" chapter in the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)*. When using Oracle Authorization Policy Manager, in the **Search** area of the **Manage Database Resources and Policies** tab, search for the ESS_REQUEST_HISTORY database resource.

3. Search the policies for the application role that is expected to receive the grants. See the "Searching Database Resources" chapter in the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)*.

4. If the policy or condition is missing, use Oracle Authorization Policy Manager to add them. See the "Managing Database Resource Conditions" chapter in the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)*.

5. Use Oracle Authorization Policy Manager to check if the submitter is assigned privileges on the request. See "Searching Database Resources" chapter in the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)*.

20.4.22 Slow Display of Job Definition List of Values in Schedule New Process Dialog

Problem

When an application user logs into an Oracle Fusion application to schedule a new Oracle Enterprise Scheduler job by using the following procedure and finding slow performance in the list of values.

1. Click the **Navigator** link.

2. Choose **Tools** > **Scheduled Processes**.

3. In the Scheduled Processes page, in the **Search Results** section, click **Schedule New Process**.

4. In the Schedule New Process dialog, click the arrow button next to the **Process Name** list and discover the performance is unacceptably slow.

Solution

To resolve this issue, update the MDS table statistics for the query optimizer to use optimal query plans:

1. Generate Oracle Enterprise Scheduler MDS schema statistics:

```
BEGIN
dbms_stats.gather_schema_stats(
ownname =>'FUSION_MDS',
METHOD_OPT =>'FOR ALL COLUMNS SIZE AUTO',
CASCADE =>TRUE,
ESTIMATE_PERCENT =>NULL);
END;
```

2. Flush the shared pool:

```
SQL> ALTER SYSTEM FLUSH SHARED_POOL;
```

20.4.23 Job Queue Full Due to a Hanging Job

Problem

At times, some jobs may not behave as expected. The job queue may become large due to some job spinning or hanging or if a job has memory leaks.

Solution

Here are some typical scenarios:

- Java job goes into an infinite loop and there are other jobs waiting in the queue for it to finish.

  To resolve this problem, perform the following with Fusion Applications Control:

  1. From the navigation pane, expand the farm, and then **Scheduling Services**.

  2. Select the **ESSAPP** application for the appropriate Managed Server.

3. From the **Scheduling Service** menu, choose **Request Processor** > **Configure** to review the **Thread Count** field in the Configure Request Processor page. See Section 5.4.2.

If the **Thread Count** field is set to 25 synchronous Java jobs, only 25 Java jobs are permitted to be in the RUNNING state, and all other Java jobs are in the queue have to wait to be processed.

If some jobs are performing heavy processing or seem to hang, you can isolate such jobs by defining dedicated work assignments to process them on a specific Oracle Enterprise Scheduler server, leaving the other servers to process the rest of the jobs. See Section 5.11.1.1.

4. Restart the Oracle Enterprise Scheduler server:

    a. From the navigation pane, expand the farm and then **WebLogic Domain**.

    b. Expand the Oracle Enterprise Scheduler cluster and select the Oracle Enterprise Scheduler server.

    c. In the WebLogic Server home page, from the **WebLogic Server** menu, choose **Control** > **Shut Down**.

    d. After the server is shut down, from the WebLogic Server home page, from the **WebLogic Server** menu, choose **Control** > **Start Up**.

    After restarting the server, the Oracle Enterprise Scheduler moves RUN-NING jobs to the ERROR state and starts processing the next batch of jobs in the queue.

- PL/SQL job goes into infinite loop and never exits. Many other PL/SQL jobs are also submitted.

  Cancel the PLSQL job to move the job to CANCELLED state. See Section 5.7.5.

- At times, heavy database inserts and updates will cause jobs to wait.

  For example, consider a PL/SQL procedure that performs 1 million inserts, updates, and deletes on a table. Such as PL/SQL procedure can take about 20 minute to two hours to complete, depending on the DBMS load. If a work Assignment named wa1 has been created with a PLSQL concurrency limit setting of 25. There are many PL/SQL jobs and each job runs this PL/SQL procedure.

  In this case, the Oracle Enterprise Scheduler server will process only 25 PL/SQL jobs concurrently. Since each PL/SQL job is taking one to two hours to complete, the remaining PL/SQL jobs will be in the WAIT state. After the completion of a RUNNING job, the first job from the WAIT queue is picked and is scheduled to run.

20.5 Using My Oracle Support for Additional Troubleshooting Information

You can use My Oracle Support (formerly MetaLink) to help resolve Oracle Fusion Applications problems. My Oracle Support contains several useful troubleshooting resources, such as:

- Knowledge Base articles

- Community Forums and Discussions

- Patches and Upgrades

- Certification information

You can access My Oracle Support at https://support.oracle.com.

21

Troubleshooting Oracle SOA Suite

This chapter describes common problems that you might encounter when using Oracle SOA Suite and explains how to solve them.

This chapter contains the following topics:

- Introduction to Troubleshooting Oracle SOA Suite
- Getting Started with Troubleshooting and Logging Basics for Oracle SOA Suite
- Runtime Diagnostics
- Security and Oracle WSM Policy Manager Configuration
- Human Workflow
- Patching and Deployment
- Performance
- Maintenance
- Custom Development (Extensibility)
- Using My Oracle Support for Additional Troubleshooting Information

To gain insight into the log details generated by Oracle Fusion applications, see Chapter 17. Also, review the *Oracle Fusion Middleware Error Messages Reference* for information about the error messages you may encounter.

> **Note:** Some sections of this chapter describe how to set properties in the System MBean Browser of Fusion Applications Control. Some MBean properties are applicable to Oracle WebLogic Server and others are applicable to the SOA Infrastructure. Oracle WebLogic Server Mbeans are properties that impact the JVM server level process that runs on the operating system. Example properties include the port on which the server is listening, and so on. These properties are agnostic to the type of application running on that server. SOA Infrastructure level Mbeans are properties that impact the SOA Infrastructure application running on the same Oracle WebLogic Server. Example properties include audit levels, transaction retries, and so on.

21.1 Introduction to Troubleshooting Oracle SOA Suite

This section provides guidelines and a process for using the information in this chapter. Using the following guidelines and process will focus and minimize the time you spend resolving problems.

Guidelines

When using the information in this chapter, Oracle recommends:

- After performing any of the solution procedures in this chapter, immediately retrying the failed task that led you to this troubleshooting information. If the task still fails when you retry it, perform a different solution procedure in this chapter and then try the failed task again. Repeat this process until you resolve the problem.

- Making notes about the solution procedures you perform, symptoms you see, and data you collect while troubleshooting. If you cannot resolve the problem using the information in this chapter and you must log a service request, the notes you make will expedite the process of solving the problem.

Process

Follow the process outlined in Table 21–1 when using the information in this chapter. If the information in a particular section does not resolve your problem, proceed to the next step in this process.

Table 21–1   Process for Using the Information in this Chapter

Step	Section to Use	Purpose
1	Section 21.2	Get started troubleshooting Oracle SOA Suite. The procedure in this section quickly addresses a wide variety of problems.
2	Section 21.3 through Section 21.9	Perform problem-specific troubleshooting procedures. These sections describe: ■ Possible causes of the problems ■ Solution procedures corresponding to each of the possible causes
3	Section 21.10	Use My Oracle Support to get additional troubleshooting information about Oracle Fusion Applications or Oracle SOA Suite. My Oracle Support provides access to several useful troubleshooting resources, including Knowledge Base articles and Community Forums and Discussions.
4	Section 21.10	Log a service request if the information in this chapter and My Oracle Support does not resolve your problem. You can log a service request using My Oracle Support at `https://support.oracle.com`.

21.2 Getting Started with Troubleshooting and Logging Basics for Oracle SOA Suite

SOA troubleshooting must be performed in the following scenarios:

1. The Oracle Fusion application transaction flow has an error and it has been identified as an error in SOA (or you want to check if it is a SOA-related error).

2. The Oracle Fusion application transaction flow has not completed and is taking longer than it should and it has been identified as stuck in SOA (or you want to check if it is stuck in SOA).

3. The Oracle Fusion application transaction flow is not working as expected, and the wrong data is passed.

The recommended approach to troubleshooting SOA issues is as follows:

1. Check the server and composite status log in Oracle WebLogic Server Administration Console and check if the SOA server is up and running successfully (that is, not in a failed state).

    a. Log in to Oracle WebLogic Server Administration Console.

    b. In the **Domain Structure**, click **Deployments**.

    c. In the **Name** column of the **Deployments** section, find **soa-infra**.

    d. Check that the **State** column is set to **Active** and the **Health** column is set to **OK**.

    e. Click **soa-infra** > **Monitoring** tab.

    f. Expand to ensure that all components have **Health** set to **OK**.

    g. Note the server name on which SOA is deployed.

    h. Go to the **Domain Structure** and click **Environment** > **Servers** to verify that the **State** column is set to **RUNNING** and the **Health** column is set to **OK**.

    i. Note the **Listen Port** column value, as this is the SOA runtime port.

2. Check if any Oracle Fusion application incident was created for the problem encountered.

For information about Oracle Fusion application incident processing, see Section 17.4.1.

3. Check whether the event was delivered to the SOA cluster by checking the log messages in Fusion Applications Control.

This requires the **oracle.integration.platform.blocks.event** logger to be set to **TRACE:32 (FINEST)**. There are potentially many types of exceptions that can appear in the log messages on the client side that raised the event. For example:

- If a Java class is not in the class path, then a `ClassNotFoundException` may appear.

- If the JNDI for the context or connection factory is not configured properly, then a `NamingException` or `FabricException` may appear.

- If the event payload XML that you create is invalid or contains invalid characters, you may receive a `FabricException` indicating that enqueueing of the event fails at the PL/SQL API with invalid characters, and so on.

For information about enabling this log and troubleshooting business events, see Section 21.3.2.

4. Check the composite instance flow. If the event was fired, then find your composite instance based on the application data.

See Section 21.3.1 for primary key mapping to the composite instance ID.

5. Review the composite instance flow and fault details. If the instance is in error and it is recoverable, attempt recovery using the **Recovery** tab of the BPEL process service engine in Fusion Applications Control.

6. View the logs.

If the problem still cannot be solved, increase the log level of the system to debug the transactions. See the list of loggers and log levels in the "Troubleshooting Oracle SOA Suite" appendix of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*. To simplify troubleshooting, it is recommended that you enable the following parent loggers at the **TRACE32 (FINEST)** level in Fusion Applications Control:

- **oracle.soa**

- **oracle.fabric**

- **oracle.integration**

The **oracle.wsm logger** can remain set to the **ERROR** level where it logs the required error messages. The **oracle.apps** logger should be set to **ALL**. To change logger levels, perform the following steps:

a. Go to Fusion Applications Control.

b. In the navigation pane, select **WebLogic Domain**.

c. Right-click a Managed Server from within the domain (each server's log levels can be independently set).

d. Choose **Logs** > **Log Configuration**.

e. In the **Logger Name** column, expand the **oracle** runtime loggers to display loggers such as **oracle.soa**.

f. Change the logging level to **TRACE:32**. The change should take effect within a few minutes.
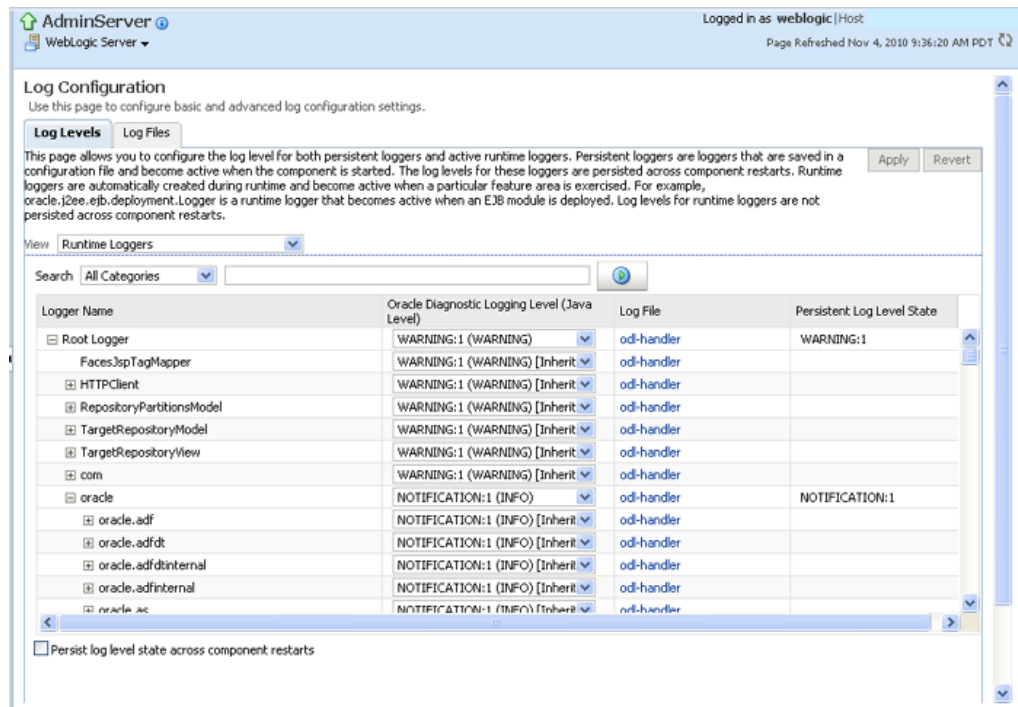
Note that in a production system, setting the trace at a fine-grained level can result in a large amount of output that must be diagnosed. You can alternately use selective tracing that provides a way to get a detailed, on-disk trace selectively (for example, by user name, thereby eliminating trace output for other users).

g. To activate selective tracing, right-click the domain under **WebLogic Domain** and choose **Logs** > **Selective Tracing**.

Note that **Selecting Tracing** does not display as an option when you right-click an Administration Server or Managed Server and choose **Logs**.

h. From the **Option Name** list, choose the type of selective trace (for example, based on user name), and start the trace.

i. When the problem has been reproduced, disable the trace and view the output to narrow down the issue.

For more information on selective tracing, see the "Configuring and Using Selective Tracing" section of *Oracle Fusion Middleware Administrator's Guide*.

j. Review the error logs (from Fusion Applications Control) for more information on the error.

Cross layer, server, and family functionality can be correlated through the execution context ID (ECID) (for example, you can look up the composite instance for a given expense report by correlating all the log entries with the ECID associated with that expense report transaction). For more information, see the "Correlating Messages Across Log Files and Components" section of *Oracle Fusion Middleware Administrator's Guide*.

7. Check for incorrect configurations and any networking issues, especially around potentially incorrect settings of the following:

- External load balancers
- Oracle HTTP Server
- Virtual host/IP address
- Oracle WebLogic Server front-end URL
- SSL host name verification settings

For more details, see the *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle SOA Suite*.

8. Verify that custom applications have been deployed on the correct servers using the correct configuration plan during deployment.

For more details, see the "Customizing SOA Composite Applications for the Target Environment" section of *Oracle Fusion Middleware Developer's Guide for Oracle SOA Suite*.

9. Contact Oracle Support Services.

If the error still cannot be resolved, file a ticket with Oracle Support Services and provide the logs and information shown in Table 21–2.

Table 21–2 SOA Log Information for Oracle Support Services

Log	Description
Application diagnostic log	View the following log: `$DOMAIN_HOME/servers/domain_name/logs/apps/server_name-diagnostic.log`
Server diagnostic log	View the following log: `$DOMAIN_HOME/servers/server_name/logs/server_name-diagnostic.log` For example, `soa_server1-diagnostic.log`, if `server_name` is `soa_server1`. This is where the log output is available. By default, only the last 100 MB of the diagnostic logs are retained.
Server log	View the following log: `server_name.log` For example, `soa_server1.log`, if `server_name` is `soa_server1`.
Server console output	`<stdout>` is also helpful, especially for deployment or patching issues.

Table 21–2   (Cont.) SOA Log Information for Oracle Support Services

Log	Description
Server thread dump	Enter the following at the operating system command prompt: `kill -3 managed_server_process_ID` You can also use Oracle WebLogic Server Administration Console. **1.** In the navigation tree of Oracle WebLogic Server Administration Console, select **Environment** > **Servers**. **2.** In the table, select the server. **3.** Select the **Monitoring** tab. **4.** Select the **Threads** tab. **5.** Click **Dump Thread Stacks**. The output is in the console logs.
OPatch thread dump	`kill -3 opatch_client_process_ID`
OWSM message log	The following log captures all SOAP messages on the wire. `$DOMAIN_HOME/servers/server_name/logs/owsm/msglogging/diagnostic.log` This log is not enabled by default. To enable this log: **1.** Go to **Fusion Applications Control** > **Weblogic Domain** > **Web Services** > **Policies**. **2.** Choose the security level for which to enable logging. **3.** Edit the policy to enable the log assertion.

For more information about setting logging levels for SOA components, see the "Setting Logging Levels for Troubleshooting" section in the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

10. Send Oracle Support Services information about the error packaged with Automatic Diagnostic Repository-based incident processing.

For information about how to investigate, report, and, in some cases, resolve a problem, see the "Investigating, Reporting, and Solving a Problem" section of *Oracle Fusion Middleware Administrator's Guide*.

21.3 Runtime Diagnostics

This section contains the following topics that describe common problems and solutions for Oracle SOA Suite runtime:

- Correlating Application Issues to the SOA Composite IDs

- Business Event Subscriptions Are Not Firing

- Events Are Consumed by Multiple Revisions of the Same Composites

- Rolled Back OAOO Event Delivery Messages are Displayed in the Log Files

- Application Transaction Does Not Complete and the Underlying Composite Is Stuck in a Running State

- BPEL Process Received No Response from an ADF Asynchronous Service

- Business Event Is Picked Up Twice (Or More) By SOA Server

- Long Running, Synchronous Calls To Remote Web Services Error Out or Asynchronous Transactions Return with an Error after a Long Time

- Some Messages Are Lost Between EDN and Composites or Composites Across Clusters

- Some Composites Are Being Retried Multiple Times on Failure

- Some Fusion Applications Control Features are Missing the No Recover Button or Export Composite Capability

- Automatic Recovery of BPEL Instances is Not Recovering A Specific Instance

- SOA Runtime Fails with a "Cannot read WSDL" Error

- Uploading a Composite for Oracle Support Services

- Confirming SOA Component Configuration Properties for Oracle Support Services

21.3.1 Correlating Application Issues to the SOA Composite IDs

Problem

How does an administrator translate an issue faced in the application to a composite instance in Fusion Applications Control for debugging purposes? For example, a purchase order transaction is stuck in processing for days and the administrator wants to see what the composite instance is doing. How does the administrator map the purchase order ID to the composite instance ID?

Solution

Oracle Fusion Applications modules write context-specific and business user-identifiable keys to the log files. This includes primary keys for any of the logical entities processed by the composite. Once the primary key and other context information is included in the log messages, the administrator can search for that context (for example, purchase order ID) in the log file and derive the composite instance ID and ECID. From there, they can diagnose the issue in Fusion Applications Control. Usually, the primary key information is written to the logs in case of any error/incident/message level.

Cross layer, server, and family functionality are all designed to be correlated through the ECID. For example, you can look up the composite instance for a given expense report by correlating all the log entries with the ECID associated with that expense report transaction.

During a specific time interval, you may create several hundred instances., This is a fairly common use case. If there is no error, or if the application has not implemented a BPEL process with sensor logging (should not be common), then instead of searching for the primary key, perform one of the following tasks:

- Try using the approximate timestamp of the transaction.

- Inspect the input payload of the audit trail flow (until you find the right one) to identify the right composite.

Any extensions/customizations are also expected to log the same context specific keys. You can set `AF_LOG` and `AF_LOG_MODULE` properly (such as setting log channel levels) to specify finer-grained logging from a particular family/logical business area (LBA)/composite.

21.3.2 Business Event Subscriptions Are Not Firing

Problem

When a business event is published, the business event subscription defined in the composite does not fire and a composite instance is not created.

This can occur, for example, with an ADF application invoking a SOA composite.

Solution

To resolve this problem.

1. Enable EDN logging in Fusion Applications Control. The following server loggers specific to EDN are available for selection:

   - **oracle.integration.platform.blocks.event**

   - **oracle.integration.platform.blocks.event.saq**

   - **oracle.integration.platform.blocks.event.jms**

   Note that **oracle.integration.platform.blocks.event.jms** only appears if EDN is running in EDN-JMS mode instead of the default EDN-DB mode.

   You can set the server loggers to one of the following levels:

   - **TRACE:1 (FINE)**

   - **TRACE:16 (FINER)**

   - **TRACE:32 (FINEST)**

   You can alternately use selective tracing to get detailed, on-disk trace selectively (for example, by user name, thereby eliminating trace output for other users). See Section 21.2 for details.

   Detailed logging goes into SOA server's `diagnostic.log` file configured in Fusion Applications Control. To set the log level for the loggers:

   a. Go to the navigation pane.

   b. Right-click **soa-infra**.

   c. Choose **Logs** > **Log Configuration**.

   d. Expand **oracle.integration** > **oracle.integration.platform** > **oracle.integration.platform.blocks** > **oracle.integration.platform.blocks.event**.

   e. Set the loggers described in step 1 to an appropriate logging level.

   The following sample shows a portion of a server log file:

```
[2011-03-22T11:52:37.038-07:00] [soa_server1] [TRACE] [SOA-31010]
[oracle.integration.platform.blocks.event.saq] [tid:
weblogic.work.j2ee.J2EEWorkManager$WorkWithListener@96bab0] [userId:
<anonymous>] [ecid:
5fc0ca821d51e919:3d214296:12ebaa2c996:-8000-00000000000c2b83,0:1:100000140]
[SRC_CLASS:
oracle.integration.platform.blocks.event.saq.SAQBusinessEventBusMessages]
[APP: soa-infra] [composite_name: MediatorPubSub] [component_name:
PublishEvent] [component_instance_id: 8E0411C054B511E0AF455DABE1395E7B] [SRC_
METHOD: fineDequeuedEvent] [composite_instance_id: 90087] Dequeued event,
Subject: null [source type "J"]:[[
    <business-event
xmlns:ns="http://schemas.oracle.com/events/edl/MyEventDefn"
```

```
xmlns="http://oracle.com/fabric/businessEvent">
      <name>ns:MyEvent</name>
      <id>a7ae9d28-9530-4049-8385-e0ebfb0eea50</id>

<source>default/MediatorPubSub!2.0*soa_
58454bff-3b36-4e29-a883-949085d85719/default/MediatorPubSub!2.0*soa_
58454bff-3b36-4e29-a883-949085d85719/PublishEvent</source>
    <tracking>
      <ecid>5fc0ca821d51e919:3d214296:12ebaa2c996:-8000-00000000000c2b83</ecid>
      <compositeInstanceId>90087</compositeInstanceId>
<parentComponentInstanceId>mediator:8E0411C054B511E0AF455DABE1395E7B</parentCom
ponentInstanceId>
    </tracking>
    <content>
        <inp1:singleString
xmlns:inp1="http://xmlns.oracle.com/singleString">PO123</inp1:singleString>
      </content>
</business-event>

  ]]
  [2011-03-22T11:52:37.042-07:00] [soa_server1] [TRACE] [SOA-31011]
[oracle.integration.platform.blocks.event.saq] [tid:
weblogic.work.j2ee.J2EEWorkManager$WorkWithListener@96bab0] [userId:
<anonymous>] [ecid:
5fc0ca821d51e919:3d214296:12ebaa2c996:-8000-00000000000c2b83,0:1:100000140]
[SRC_CLASS:
oracle.integration.platform.blocks.event.saq.SAQBusinessEventBusMessages]
[APP: soa-infra] [composite_name: MediatorPubSub] [component_name:
PublishEvent] [component_instance_id: 8E0411C054B511E0AF455DABE1395E7B] [SRC_
METHOD: fineFilterResults] [composite_instance_id: 90087] Filter [XPath Filter:
/be:business-event/be:content/ns1:singleString = 'PO123'] for subscriber
"default/MediatorPubSub!2.0*soa_
58454bff-3b36-4e29-a883-949085d85719/SubsEvent" returned true
. . .
. . .
```

For more information about setting logging levels for Oracle SOA Suite, see the "Setting Logging Levels for Troubleshooting" section of the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

2. Check whether the event publication is being delivered to the SOA cluster.

One way to do this is to use the log messages in Fusion Applications Control. The log messages include information about the incoming events and the event subject and payload. Before you can view the event information in the log messages, you must first set the **oracle.integration.platform.blocks.event** logger to **TRACE:32 (FINEST)**, as described in step 1, and then raise the event again. To search for the event-related messages:

a. Right-click **soa-infra**, and choose **Logs** > **View Log Messages**.

b. Specify the date range.

c. In the **Message Types** section, select the **Notification** and **Trace** checkboxes.

d. To search for all events delivered to the SOA cluster, specify **Message contains** `Dequeued event` as the search criteria, and click **Search**.

The subject associated with the event is displayed in the **Message** field and the business event namespace and local name, payload, ECID, and message details are displayed in the **Supplemental Detail** field.

  e. To search for a particular event by event name, click **Add Fields**.

  f. Select **Supplemental Detail**, and click **Add**.

  g. Specify that **Supplemental Detail contains** either the event namespace or event local name as the search criteria.

  h. Click **Search**.

  i. Use additional information such as the time interval or data in the payload to identify the specific event message of interest.

    The log message is generally a good way to determine whether the issue is with the publisher or subscriber. For example, if the subject is missing or the `fmw-context` is missing, it typically means there is an issue with the publisher. If the information reported in the log message is correct, then it typically means there is a problem with the subscriber.

    If you see the `Dequeued event` log message for your business event, then go to step 9.

    If you do not see the `Dequeued event` log message, there may be issues with the following:

For This Issue...	See...
EDN is in paused mode and temporarily stopped delivering events to its subscribers.	Step 3
Events are stuck in the advanced queue.	Step 4
More than one SOA cluster is pointing to the same advanced queue.	Step 5
There are internal SOA issues in the database layer.	Step 6
The data source is configured to point to the wrong `SOAINFRA` schema.	Step 7
There are issues with the Java, SOA, or PL/SQL code raising the event.	Step 8

3. When patches are applied to the SOA cluster, EDN is automatically placed in paused mode to prevent delivery of events during patching. Once patching is complete, EDN exits paused mode and resumes event processing. When a patch fails, EDN may remain in paused mode to prevent event subscriptions from firing after the failure. Follow these steps to determine if EDN is in paused mode and restart it:

  a. From the navigation pane, expand **SOA** and right-click **soa-infra**.

  b. Choose **Administration** > **System MBean Browser**.

  c. Go to **Application Defined MBeans** > **oracle.as.soainfra.config** > **Server:** *SOA_cluster_name* > **EDNConfig** > **edn**.

  d. In the **Attributes** column, locate the **Paused** property.

  e. If **Paused** is set to **true**, the listener threads are decreased to zero, which temporarily stops event delivery.

  f. To restart event delivery, select **false**.

  g. Click **Apply**.

4. At times, the message may be stuck in the advanced queue in different stages of delivery. Check the count of (potentially stuck) events currently in the following queue tables:

- `EDN_EVENT_QUEUE_TABLE`: Every event published is temporarily enqueued into this table (for `edn_event_queue`).

- `EDN_OAOO_DELIVERY_TABLE`: Only events with a one-and-only-one (OAOO) delivery target(s) are temporarily enqueued into this table (for `edn_ oaoo_queue`).

You can also check the count in the `edn-db-log`. The total number of messages in the `EDN_EVENT_QUEUE_TABLE` and `EDN_OAOO_DELIVERY_TABLE` are displayed at the top of the page.

a. Navigate to the following URL to see the total number of messages at the top of the page:

```
http://host_name:port/soa-infra/events/edn-db-log
```

You do not need to enable the `edn-db-log` to view these messages.

5. There may be more than one SOA cluster pointing to the same `SOAINFRA` schema. This is highly unlikely in a provisioned environment, but for environments that are created by developers, it may be an issue. In this case, the expected SOA runtime environment may not be receiving the events properly. It is expected that only one SOA runtime environment listens for business events on a `SOAINFRA` schema. Use the following query to identify the list of SOA runtime environments subscribing to business events from a `SOAINFRA` schema. Have only the required SOA runtime environment up, which points to the `SOAINFRA` schema, and shut down the others.

a. Navigate to the `edn-db-log` to see log messages that display each SOA cluster that is connecting to the EDN advanced queue:

```
http://host_name:port/soa-infra/events/edn-db-log
```

You do not need to enable the `edn-db-log` to view these messages. The messages continue to be displayed in the `edn-db-log` until the next time you click the **Clear** link. Here is a sample message that you see in the `edn-db-log`:

```
Starting EDN bus. Timestamp=Wed Feb 23 10:11:26 PST 2011.
    Parameters: platform="weblogic", cluster="", server="soa_server1",
    domain="fusion_domain", admin server="AdminServer", host="adc2180440".
```

b. Alternatively, you can run the following `select` statement against the `SOAINFRA` schema. For example, use `FIN_FUSION_SOAINFRA` for the SOA cluster in the Financials domain.

```
select * from V$SESSION where username like 'family_FUSION_SOAINFRA';
```

6. Verify that the event is properly being enqueued and dequeued by the underlying EDN PL/SQL procedures. You can do this using the `edn-db-log` that displays all events, including their namespace, names, payload, and subject, which are enqueued and dequeued by the PL/SQL EDN procedures.

a. Navigate to the following URL:

```
http://host_name:port/soa-infra/events/edn-db-log
```

b. Click the **Enable** link.

c. Raise the business event again.

d. Click the **Refresh** link.

> **Note:**
>
> - You must have the administrator privilege to enable/disable the `edn-db-log`.
>
> - You should always disable the `edn-db-log` after debugging to disable logging. This prevents excessive database growth in the EDN database log table. If the `edn-db-log` remains enabled, then debugging messages related to events that are published/enqueued into the database and subscribed to/dequeued from the database continue to be persisted into certain EDN database log tables. This causes the table to grow indefinitely.

7. The `EDNSource` and `EDNLocalTxSource` data sources may be pointing to a different database than the `EDNDataSource` and `EDNLocalTxDataSource` data source connections of the SOA server. Ensure that these four data source connections on ADF and the SOA server are pointing to the same database schema.

   Check or set the data sources in the Oracle WebLogic Server Administration Console:

   a. In the **Domain Structure**, select **Services** > **Data Sources**.

   b. In the **Name** column of the table, select *data_source_name*.

      The connection pool and driver name are only available for generic data sources, and not for multidata sources.

   c. Select the **Connection Pool** tab.

   d. In the **Driver Class Name** field, set the data source.

   For more information, see the "Configuring the SOA Infrastructure" chapter in the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

8. Check the log files in Fusion Applications Control for the code that raises the event to see if there were failures in raising the event or failures in the code that prevented the event from being raised.

   For ADF, check the log files. Look for the ECID corresponding to the business key (for example, purchase order).

   To enable event-related log messages in the server running the ADF application that raises the event, set the log level for the **oracle.integration.platform.blocks.event** logger to **TRACE (32)** on that server.

9. By the time you reach this step, you have confirmed that the business event is properly delivered to the SOA cluster. To determine whether any composite instances were created, perform the following steps:

   a. Because the ECID is different for a subscription with a guaranteed consistency level and a subscription with an OAOO level defined for the same event, you must first determine the consistency level for the event subscription. To do this, right click **soa-infra**, choose **Business Events**, and click the **Subscriptions** tab. Find the composite that subscribes to the business event in the **Components Subscriptions** table and get the value from the **Consistency Level** column.

b. If the value is set to **Guaranteed**, then get the ECID from the `Dequeued event` log message.

c. If the value is set to **One And Only One**, then search for the log message with **Message contains** `Dequeued OAOO event` as criteria and use the ECID from that message.

d. Locate the composite instance by right-clicking **soa-infra**, selecting the **Instances** tab, entering the ECID in the search criteria, and clicking **Search**.

If you see composite instances, then continue to step 12.

If you do not see any composite instances, this may be due to the following issues:

- There are issues with the subscription logic in the composite, possibly a mismatch in the event namespace or event name or error in the filter logic. See step 10.

- There may be a SOA issue. See step 11.

10. If a filter is being used, there may be an issue with the filter logic. Use Fusion Applications Control to review the event subscriptions and filters defined for the composite and the payload content.

a. Log in to Fusion Applications Control in the domain where the composite with the subscription is deployed.

b. Navigate to *domain_name* > **SOA**.

c. Right-click **soa-infra (*SOA_cluster_name*)**, and choose **Business Events**.

d. Click the **Subscriptions** tab.

e. Review the information in the **Event Name**, **Namespace**, and **XPath Filter** columns of the **Component Subscriptions** table.

Alternatively, you can also check the SOA server `diagnostic.log` file for filter-related logging. You must set the log level to at least **TRACE:1 (FINE)** for the **oracle.integration.platform.blocks.event.saq** logger. The following is an example log message that has the following statement pattern: `Filter [XPath Filter: ….] for subscriber "…" returned true/false`.

```
[2011-03-22T11:52:36.976-07:00] [soa_server1] [TRACE] [SOA-31011]
[oracle.integration.platform.blocks.event.saq] [tid:
[ACTIVE].ExecuteThread: '2' for queue: 'weblogic.kernel.Default
(self-tuning)'] [userId: <anonymous>] [ecid:
5fc0ca821d51e919:3d214296:12ebaa2c996:-8000-00000000000c2b83,0:1] [SRC_
CLASS:
oracle.integration.platform.blocks.event.saq.SAQBusinessEventBusMessages]
[WEBSERVICE_PORT.name: execute_pt] [APP: soa-infra] [composite_name:
MediatorPubSub] [component_name: PublishEvent] [component_instance_id:
8E0411C054B511E0AF455DABE1395E7B] [J2EE_MODULE.name: fabric] [SRC_METHOD:
fineFilterResults] [WEBSERVICE.name: PublishEvent_ep] [J2EE_APP.name:
soa-infra] [composite_instance_id: 90087] Filter [XPath Filter:
/be:business-event/be:content/ns1:singleString = 'PO123'] for subscriber
"default/MediatorPubSub!2.0*soa_
58454bff-3b36-4e29-a883-949085d85719/SubsEvent" returned true
```

11. Check the SOA diagnostic log for stack traces related to the issue. For example, the SOA instance may have been rolled back due to an error.

12. By the time you reach this step, you have confirmed that the business event has been delivered to the SOA cluster and a composite instance has been created. However, you may be encountering issues with the composite components:

- If you see a composite instance, but not a BPEL instance, check if the BPEL process can be executed independently by a client, as opposed to the event subscription.

- If you see a composite instance, but not a BPEL instance, this symptom may indicate a stuck thread on a service call. In this scenario, the BPEL process service engine is hung on a service call and has not created the audit trail. To verify the behavior, take three thread dumps approximately 30 seconds apart on all servers in the cluster. See Table 21–2 for more information on taking thread dumps. If the same thread shows it is stuck after the three thread dumps (that is, after 1.5 minutes), the thread is stuck. In this case, wait for the call to time out, which then enables you to recover the instance.

- If the audit trail shows the following message:

  ```
  Invoked 1-way operation "initiate" on target service service_name
  ```

  There may be an issue with an incorrect driver type for the data sources. Note that the driver XA configuration should already be created by the SOA/provisioning template. If the driver type is changed or is not set properly, routing from Oracle Mediator does not occur. Instead, it tries to route and fails several times. This can be confirmed by instances in Fusion Applications Control. As per Oracle Fusion Applications standards, the `SOADataSource` and `EDNDataSource` data sources should use the XA driver type; that is:

  - `EDNDataSource`: `oracle.jdbc.xa.client.OracleXADataSource`

  - `SOADataSource`: `oracle.jdbc.xa.client.OracleXADataSource`

21.3.3 Events Are Consumed by Multiple Revisions of the Same Composites

Problem

Events are consumed by multiple revisions of the same composites.

Solution

By design, different composites can subscribe to the same event. However, it is an error if multiple revisions of the same composite subscribe to the same event. This occurs when you have multiple active versions of the composite that may be a result of a patch failure.

Perform the following steps to determine whether there are multiple, active revisions of the same composite and to retire the composite revision that should not be active:

1. Go to Fusion Applications Control.

2. Click **SOA** > **soa-infra**.

   In the **Deployed Composites** section of the **Dashboard** tab, you see a list of deployed composite names, revisions, and modes (for example, active).

3. Identify composites with the same name and with an active mode, but with different revisions.

4. Click the composite revision that should not be active.

5. Click the **Retire** button.

21.3.4 Rolled Back OAOO Event Delivery Messages are Displayed in the Log Files

Problem

If the Java debug level is set to **TRACE:16 (FINER)** or a lower value, you may see log messages such as the following:

```
Began XA for OAOO
Rolled back XA for OAOO
```

Solution

These are normal messages of OAOO event delivery when there are no events waiting to be delivered. These are *not* error conditions. You can turn off these messages by setting the Java logging level to **TRACE:1 (FINE)** or a higher value.

21.3.5 Application Transaction Does Not Complete and the Underlying Composite Is Stuck in a Running State

Problem

An application transaction is not completing. For example, a purchase order status may remain processing. Checking the composite instance shows that the composite is stuck in the running state. In this case, the component is probably not running. Instead, it has likely faulted and may need recovery.

Faults may occur for various reasons:

- A BPEL activity faulted with an error (for example, a business error, security authorization error, or some other error).

- A BPEL activity invoked an external web service that was unavailable.

- A BPEL activity has already been terminated by the administrator using Fusion Applications Control.

- A BPEL activity invoked an asynchronous ADF service and the message is stuck in the AQ/JMS queue.

- A BPEL activity invoked an asynchronous ADF service, but because SOA was unavailable, the callback message did not arrive.

- A BPEL activity invoked a synchronous ADF service, which is taking a long time (or is hanging).

- A network error occurred.

Solution

1. Log in to Fusion Applications Control.

2. In the navigator pane, go to *domain_name* > **SOA**.

3. Click **soa-infra (*SOA_cluster_name*)**.

4. Click the **Instances** tab.

5. Search for the composite instance, and click the instance ID.

   The Flow Trace page appears.

   If the instance is not visible (and the **Audit Level** is not set to **Off** in the SOA Infrastructure Common Properties page), this implies that the message is stuck

outside of SOA. Check the ADF log (if invoked through an ADF service) to find out if it is stuck in the AQ/JMS queue.

If the message has reached Oracle Mediator, but not instantiated the BPEL flow, the BPEL instance may have been rolled back from the start due to an error.

If the BPEL flow exists, the **Faults** section of the Flow Trace page typically shows the faulted service that can trace the root cause.

6. In the **Trace** section, click the BPEL process.

7. Expand the BPEL audit trail to see the exact point at which the service faulted.

   This information is also available in the **Faults** tab of the BPEL flow trace. It also indicates whether the fault can be recovered.

8. Click the **View Raw XML** link.

   The same information is also available through this link, where you can see the error. For example:

```
...
<message>Faulted while invoking operation "modifyUserRoles" on provider
 "UserService". </message><details>
...
<tns:message>JBO-27023: Failed to validate all rows in a
 transaction.</tns:message>
<tns:severity>SEVERITY_ERROR</tns:severity>
...
<message>The transaction was rolled back. The work performed for bpel instance
 "451042" was rolled back to the previous dehydration point, but the audit
 trail has been saved.
You can recover the instance from the recovery console by resubmitting the
 callback message or activity for execution.</message
```

   Since the instance was rolled back to its previous dehydration point, the status remains as **Running**.

9. In the **Audit Trail** and **Faults** tabs, make a note of the following:

   ■ Composite name (for example, **HcmUsersSpmlComposite**)

   ■ Component (for example, **UpdateGuid** BPEL process)

   ■ BPEL instance ID (for example, **bpel:451042**)

   This is all used in the recovery of the instance, if it is recoverable. Note that the Audit Trail window may mark the error as a nonrecoverable business fault, but the recoverability of the message can be found in the **Recovery** tab of the BPEL process service engine.

10. To attempt to recover the instance, right-click **soa-infra (***SOA_cluster_name***)**, and choose **Service Engines** > **BPEL**.

11. Click the **Recovery** tab.

12. From the **Type** list, select **Activity**.

13. Specify the composite and component names captured in step 9, and click **Search**.

14. Find the specific BPEL instance ID. You can recover faults marked as **Recoverable**.

15. Check the other recovery options in the **Type** list (for example, **Invoke** and **Callback**), if they exist.

> **Note:** You can also search for recoverable messages from the **Faults and Rejected Messages** tab of the SOA Infrastructure by clicking the message and selecting the appropriate action from the **Recovery Actions** list.

16. If the instance is not marked as recoverable, then reinvoking the service is not allowed (most probably because it is not idempotent). In some cases, you may need to provide diagnostic information to Oracle Support Services to resolve issues with nonrecoverable, nonidempotent transactions.

17. If the BPEL activity has invoked an asynchronous ADF service and the message is stuck in the AQ JMS queue, you can view the `server.log` and `server-diagnostic.log` files to see the logging of the message metadata logged by the JRF web services infrastructure.

   In addition, the ADF diagnostic logs are also available to debug, if needed.

   a. Use the ECID field to correlate and track ADF service logging corresponding with the SOA composite that invoked it.

   b. When viewing the log in Fusion Applications Control, click the **Broaden Target Scope** dropdown list and select the *farm_name/domain_name* (Oracle WebLogic Domain) to view messages across the domain.

   c. In the Log Messages page for the Oracle WebLogic Server domain, in the **Selected Targets** section, ensure that the search includes the **ECID** field with the value noted in step a and the **Component Name** field is set to **adf-bc**.

   d. Search and view log records for ADF-BC and the ECID and note any issues.

   For a specific ECID, you find several root instances (top level clients). You must drill down to the appropriate instance to find a specific fault.

   e. Observe if the Oracle ADF Business Components completed successfully or completed with an error. See the "Viewing and Searching Log Files" section in the *Oracle Fusion Middleware Administrator's Guide*.

   For more details on diagnosing ADF-BC asynchronous web service calls, see Section 21.3.6.

See the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite* for the various ways to recover transactions and messages:

- "Managing SOA Composite Applications" chapter to recover from the SOA composite application page in Fusion Applications Control.

- "Managing BPEL Process Service Components and Engines" chapter to recover from the BPEL process service component and BPEL process service engine message recovery pages.

21.3.6 BPEL Process Received No Response from an ADF Asynchronous Service

Problem

A BPEL activity has invoked an asynchronous ADF service and not received a response. The message may be stuck in the request or response queues.

Solution

1. Check if the reply address is valid, the server is running, and the name of the server and port match with the server on which the BPEL process is running.

2. Verify that the policy on the callback receive activity matches with the policy advertised in the WSDL for the response port type/binding.

3. Look at the *server-name*_diagnostic.log file and check for all status messages of the request.

   a. Go to Fusion Applications Control.

   b. Right-click **soa-infra**.

   c. Choose **Logs** > **View Log Messages**.

   d. Click **Target Log Files**.

   e. Select the server diagnostic file to view, and click **View Log File**.

   For more information about the different types of loggers to set, see the "Setting Logging Levels for Troubleshooting" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

   Note that every log contains the message ID and other details, such as ECID, that can help isolate it from other messages. Once you find the message ID (that is, MessageID: urn:uuid:ac1a4a81-39df-45b1-a741-e16e752d5d33), search for the same in the log file. You should see messages such as the following for each successful asynchronous operation.

   If the following log is not available, this means the asynchronous operation was never called on this server:

   ```
   [2010-12-17T12:27:13.537-08:00] [AdminServer] [NOTIFICATION] []
    [oracle.j2ee.ws.common.jaxws.JAXWSMessages]
   [tid: [ACTIVE].ExecuteThread: '2' for queue: 'weblogic.kernel.Default
    (self-tuning)'] [userId: <anonymous>]
   [ecid: 0000Inq0wfa76EWLHyo2yf1D2wUW000000,0:1]
   [WEBSERVICE_PORT.name: AsyncEjbPort] [APP: AsyncEjb] [J2EE_MODULE.name:
    AsyncEjb-ejb] [WEBSERVICE.name: AsyncEjbService] [J2EE_APP.name: AsyncEjb]
   *[MessageID: urn:uuid:ac1a4a81-39df-45b1-a741-e16e752d5d33] Sending message
    to JMS queue*
   "oracle.j2ee.ws.server.async.NonDefaultRequestQueue" for async processing
    of service "1d7551aa-73d9-4624-ad52-c15e35b5b25dRequest"
   ```

   If the following log is not available, this means the asynchronous operation was never called on this server:

   ```
   [2010-12-17T12:27:13.783-08:00] [AdminServer] [NOTIFICATION] []
    [oracle.j2ee.ws.common.jaxws.JAXWSMessages]
   [tid: [ACTIVE].ExecuteThread: '2' for queue: 'weblogic.kernel.Default
    (self-tuning)'] [userId: <anonymous>]
   [ecid: 0000Inq0wfa76EWLHyo2yf1D2wUW000000,0:1]
   [WEBSERVICE_PORT.name: AsyncEjbPort] [APP: AsyncEjb] [J2EE_MODULE.name:
    AsyncEjb-ejb] [WEBSERVICE.name: AsyncEjbService] [J2EE_APP.name: AsyncEjb]
   *[MessageID: urn:uuid:ac1a4a81-39df-45b1-a741-e16e752d5d33] An asynchronous
    request message is received*
   and successfully recorded for service "AsyncEjbService" with a replyTo
   address as
   "http://adc2180314:7001/AsyncEjbCallback/AsyncEjbResponseImplService"
   ```

```
[2010-12-17T12:27:13.986-08:00] [AdminServer] [NOTIFICATION] []
 [oracle.j2ee.ws.common.jaxws.JAXWSMessages]
[tid: [ACTIVE].ExecuteThread: '0' for queue: 'weblogic.kernel.Default
 (self-tuning)'] [userId: weblogic]
[ecid: 0000Inq0wfa76EWLHyo2yf1D2wUW000000,0] [APP: AsyncEjb]
*[MessageID: urn:uuid:ac1a4a81-39df-45b1-a741-e16e752d5d33]
 Startedasynchronous request processing for the service*
"AsyncEjbService" with the message selector
 "1d7551aa-73d9-4624-ad52-c15e35b5b25dRequest". Transaction enabled: "true"

[2010-12-17T12:27:14.296-08:00] [AdminServer] [NOTIFICATION] []
 [oracle.j2ee.ws.common.jaxws.JAXWSMessages]
[tid: [ACTIVE].ExecuteThread: '0' for queue: 'weblogic.kernel.Default
 (self-tuning)'] [userId: weblogic]
[ecid: 0000Inq0wfa76EWLHyo2yf1D2wUW000000,0] [APP: AsyncEjb]
*[MessageID: urn:uuid:ac1a4a81-39df-45b1-a741-e16e752d5d33] Completed
 asynchronous request processing. A response will be sent to the client.*

[2010-12-17T12:27:14.312-08:00] [AdminServer] [NOTIFICATION] []
 [oracle.j2ee.ws.common.jaxws.JAXWSMessages]
[tid: [ACTIVE].ExecuteThread: '2' for queue: 'weblogic.kernel.Default
 (self-tuning)'] [userId: weblogic]
[ecid: 0000Inq0wfa76EWLHyo2yf1D2wUW000000,0] [APP: AsyncEjb]
*[MessageID: urn:uuid:ac1a4a81-39df-45b1-a741-e16e752d5d33]
 Startedasynchronous response processing for the service*
"AsyncEjbService" with the message selector
"1d7551aa-73d9-4624-ad52-c15e35b5b25dResponse".

[2010-12-17T12:27:15.018-08:00] [AdminServer] [NOTIFICATION] []
[oracle.j2ee.ws.common.jaxws.JAXWSMessages]
[tid: [ACTIVE].ExecuteThread: '2' for queue: 'weblogic.kernel.Default
(self-tuning)'] [userId: weblogic]
[ecid: 77e271e3-afd5-48e8-a83d-9b326860f3f2-0000000000000018,0] [APP:
AsyncEjb] *[MessageID: urn:uuid:ac1a4a81-39df-45b1-a741-e16e752d5d33]
Completed asynchronous response processing successfully. Client must have
received the response by now.*
```

To monitor the JMS queues used for asynchronous web services in Oracle WebLogic Server Administration Console:

1. In the **Domain Structure**, select **Services**.

2. Select **Messaging** > **JMS Modules**.

3. Select *JMS_module* > **Queue Name**.

   A JMS module is defined for each ADF J2EE application cluster.

   The queue names are set up as follows:

Family AQ	Name
Oracle Fusion Customer Relationship Management Product	CRM_AsyncWS_Request and CRM_AsyncWS_Response
Oracle Fusion Human Capital Management Product	HCM_AsyncWS_Request and HCM_AsyncWS_Response
Oracle Fusion Financials Product	FIN_AsyncWS_Request and Fin_AsyncWS_Response
Oracle Fusion Procurement Product	PRC_AsyncWS_Request and PRC_AsyncWS_Response
Oracle Fusion Project Product	PRJ_AsyncWS_Request and PRJ_AsyncWS_Response

Family AQ	Name
Oracle Fusion Supply Chain Management Product	`SCM_AsyncWS_Request` and `SCM_AsyncWS_Response`
Oracle Fusion Setup Product	`COMMON_AsyncWS_Request` and `COMMON_AsyncWS_Response`
IC	`IC_AsyncWS_Request` and `IC_AsyncWS_Response`

> **Note:** Asynchronous web services support multiple JMS providers. When the Oracle WebLogic Server-JMS provider is used, the error queues are used:
>
> - `JRFWSAsyncRequestErrorQueue_ContractManagementCluster`
>
> - `JRFWSAsyncResponseErrorQueue_ContractManagementCluster`
>
> When the AQ-JMS provider is used (as is the case with Oracle Fusion Applications), the error queues are not used or required.

4. Select the appropriate queue in **Summary of Resources** > **Monitoring** tab.

    This shows the high level statistics of the messages. For example:

    - **Messages Total** shows messages that have been processed.

    - **Messages Current** shows messages currently in the queue (ideally zero).

5. If there are messages in the queue, drill down further by clicking the checkbox in the first column to select the queue.

6. Click **Show Messages**.

    You need administrator access and access to the JMS queue that can be configured through the JMS queue **Security** > **Policies** tab.

7. From the list of messages, select the message to view its details.

    By looking at the properties of the message, you can obtain useful information for debugging. For example, you can verify if the `wsaReplyTo` points to a valid URL at which the callback should be running. The service name can also narrow down the problem domain.

21.3.7 Business Event Is Picked Up Twice (Or More) By SOA Server

Problem

Business events are raised from Fusion J2EE applications and are picked up more than once. Expected behavior is that they should be picked up only once.

Solution

If a patch fails, it may be possible for two versions of a given composite to be active (that is, the older version has not been retired). When multiple versions of a composite are active, they all become subscribers and the event is picked up more than once. This should not happen under normal scenarios. Reapply the patch and ensure that the deployed composite only has one active version.

21.3.8 Long Running, Synchronous Calls To Remote Web Services Error Out or Asynchronous Transactions Return with an Error after a Long Time

Problem

Long running synchronous calls to remote web services end with JTA Transaction rolled-back errors. When executing a transaction making an asynchronous call (for example, to the SOA server), the application returns with an error. Server logs show JTA transaction timeouts, which can cause this behavior.

Solution

Check the JTA transaction timeout in Oracle WebLogic Administration Console.

1. Log in to Oracle WebLogic Server Administration Console.

2. In the **Domain Structure**, select **Services** > **JTA** to check the timeout value.

    If the transaction is always timing out beyond 30 seconds and is a custom composite synchronous client invocation, then you may need to revisit the design approach. It may be best for the external web service to be invoked as an asynchronous transaction.

    Increasing the JTA for supporting long running synchronous transactions is simply an interim mechanism.

    For information on changing the JTA transaction timeout setting, see the "Resolving Connection Timeouts" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

3. If synchronous client invocations to Oracle Fusion Applications code take a long time, check for any performance issues with the system and try to resolve them.

4. Check for appropriate values for the **syncMaxWaitTime** property and BPEL's EJB transaction timeout settings in relation to the JTA timeout settings and only then increase the value of the JTA timeout, if needed.

    For information on viewing and changing the **syncMaxWaitTime** property, see the "Configuring BPEL Process Service Engine Properties" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

5. To view and change the BPEL EJB transaction timeout settings, perform the following steps:

    a. Log in to Oracle WebLogic Server Administration Console.

    b. In the **Domain Structure**, click **Deployments**.

    c. Expand **soa-infra** > **EJBs**.

    d. Update the following EJBs:

    - **BPELActivityManagerBean**

    - **BPELDeliveryBean**

    - **BPELDispatcherBean**

    - **BPELEngineBean**

    - **BPELFinderBean**

    - **BPELInstanceManagerBean**

- **BPELProcessManagerBean**

- **BPELSensorValuesBean**

- **BPELServerManagerBean**

    e. Click **Save**.

    f. Restart Oracle WebLogic Server.

6. For asynchronous transactions, check the values for both the BPEL EJB transaction timeout and the JTA transaction timeout and adjust as needed.

21.3.9 Some Messages Are Lost Between EDN and Composites or Composites Across Clusters

Problem

Messages primarily get lost for the following reasons:

1. The EDN message was not delivered.

2. The target asynchronous service did not respond.

3. The message was roll backed (though it was not lost; you still see it in Fusion Applications Control).

Solution

This issue typically requires you to identify where the messages are supposed to be and to diagnose the path.

- If the calling component is an ADF web service invoking a SOA composite, then check for any client errors in the ADF logs.

- Because ADF invokes SOA composites using EDN, check the JMS/AQ and XA data source configurations to ensure that events are being fired. For more information about EDN issues, see Section 21.3.2.

- Check if the Oracle WSM Policy Manager security configurations allow the client to invoke the SOA composite.

- Verify that the transaction is not transient (nonpersistent).

    Oracle BPEL Process Manager uses the dehydration store database to maintain long-running, asynchronous processes and their current state information in a database while they wait for asynchronous callbacks. Storing the process in a database preserves the process and prevents any loss of state or reliability if a system shuts down or a network problem occurs. There are two types of processes in Oracle BPEL Process Manager. These processes impact the dehydration store database in different ways.

    – Transient processes: This process type does not incur any intermediate dehydration points during process execution. If there are unhandled faults or there is system downtime during process execution, the instances of a transient process do not leave a trace in the system. Instances of transient processes cannot be saved in-flight (whether they complete normally or abnormally). Transient processes are typically short-lived, request-response style processes. The synchronous process you design in JDeveloper is an example of a transient process.

    – Durable processes: This process type incurs one or more dehydration points in the database during execution because of the following activities:

> * Receive activity
>
> * OnMessage branch of a pick activity
>
> * OnAlarm branch of a pick activity
>
> * Wait activity
>
> Instances of durable processes can be saved in-flight (whether they complete normally or abnormally). These processes are typically long-living and initiated through a one-way invocation. Because of out-of-memory and system downtime issues, durable processes cannot be memory-optimized.

- If a composite instance is not visible in Fusion Applications Control and the SOA Infrastructure is running, check that the **Audit Level** is not set to **Off** in the SOA Infrastructure Common Properties page. This can be checked in Fusion Applications Control:

  1. Right-click **soa-infra (***SOA_cluster_name***)**.

  2. Choose **SOA Administration** > **Common Properties**.

- Check the ADF logs to see if there was a bad/malformed message or incorrect SOAP headers (if invoking from an external web service). The transaction may have rolled back before being dehydrated due to an internal error. Again, the SOA logs can identify the cause of this issue.

- If the composite instance is available, check the **bpel.config.oneWayDeliveryPolicy** BPEL property value. You can check the value in Fusion Applications Control:

  1. In the navigation pane, expand **soa-infra (***SOA_cluster_name***)**.

  2. Expand the partition, and select the composite.

     The Dashboard page for the composite is displayed.

  3. In the upper right corner, click the **Show XML Definition** icon.

     The contents of `composite.xml` for that composite are displayed.

  If this is set to **async.cache**, you may lose messages. Set it to **async.persist** for reliable messages. This is typically specified in the BPEL process service component section of the `composite.xml` file, so this can be set for custom composites. If the value is not set in `composite.xml`, the value for **oneWayDeliveryPolicy** in the System MBean Browser in Fusion Applications Control is used. The following values are possible:

  - **async.persist**: Messages are persisted in the database hash map.

  - **async.cache**: Messages are stored in memory.

  - **sync**: Direct invocation occurs on the same thread.

  For more information about these settings, see the "Deployment Descriptor Properties" appendix and the "Transaction and Fault Propagation Semantics in BPEL Processes" chapter of the *Oracle Fusion Middleware Developer's Guide for Oracle SOA Suite*.

  For information about setting the **oneWayDeliveryPolicy** property in the System MBean Browser, see the "Configuring BPEL Process Service Engine Properties" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

It is also possible that the transaction has invoked a target asynchronous service that has not responded back. In this case, the composite instance flow shows the call to the target asynchronous service.

■ Check if the instance has rolled back and the message is in recovery.

1. Log in to Fusion Applications Control.

2. Right-click **soa-infra (*SOA_cluster_name*)**, and choose **Service Engines** > **BPEL**.

3. Click the **Recovery** tab.

This may occur if any external references receiving the message are not reachable (for example, an external web service, enterprise applications such as Siebel, and so on).

21.3.10 Some Composites Are Being Retried Multiple Times on Failure

Problem

When a BPEL process flow errors out, it is retried with all its invocations. This is undesirable in some cases.

Solution

The property **GlobalTxMaxRetry** (default value is 3) specifies how many retries are performed if an error is identified as a retriable one. For example, after several web service invocations, if dehydration fails due to a data source error, then this is identified as a retriable error and all activities from the prior dehydration state are retried. If the activities being retried are not idempotent (that is, their state can change with each retry and is not guaranteed to give the same behavior), then multiple retries can be problematic.

To rectify this, customize the composite by specifically marking the nonidempotent activities with `idempotentset` set to `false` in the partner link settings section of the `composite.xml` file to prevent retries.

```
<property name="bpel.partnerLink.partner_link_name.idempotent">false</property>
```

You cannot set the `idempotentset` property in Fusion Applications Control.

For more information about Oracle SOA Suite customizations in JDeveloper, see the "Customizing and Extending SOA Components" chapter of the *Oracle Fusion Applications Extensibility Guide*.

You can also set **GlobalTxMaxRetry** to `0` in the Systems MBean Browser.

1. Right-click **soa-infra (*SOA_cluster_name*)**.

2. Choose **SOA Administration** > **Common Properties**.

3. Click **More SOA Infra Advanced Configuration Properties**.

4. Click **GlobalTxMaxRetry**.

5. In the **Value** field, enter an appropriate value.

6. Click **Apply**.

For information about the `idempotent` property, see the "BPEL Properties Set Inside a Composite" section of *Oracle Fusion Middleware Performance and Tuning Guide*.

21.3.11 Some Fusion Applications Control Features are Missing the No Recover Button or Export Composite Capability

Problem

Fusion Applications Control is missing features needed for diagnostics and recovery. For example, the **Recovery** tab for the BPEL process service engine shows a BPEL instance as recoverable along with instructions to click the **Recover** button. However, there is no **Recover** button available. In addition, the ability to export a composite is not available.

Solution

Check that you are logged in as a user with administrator privileges and are not logged in as a read-only user. Read-only users do not see administrative features such as the **Recover** button or the ability to export composites.

See the "Oracle Enterprise Manager Roles" appendix of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

21.3.12 Automatic Recovery of BPEL Instances is Not Recovering A Specific Instance

Problem

BPEL processes have an automatic recovery feature that attempts to automatically recover activities that are recoverable such as unresolved invoke/callback messages, activities not completed over a provided threshold time, and so on. However, the automatic recovery feature only tries to recover a few instances and only retries a fixed number of times. If some instances are not being automatically recovered, they are likely not being picked up because of the configuration of the automatic recovery parameters.

Solution

- Set the maximum number of messages to automatically recover.

  By default, the automatic recovery feature of Oracle BPEL Process Manager processes fifty messages to submit for each recovery attempt. This is controlled by the **maxMessageRaiseSize** property.

  1. In the navigation pane, right-click **soa-infra (***SOA_cluster_name)*.

  2. Choose **SOA Administration** > **BPEL Properties** > **More BPEL Configuration Properties** > **RecoveryConfig**.

  3. Expand both **RecurringScheduleConfig** > **maxMessageRaiseSize** and **StartupScheduleConfig** > **maxMessageRaiseSize**.

     The default value is 50 for each. A negative value causes all messages selected from the database to be submitted for recovery. A value of 0 causes no messages to be selected from the database (effectively disabling recovery). To recover more than fifty messages, set the property value to that value. Use this property to limit the impact of recovery on the server.

- Set the maximum number of automatic recovery attempts on a given message.

  You can also configure the number of automatic recovery attempts to submit in the same recoverable instance. The value you provide specifies the maximum number of times that invoke and callback messages are recovered. If the value is 0 (the default value), it recovers all messages. Once the number of recovery attempts on a message exceeds the specified value, a message is marked as nonrecoverable.

Follow these steps to configure automatic recovery attempts for invoke and callback messages in Fusion Applications Control.

1. In the navigation pane, right-click **soa-infra (***SOA_cluster_name***)**.

2. Choose **SOA Administration** > **BPEL Properties** > **More BPEL Configuration Properties**.

3. Select **MaxRecoverAttempt**, and enter a value in the **Value** field.

4. Click **Apply**.

   It may not be desirable in all cases to use automatic recovery. If services are not idempotent, then corruption can occur. Moreover, the automatic recovery restores the composite to the last save point that can be immediately after an asynchronous invoke, wait, and so on. Therefore, it is important to understand the process behavior and what it does next before performing mass recoveries. Attempt mass automatic recovery only after the root cause of the composite failures has been fixed (for example, a service that was unavailable is now available, a database running out of space was fixed, and so on). Automatic recovery can also trigger an unexpected load during failure scenarios. This causes more threads to block on a remote server that can induce hangs in the SOA server in a cascading fashion.

   For more information, see Section 21.3.10 and the "Configuring Automatic Recovery for Oracle BPEL Process Manager" section in the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

21.3.13 SOA Runtime Fails with a "Cannot read WSDL" Error

Problem

The following error is displayed if either the endpoint is not available for a reference or a composite is deployed with an incorrect deployment plan file.

```
oracle.fabric.common.FabricException: Cannot read WSDL
```

Solution

1. Ensure that the endpoint for the reference is up and running.

   The WSDL or endpoint is stated in the error. The service can then be looked up from Fusion Applications Control to check if it is active. Once the service is active and the endpoint is reachable, search for the SOA instance and retry through the **Recovery** tab for the BPEL process service engine in Fusion Applications Control.

2. For extensions/customizations, ensure that the correct URL is updated in the deployment configuration plan.

   For information about deployment plans, including examples of using `sca_extractPlan` to extract plans, see the "Customizing Your Application for the Target Environment Prior to Deployment" section of *Oracle Fusion Middleware Developer's Guide for Oracle SOA Suite* and the "Customizing SOA Composites with JDeveloper" section of *Oracle Fusion Applications Extensibility Guide*.

21.3.14 Uploading a Composite for Oracle Support Services

Problem

Oracle Support Services has requested you to upload your composite to review the flow and associated configuration files.

Solution

1. Log in to Fusion Applications Control as a user with administrator privileges, and not as a read-only user.

2. In the navigation pane, expand **soa-infra (***SOA_cluster_name)***.

3. Expand the partition in which the composite is located (for example, **default**).

4. Right-click the composite and choose **Export**.

5. Accept the default selections, and save the resulting JAR file.

6. Provide the JAR file to Oracle Support Services so that it can be imported for review in JDeveloper by selecting **File** > **Import**.

21.3.15 Confirming SOA Component Configuration Properties for Oracle Support Services

Problem

Oracle Support Services has requested you to confirm configuration properties of various SOA components.

Solution

1. In the navigation pane, right-click **soa-infra (***SOA_cluster_name)***.

2. Choose **Administration** > **MDS Configuration** > **Export**.

3. Upload the resulting file to Oracle Support Services for review.

21.4 Security and Oracle WSM Policy Manager Configuration

For information about troubleshooting Oracle WSM Policy Manager, see the "Diagnosing Problems" chapter in the *Oracle Fusion Middleware Security and Administrator's Guide for Web Services*.

21.5 Human Workflow

This section contains the following topics that describe common problems and solutions for human workflow:

- Task Detail Page is Not Available

- Task Chooses an Incorrect User If a Number of Rules Are Defined or Errors with "Ruleset returned lists with different list builder"

- Document Is Stuck in Pending Approval and Logs Show Error "Unable to walk hierarchy for user user_name"

- Task History: Cannot Remove or Move a Participant After Yourself

- Hierarchy List Builders Display an Error Due to Security Configuration Issues and Logs Report Errors in Authentication or Authorization with HCM Services

- HCM Services are Timing Out or Getting Stuck
- Error/Warning Reported in Retrieving the Position for a User
- Users See No Approvals in the Inbox when Using Position Hierarchy
- Task Escalation Does Not Follow Job Level, Supervisory, or Position Hierarchies
- User Is Not Allowed to Perform an Action on a Human Workflow Task
- Email Notification Is Not Being Sent Out
- Notifications Are Sent But Are Not Actionable
- Actionable Notifications Are Sent But No Action is Taken After Responding
- After LDAP Configuration with Oracle WebLogic Server, Users Appear in the Console, But a "No Role found matching the criteria" Error Appears
- Only a Subset of Users in LDAP Can Log In to Oracle BPM Worklist
- "Unknown Macro" Exception Appears when Myrealm Is Passed as a Parameter to the Identity Context To Identity Service APIs
- Add Participant Button in Oracle BPM Worklist Is Disabled
- Task Modifications Made in a Runtime Tool Are Not Appearing for the Task
- How Can I Set Commonly Used Human Workflow Configuration Parameters?
- How Do I Set Human Workflow Configuration Parameters Not Available in the Fusion Applications Control Properties Pages?
- System MBean Browser Does Not Reflect My Changes After Editing the Human Workflow Config MBeans
- Human Workflow Services Are Not Locating Resource Bundles or Classes Located at the Workflow Customizations Class Path URL
- How Do I Manually Set the URL for Displaying Task Details for a Particular Human Workflow Task Component?
- How to Test the Health of the Installed Server
- User Authentication Error (in this example, the user is FMW_USERID and identityContext is jazn.com)
- Error During Import of Task Payload Mapped Attribute Mappings
- Error During Rule Migration
- When Defining a Rule Based on the Approval Group List Builder During Runtime, a Message Says the Group Does Not Exist

21.5.1 Task Detail Page is Not Available

Problem

When you click the human task in Oracle BPM Worklist, the following message is displayed:

```
Task detail is not available
```

Solution

This problem typically occurs during the deployment of the ADF J2EE application containing the task details page definition. As a result, the URI for the detail page is not registered with the SOA cluster.

To confirm whether the task details page was registered in Fusion Applications Control:

1. Log in to the domain where the composite with the subscription is deployed.

2. In the navigation pane, right-click **SOA** > **soa-infra** (*SOA_cluster_name*).

3. Choose **Service Engines** > **Human Workflow**.

4. In the **Components** section, click the appropriate task.

5. Click the **Administration** tab to see a URI for the task details page.

   If there is no URI displayed, then check the log files for any errors during deployment.

   A possible problem may be incorrect entries in the `wf_client_config.xml` file. The `wf_client_config.xml` file resides in the exploded EAR file's `APP-INF/classes` directory. For example:

   ```
   APPLICATIONS_BASE/fusionapps/applications/fin/deploy/EarFinPayables.ear/
   APP-INF/classes/wf_client_config.xml
   ```

   The file contains the cluster URL information for the local SOA runtime of that particular Oracle WebLogic Server domain. If the cluster name is wrong or that cluster does not exist in the domain, then a problem exists with the deployment and configuration of the environment.

   As a workaround, you can directly enter the task detail page URI in Fusion Applications Control. See the "Managing the URI of the Human Task Service Component Task Details Application" section of the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*. For the values, enter the following:

Field	Value
Application Name	`worklist`
Host Name	Your server host name for SOA
HTTP Port	Your HTTP server port used for SOA if SSL is disabled.
HTTPS Port	Your HTTPS server used for SOA if SSL is enabled.
URI	For example, for financials:

	`/payables/faces/adf.task-flow?_id=FinApInvoiceApprovalTaskFlow&_`
	`document=WEB-INF/oracle/apps/financials/payables/invoices/transact`
	`ions/ui/invoiceApprovalTask/flow/FinApInvoiceApprovalTaskFlow.xml`

21.5.2 Task Chooses an Incorrect User If a Number of Rules Are Defined or Errors with "Ruleset returned lists with different list builder"

Problem

At runtime, when a human workflow task tries to fetch the list of users, it errors with the following message:

```
Ruleset returned lists with different list builder
```

This can be seen in the **Task Detail Comments** field. Alternately, it may select a user or approver, which may not seem to be the correct or expected one. This is primarily caused by having overlapping rules with different list builders. When participants of a task are specified using business rules, it is expected that business rules return list builders of the same type.

Moreover, only one rule from a ruleset should be applicable for a transaction. In case a number of rules are true, the actions associated with the applicable rule with the highest priority get executed. If multiple applicable rules have the same priority, then the first rule in the list is picked and its actions executed.

Solution

To resolve this issue, avoid writing overlapping rules. Constraints from different list builders are different and cannot be mixed. If more than one rule gets triggered with a different list builder, this error occurs. Moreover, only one set of constraints is honored.

Check that all rules in the ruleset have priorities defined so that multiple rules with the same priority are not applicable for the same transaction. For more details, see the "Using Approval Management" chapter in the *Oracle Fusion Middleware Modeling and Implementation Guide for Oracle Business Process Management*.

21.5.3 Document Is Stuck in Pending Approval and Logs Show Error "Unable to walk hierarchy for user *user_name*"

Problem

The document is stuck pending approval and the logs show the following error:

```
"Unable to walk hierarchy for user username'
```

This occurs when the hierarchy for the given user has not been set up correctly. For example, assume you have approval rules set up so that the rule getting fired uses the supervisory hierarchy and you expect it to route to USER01 > USER02 > USER03 > . . . > USER10 (ten level hierarchy). If the supervisory/management hierarchy of USER01 has not been set up, then the approval task throws an error stating the inability to go through the hierarchy for USER01.

Solution

Since the user level hierarchy does not exist for the `user_name`, verify the Human Capital Management (HCM) setup by manually running the HCM service to which `workflow-identity-config.xml` is pointing. If it does not return any result, then configure the correct hierarchy for the user in HCM.

Run the HCM service, select the `fetchManager`(s) API (for supervisory), and provide the following payload, where `<ns2:Id>DD7A1614BBFAA0F0A4511ACD96D2C88D</ns2:Id>` is the GUID of the user whose manager you are trying to find. The supervisory hierarchy and position hierarchy services can be found under `HcmCore`. The name is `HierarchyProviderService`. The corresponding WSDLs are as follows:

- Supervisory hierarchy:

  ```
  http://host:port/hcmEmploymentCoreApprovalHierarchy/HierarchyProviderService?ws
  dl
  ```

- Position hierarchy:

  http://*host*:*port*/hcmTreesModel/HierarchyProviderService?wsdl

```
<soap:Envelope xmlns:soap="http://schemas.xmlsoap.org/soap/envelope/">
<soap:Header>
<wsse:Security soap:mustUnderstand="1"
xmlns:wsse="http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-wssecurity-sec
ext-1.0.xsd">
<wsse:UsernameToken>
<wsse:Username>FUSION_APPS_AMX_APPID</wsse:Username>
<wsse:Password
Type="http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-username-token-profi
le-1.0#PasswordText">Welcome1</wsse:Password>
</wsse:UsernameToken>
</wsse:Security>
</soap:Header>
<soap:Body
xmlns:ns1="http://xmlns.oracle.com/apps/hcm/employment/core/approvalHierarchyServi
ce/types/">
<ns1:fetchManagers>
<ns1:principal
xmlns:ns2="http://xmlns.oracle.com/apps/hcm/employment/core/approvalHierarchyServi
ce/">
<ns2:Id>DD7A1614BBFAA0F0A4511ACD96D2C88D</ns2:Id>
</ns1:principal>
<ns1:level>2</ns1:level>
</ns1:fetchManagers>
</soap:Body>
</soap:Envelope>
```

The following example shows the `workflow-identity-config.xml` file used in a development environment:

```
<?xml version='1.0' encoding='UTF-8'?>
<ISConfiguration xmlns="http://www.oracle.com/pcbpel/identityservice/isconfig">
<configurations>
<configuration realmName="myrealm">
<provider name="JpsProvider" providerType="JPS" service="Identity">
<property name="jpsContextName" value="default"/>
<property name="IdentityServiceExtension" value="HCMIdentityServiceExtension"/>
</provider>
</configuration>
</configurations>
<property name="caseSensitive" value="false"/>
<property name="caseSensitiveGroups" value="true"/>
<serviceExtensions>
<serviceExtension name="HCMIdentityServiceExtension">
<serviceProvider
classname="oracle.bpel.services.identity.hierarchy.providers.hcm.HCMHierarchyProvi
der" type="supervisoryHierarchyProvider">
<initializationParameter name="wsdlUrl"
value="http://iashaqa05.us.oracle.com:20619/hcmEmploymentCoreApprovalHierarchy/Hie
rarchyProviderService?WSDL"/>
<initializationParameter name="csf-key-name" value="FUSION_APPS_AMX_APPID-KEY"/>
<initializationParameter name="http-read-timeout" value="30000"/>
<initializationParameter name="securityPolicyName" value="oracle/wss_username_
token_client_policy"/>
</serviceProvider>
<serviceProvider
classname="oracle.bpel.services.identity.hierarchy.providers.hcm.HCMHierarchyProvi
```

```
der" type="jobLevelHierarchyProvider">
<initializationParameter name="wsdlUrl"
value="http://iashaqa05.us.oracle.com:20619/hcmEmploymentCoreApprovalHierarchy/Hie
rarchyProviderService?WSDL"/>
<initializationParameter name="csf-key-name" value="FUSION_APPS_AMX_APPID-KEY"/>
<initializationParameter name="http-read-timeout" value="30000"/>
<initializationParameter name="securityPolicyName" value="oracle/wss_username_
token_client_policy"/>
</serviceProvider>
<serviceProvider
classname="oracle.bpel.services.identity.hierarchy.providers.hcm.HCMPositionHierar
chyProvider" type="positionHierarchyProvider">
<initializationParameter name="wsdlUrl"
value="http://iashaqa05.us.oracle.com:20619/hcmTreesModel/HierarchyProviderService
?WSDL"/>
<initializationParameter name="csf-key-name" value="FUSION_APPS_AMX_APPID-KEY"/>
<initializationParameter name="http-read-timeout" value="30000"/>
<initializationParameter name="securityPolicyName" value="oracle/wss_username_
token_client_policy"/>
</serviceProvider>
<serviceProvider
classname="oracle.bpel.services.identity.position.provider.hcm.PositionLookupServi
ceProvider" type="positionLookupProvider">
<initializationParameter name="wsdlUrl"
value="http://iashaqa05.us.oracle.com:20619/hcmEmploymentCore/positionLookupServic
e?WSDL"/>
<initializationParameter name="csf-key-name" value="FUSION_APPS_AMX_APPID-KEY"/>
<initializationParameter name="http-read-timeout" value="30000"/>
<initializationParameter name="securityPolicyName" value="oracle/wss_username_
token_client_policy"/>
</serviceProvider>
<serviceProvider
classname="oracle.bpel.services.identity.position.provider.hcm.PositionDisplayName
Provider" type="positionDisplayNameProvider">
<initializationParameter name="wsdlUrl"
value="http://iashaqa05.us.oracle.com:20619/hcmTreesModel/HierarchyProviderService
?WSDL"/>
<initializationParameter name="csf-key-name" value="FUSION_APPS_AMX_APPID-KEY"/>
<initializationParameter name="http-read-timeout" value="30000"/>
<initializationParameter name="securityPolicyName" value="oracle/wss_username_
token_client_policy"/>
</serviceProvider>
</serviceExtension>
</serviceExtensions>
</ISConfiguration>
```

The `workflow-identity-config.xml` file is in the MDS repository. For information about how to export the `workflow-identity-config.xml` file from MDS, see Section 21.5.5.

21.5.4 Task History: Cannot Remove or Move a Participant After Yourself

Problem

You cannot remove or move a participant that was added through future participant editing. This is because the participant was anchored to you and you are the current assignee.

Solution

This feature is not supported.

21.5.5 Hierarchy List Builders Display an Error Due to Security Configuration Issues and Logs Report Errors in Authentication or Authorization with HCM Services

Problem

Hierarchy list builders (for example, supervisory, job-level, and position) used in rules are not working as expected and are throwing security configuration-related errors. For example:

```
SOAPFaultException: FailedAuthorization
```

Hierarchy services are protected services using Oracle Web Services Manager policies (for example, supervisory, job-level, and position services). Elevated privileges are used for authentication. These privileges must be configured in Fusion Applications Control and in configuration files.

Solution

To resolve this problem:

1. Check the MDS store to see if `workflow-identity-config.xml` and `workflow-config.xml` are set up with the correct information (pointing to HCM services and the JAZN name).

   1. Run the following script to connect to the `WLST` shell on Linux operating systems. There is no equivalent script on Windows operating systems.

      ```
      $MW_HOME/oracle_common/common/bin/wlst.sh
      ```

   2. Run `connect()`. This takes you to the prompt as shown in the following example:

      ```
      wls:/offline> connect()
      Please enter your username [weblogic] :weblogic
      Please enter your password [weblogic] :
      Please enter your server URL [t3://localhost:7001]
      :t3://adcdai01.us.oracle.com:9401
      ```

   3. Run the following command after connecting:

      ```
      exportMetadata(application='soa-infra',server='soa_server1',toLocation='any
      Location on
      server',docs='/soa/configuration/default/workflow-identity-config.xml')
      ```

      The file is stored in the `soa/configuration/default` directory in the location given in the `toLocation` attribute. After verifying, you can upload the file with the following command.

      ```
      importMetadata(application='soa-infra',server='soa_
      server1',fromLocation='any Location on
      server',docs='/soa/configuration/default/workflow-identity-config.xml')
      ```

2. Verify that the `csf-keys` are specified as initialization parameters.

   a. Ensure that the `csf-key-name` value is `FUSION_APPS_AMX_APPID-KEY` in `workflow-identity-config.xml`, as described in Section 21.5.3.

3. Check in Fusion Applications Control that the keys defined for the services have the correct values for authentication.

   a. Run the following script to connect to the `WLST` shell:

   ```
   $MW_HOME/oracle_common/common/bin/wlst.sh

   wls:/offline> connect()
   Please enter your username [weblogic] :weblogic
   Please enter your password [weblogic] :
   Please enter your server URL [t3://localhost:7001]
   :t3://myhost.us.oracle.com:9401
   ```

   b. Run the following command after connecting:

   ```
   listCred(map="oracle.wsm.security", key="FUSION_APPS_AMX_APPID-KEY")
   ```

   Output similar to the following is displayed:

   ```
   [Name : FUSION_APPS_AMX_APPID,
   Description : HCM Hierarchy Service, expiry
   Date : null]
   PASSWORD:password
   ```

21.5.6 HCM Services are Timing Out or Getting Stuck

Problem

Hierarchy services sometimes get overloaded and do not respond (for example, supervisory, job-level, and position services).

Solution

Check that the `workflow-identity-config.xml` file for the problematic service has correct timeout settings defined as initialization parameters. The timeout parameter is `http-read-timeout`. For an example of the configuration file, see Section 21.5.3. Note that increasing the timeout may cause errors, such as other timeouts and only puts more load on an already overloaded server.

21.5.7 Error/Warning Reported in Retrieving the Position for a User

Problem

Not all LDAP users are defined as people within Oracle Fusion Human Capital Management. For example, users may be defined in LDAP for Oracle Fusion Supply Chain Management that are not defined as people with positions in Oracle Fusion Human Capital Management. Therefore, the position lookup can generate a message that the user does not have positions.

Solution

This is a tolerated condition and the warning can be ignored if the user name specified in the warning is not an HCM user. If this is an HCM user, the problem is on the HCM side. Add a position to these users in the HCM application.

21.5.8 Users See No Approvals in the Inbox when Using Position Hierarchy

Problem

This issue may be noticed by organizations using position hierarchy such as governments and universities. After users log in, they do not see any approvals in their inbox. This is caused by a failure to look up the position. This can be caused by network issues when accessing the HCM services. Sometimes due to load balancing or network setup, HCM services may not be accessible and Oracle BPM Worklist has errors such as the following:

```
WSDL port is not accessible
```
or
```
Could not find the host
```

Solution

1. Check the network setup and ensure that the HCM services are accessible and WSDL ports can be accessed in the browser. Tools like `ping` and `traceroute` can also diagnose the problem.

2. If the services are not available, bring up the services and retry.

3. Restart the Human Capital Management server and SOA server. For information, see the following documentation:

   - "Starting Managed Servers with the java weblogic.Server Command" section in *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server*

   - "Starting and Stopping Managed Servers Using Fusion Middleware Control" section in *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server*

   - "Start and Stop Servers" and various startup and shutdown procedures in the Cluster section of the *Administration Console Online Help*.

21.5.9 Task Escalation Does Not Follow Job Level, Supervisory, or Position Hierarchies

Problem

This feature is not supported. Escalation does not follow the job level, supervisory, or position hierarchies.

Solution

Build your own custom plug-in. For more details, see *Oracle Fusion Middleware Developer's Guide for Oracle SOA Suite*.

21.5.10 User Is Not Allowed to Perform an Action on a Human Workflow Task

For troubleshooting information about this issue, see the "Task Action Issues" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

21.5.11 Email Notification Is Not Being Sent Out

For troubleshooting information about this issue, see the "Notification Issues" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

For additional details about configuring email driver properties, see the "Configuring Email Notifications for Oracle Fusion Applications" section of the *Oracle Fusion Applications Postinstallation Guide*.

21.5.12 Notifications Are Sent But Are Not Actionable

For troubleshooting information about this issue, see the "Notification Issues" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*. This section describes how to make the task actionable in the Human Task Editor during design time.

In addition, if notifications are sent, but are not actionable, the administrator may have configured the notification to be nonactionable through Oracle BPM Worklist. To make notifications actionable, select the **Make notification actionable** checkbox under the **Event Driven** subtab of the **Task Configuration** tab.

For additional details, see the "Configuring Email Notifications for Oracle Fusion Applications" section of the *Oracle Fusion Applications Postinstallation Guide*.

21.5.13 Actionable Notifications Are Sent But No Action is Taken After Responding

For troubleshooting information about this issue, see the "Notification Issues" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

21.5.14 After LDAP Configuration with Oracle WebLogic Server, Users Appear in the Console, But a "No Role found matching the criteria" Error Appears

For troubleshooting information about this issue, see the "Identity Service Issues" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

21.5.15 Only a Subset of Users in LDAP Can Log In to Oracle BPM Worklist

For troubleshooting information about this issue, see the "Identity Service Issues" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

21.5.16 "Unknown Macro" Exception Appears when Myrealm Is Passed as a Parameter to the Identity Context To Identity Service APIs

For troubleshooting information about this issue, see the "Identity Service Issues" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

21.5.17 Add Participant Button in Oracle BPM Worklist Is Disabled

For troubleshooting information about this issue, see the "Task History Issues" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

21.5.18 Task Modifications Made in a Runtime Tool Are Not Appearing for the Task

For troubleshooting information about this issue, see the "Design Time at Runtime Issues" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

21.5.19 How Can I Set Commonly Used Human Workflow Configuration Parameters?

For troubleshooting information about this issue, see the "Human Workflow Service/System MBean Browser Issues" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

21.5.20 How Do I Set Human Workflow Configuration Parameters Not Available in the Fusion Applications Control Properties Pages?

For troubleshooting information about this issue, see the "Human Workflow Service/System MBean Browser Issues" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

21.5.21 System MBean Browser Does Not Reflect My Changes After Editing the Human Workflow Config MBeans

For troubleshooting information about this issue, see the "Human Workflow Service/System MBean Browser Issues" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

21.5.22 Human Workflow Services Are Not Locating Resource Bundles or Classes Located at the Workflow Customizations Class Path URL

For troubleshooting information about this issue, see the "Human Workflow Service/System MBean Browser Issues" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

21.5.23 How Do I Manually Set the URL for Displaying Task Details for a Particular Human Workflow Task Component?

For troubleshooting information about this issue, see the "Human Workflow Service/System MBean Browser Issues" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

21.5.24 How to Test the Health of the Installed Server

For troubleshooting information about this issue, see the "Test-to-Production Issues" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

21.5.25 User Authentication Error (in this example, the user is FMW_USERID and identityContext is jazn.com)

For troubleshooting information about this issue, see the "Test-to-Production Issues" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

21.5.26 Error During Import of Task Payload Mapped Attribute Mappings

For troubleshooting information about this issue, see the "Test-to-Production Issues" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

21.5.27 Error During Rule Migration

For troubleshooting information about this issue, see the "Test-to-Production Issues" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

21.5.28 When Defining a Rule Based on the Approval Group List Builder During Runtime, a Message Says the Group Does Not Exist

For troubleshooting information about this issue, see the "AMX Extension Issues" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

21.6 Patching and Deployment

For troubleshooting information about patching and deployment issues, see the "Troubleshooting Patching Sessions for Service Oriented Architecture (SOA) Composites" in the *Oracle Fusion Applications Patching Guide*.

21.7 Performance

This section contains the following topics that describe common problems and solutions for Oracle SOA Suite performance:

- SOA Transactions are Failing and Logs Indicate the SOA Database is Running Out of Space

- Slow Application Performance Such as Longer Time to Serve Pages or Finish Transactions

- Slow Fusion Applications Control Performance

21.7.1 SOA Transactions are Failing and Logs Indicate the SOA Database is Running Out of Space

Problem

SOA transactions are failing and logs indicate that the database is running out of space. The problem may occur with the SOA dehydration store or MDS store running out of space due to a high volume of transactions. In the latter example, you see errors such as the following:

```
java.sql.SQLException: ORA-01653: unable to extend table SH_MDS.CUBE_INSTANCE by
16 in tablespace FUSION_TS_TOOLS
```

This indicates that the tablespace is full and the database cannot extend it.

Solution

1. Purge the SOA dehydration store tables periodically, taking into account the appropriate record retention policies and ensuring that the applications have no dependencies on runtime data.

    The purge should be followed by commands to coalesce the space. For the purging strategy to work, it is important to understand how long to retain the data in the database. Factors that drive the retention policy include the following:

    -   Legal requirements

- Line of business requirements

- Overall company policy on retention of data

The longer the retention policy, the greater the volume of data that must be stored and, correspondingly, the higher the disk capacity requirements.

For details on creating a purging strategy, see the "Managing Database Growth" chapter of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

2. Ensure that the database hardware has sufficient resources to handle the demands of Oracle Database partitioning prior to configuring your tables for partitioning.

   For the SOA dehydration store, database partitioning using range partitioning and hash partitioning is an optimal solution. Partitioning by definition means storing data in multiple tables to reduce bigger data sets into smaller, more manageable data sets. Partitioning strategies play a large role in easing maintenance overheads (dropping and pruning the partition) and improving performance. Partitioning should at least be done for tables having high activity. This plays a large role in balancing disk I/O and preventing hot disks. One important requirement that you must meet prior to configuring your tables for partitioning is to ensure that the database hardware has sufficient resources to handle the demands of Oracle Database partitioning. If preproduction testing has indicated that the installation is large, Oracle expects that you have sized your environment (CPU, memory, and disk space) correctly to take advantage of the partitioning features.

3. Tune database parameters for memory, tablespace, and partitions to get maximum performance. For more information, see the "Tune Database Parameters" section in the *Oracle Fusion Middleware Performance and Tuning Guide*.

4. For other tablespaces running out of space, use the following query to check for free tablespace:

```
SELECT Total.tablespace_name "Tablespace Name", Free_space, (total_space-Free_
space) Used_space, total_space, round((Free_space*100/total_space),2) "Free %"
FROM (select tablespace_name, sum(bytes/1024/1024) Free_Space from sys.dba_
free_space group by tablespace_name) Free,
(select tablespace_name, sum(bytes/1024/1024) TOTAL_SPACE from sys.dba_data_
files group by tablespace_name) Total
WHERE Free.Tablespace_name = Total.tablespace_name AND Total.tablespace_name =
 '<tablespacename>'
ORDER BY 5;
```

5. To increase tablespace settings, use the administrator account. For example:

```
alter tablespace tablespace_name add datafile 'datafile_name' size 500m
autoextend on;
```

For more details, see the "Resolving Message Failure Caused by Too Many Open Files" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

21.7.2 Slow Application Performance Such as Longer Time to Serve Pages or Finish Transactions

Problem

You observe slow application performance and/or memory trashing. For example, it may take longer to load and serve pages or to complete composite transactions. Response time may seem slower compared to normal behavior.

There are various reasons for slow performance. It may be due to a large number of servers running on the same host, or there may be a large number of records/sessions/locking at the database. Thread contention can also be a reason for poor performance.

> **Note:** All configuration changes in Fusion Applications Control may be unavailable due to this problem.

Solution

1. Check the CPU utilization to see if it is saturated due to a heavy load or an excessive number of processes in relation to CPU capacity.

   If CPU utilization is at 100% during normal load hours (the target should be 70-80%), you have no capacity to handle a peak load and the hardware resources are insufficient. Add scale-out servers to handle the additional load.

2. Check applications using Fusion Applications Control to report on performance. Check the performance of ADF services and the invoke/response times in the BPEL process audit trail to identify if the issue is in Oracle Fusion Applications.

3. To maximize performance, it is recommended that you not set the logging level higher than the default **INFO** level.

   For debugging purposes, you must set the logging level to the **FINEST** level. However once issues are resolved, reset the logging level to the default level for best performance. It is also recommended that you set **Audit Level** to **Production** in the SOA Infrastructure Common Properties page. This can be set in Fusion Applications Control as follows:

   a. Right-click **soa-infra (*SOA_cluster_name*)**.

   b. Choose **SOA Administration** > **Common Properties**.

   c. Set **Audit Level** to **Production**.

4. Purge periodically based on retention requirements to maintain the service level agreements (SLAs).

   To identify the tables where data growth may lead to performance issues, see the "Tables Impacted By Instance Data Growth" section in the *Oracle Fusion Middleware Performance and Tuning Guide*.

   For information about the use of the purge scripts, see the "Managing Database Growth" chapter in the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

5. Check the memory/IO/paging/swapping/CPU usage load stats using Top or Glance or another monitoring tool.

6. Optimize the JVM to avoid full garbage collection or out-of-memory errors.

Frequent garbage collection can be either due to higher memory usage or memory leaks.

a. Ensure that the sum of the maximum heap size of all the JVMs running on your system does not exceed the amount of available physical RAM to avoid operating system level paging activity.

b. Use the JRockit mission control memory profiling tools to get thread dumps and memory snapshots, which helps Oracle Support Services debug any code issues.

7. Optimize threads to avoid contention. Get a thread dump to investigate and submit it to Oracle Support Services.

8. Check the JVM and thread dumps for methods invoking the database to identify if database performance is a bottleneck.

9. Run database Automatic Workload Repository (AWR) snapshots to identify causes of database performance issues.

10. Ensure that database statistics are updated at regular intervals and other tunable parameters for memory, tablespace, and partitions are used effectively to obtain maximum performance.

   Here are some common tuning recommendations. For more information, see the "Tune Database Parameters" section in the *Oracle Fusion Middleware Performance and Tuning Guide*.

   ■ Put indexes and tables in as physically separate disk areas if possible.

   ■ Never put rollback segments with data or index segments.

   ■ Separate highly active tables and indexes into their own tablespaces.

   ■ Partition high activity tables and indexes to help balance disk I/O and prevent hot disks.

   ■ Have processes in place to generate database table statistics at regular intervals.

11. Tune database tables to control the high watermark (HWM) contention of large objects. Tune database advanced queues (AQ) to control high watermark (HWM) contention and ensure consistent performance of producing and consuming messages from AQ.

   Specifically, the AQs to be aware of are as follows:

Family AQ	Queue Name
EDN	`EDN_EVENT_QUEUE_TABLE` and `EDN_OAOO_DELIVERY_TABLE`
CRM	`CRM_AsyncWS_Request` and `CRM_AsyncWS_Response`
HCM	`HCM_AsyncWS_Request` and `HCM_AsyncWS_Response`
FIN	`FIN_AsyncWS_Request` and `Fin_AsyncWS_Response`
PRC	`PRC_AsyncWS_Request` and `PRC_AsyncWS_Response`
PRJ	`PRJ_AsyncWS_Request` and `PRJ_AsyncWS_Response`
SCM	`SCM_AsyncWS_Request` and `SCM_AsyncWS_Response`
COMMON	`COMMON_AsyncWS_Request` and `COMMON_AsyncWS_Response`
IC	`IC_AsyncWS_Request` and `IC_AsyncWS_Response`

Family AQ	Queue Name
Cross Family Business Event Subscriptions	`ACR_XFAMILY_EVENT_Q` and `ACR_XFAMILY_EVENT_QT`

12. Tune the BPEL and EDN thread counts to ensure optimal settings. If the thread configuration is too high, the servers run out of memory. If they are too low, the messages start backing up.

13. In case of integration with packaged applications (for example, Siebel), check if the issue lies with the legacy applications.

For more information about performance tuning the various components, see the "Top Performance Areas" chapter and the "SOA Suite Components" part in the *Oracle Fusion Middleware Performance and Tuning Guide*.

21.7.3 Slow Fusion Applications Control Performance

Problem
Fusion Applications Control pages are loading very slowly.

Solution
To improve Fusion Applications Control page load times on pages with instance and fault metrics, set the SOA Infrastructure **Display Data Counts** properties.

1. Right-click **soa-infra (*SOA_cluster_name*)**.

2. Choose **SOA Administration** > **Common Properties**.

3. In the **Display Data Counts** section, deselect the **Disable fetching of instance and fault count metrics** checkbox.

4. Reduce the **Duration Details** value.

5. Click **Apply**.

For more information, see the "Optimizing the Loading of Pages with Instance and Fault Metrics" section in the *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

21.8 Maintenance

This section contains the following topics that describe common problems and solutions for Oracle SOA Suite maintenance:

- JVM is Suspended If Any In-use Log Directories are Deleted
- Server Start Script Does Not Work for Scale-Out (Postprovisioning) Servers
- SOA Server Does Not Start
- SOA Server Is Not Responding To Administrative Commands
- Undeploying a Composite Left in an Inconsistent State Due to a Failed Patch
- Increased Load Observed on RAC Instances when Using Multiple Oracle RAC Servers
- List of SOA Composite State Values
- Log Files Have Been Deleted or Are Missing

21.8.1 JVM is Suspended If Any In-use Log Directories are Deleted

Problem

JRockit writes log information (called flight recorder files) under /tmp when using JRockit Mission Control. When it rotates to a new flight recorder file, the JVM is suspended until it completes. If the old directories in use under /tmp are deleted (for example, a cleanup or custom maintenance script to run to clear /tmp when it is close to full to avoid server downtime), the domain may go into a suspended state because the JVM may wait for a directory that no longer exists.

Solution

1. Change any cleanup scripts/procedures to first shut down the domain cleanly.

2. Perform a cleanup of /tmp.

3. Restart the domain.

To recover existing, suspended JVMs without restarting them, you must recreate the directories that the JVM was trying to use.

21.8.2 Server Start Script Does Not Work for Scale-Out (Postprovisioning) Servers

Problem

Scale-out servers, created after provisioning, cannot be started using back-end commands. The script start_fusion_Allservers.sh only starts servers that are created during provisioning time.

Solution

To start scale-out servers, use Fusion Applications Control. See Section 3.3.5.3.

21.8.3 SOA Server Does Not Start

Problem

The SOA server (soa_infra) does not start.

Solution

There can be various reasons for server startup issues. A discovery-based approach to find the root cause is required. Check the server and diagnostic logs as a first attempt to diagnose the issue. In addition, check the following:

1. Check if the database is not available or there are not enough connections available. Some failures with SOA runtime can result from database outage/connectivity issues. Perform the following steps:

   a. Log in to Oracle WebLogic Server Administration Console.

   b. In the **Domain Structure**, view the status by selecting **Services** > **Data Sources** > **SOADatasource** > **Monitoring** > **Testing** to test the data source. You can also check this from a SQL prompt.

2. Check the list of ports used for port conflicts.

3. Check if the MDS repository is unable to load shared documents (incorrect MDS configuration or the database that holds the MDS schema is not reachable).

4. Check the Coherence configuration if using a SOA cluster (use of unicast versus multicast).

This issue only applies after provisioning is complete. Clustering of SOA servers may fail if there are port conflicts (used by Coherence). For example, one SOA server may be picking up the deployment, but the other server is not. In some cases, this may only present the following error:

```
[soa_server1] [ERROR] [] [Coherence] [tid: Logger@352821903 3.6.0.4]
[ecid: 46f620208907e045:63f295ec:12dd091ec2e:-8000-0000000000000003,1:27187]
[APP: soa-infra] 2011-01-28 23:06:19.463/414.816 Oracle Coherence GE 3.6.0.4
<Error> (thread=[ACTIVE] ExecuteThread: '0' for queue: 'weblogic.kernel.Default
 (self-tuning)', member=n/a):
Error while starting cluster: com.tangosol.net.RequestTimeoutException: Timeout
 during service start: ServiceInfo(Id=0, Name=Cluster, Type=Cluster[[...
```

In this example, Coherence timeouts are prominently available. However, sometimes it presents itself with an unrelated error, such as the following:

```
Error creating bean with name 'SensorManager' defined in ServletContext
resource
```

The root cause of this is still primarily related to Coherence configuration.

5. Check if Managed Server startup failed because the Administration Server is not reachable.

6. Check network issues (for example, IP routing filtering/rules that may be causing issues).

7. Check Oracle WebLogic Server LDAP security corruption.

The Managed Server may report that policies for the application System MBeans Browser already exist or do not exist.

For more information, see the "Setting the Frontend URL for the Administration Console and Setting Redirection Preferences" section and the "Setting the Frontend HTTP Host and Port" section of *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle SOA Suite*.

21.8.4 SOA Server Is Not Responding To Administrative Commands

Problem

The SOA server (`soa_infra`) is not responding to administrative commands. This problem can be due to your environment using a load balancer, and the default listener address or channel is local-only. Therefore, the connection cannot be made from the external client or server.

Solution

To resolve this problem, perform the following steps from the Oracle WebLogic Server Administration Console:

1. Ensure that the listener address is correctly defined and accessible.

    a. From the **Domain Structure**, expand **Environment** and select **Servers**.

    b. From the **Servers** table, click **AdminServer(admin)**.

    c. From the Settings for AdminServer page, click the **Protocols** tab.

    d. Click the **HTTP** subtab.

    e. Set the **Frontend Host** property to the load balancer address.

2. Similarly, set the **Frontend HTTP Host** and **Frontend HTTP Port** properties for the SOA cluster:

    a. From the **Domain Structure**, expand **Environment** and select **Clusters**.

    b. From the **Clusters** table, click *domain*_**SOACluster**

    c. From the Settings for the *domain*_SOACluster page, click the **HTTP** subtab.

    d. Set the **Frontend Host** property to the load balancer address.

    e. Change the **Frontend HTTP Host** and **Frontend HTTP Port** properties.

For more information, see the "Setting the Frontend URL for the Administration Console and Setting Redirection Preferences" section and the "Setting the Frontend HTTP Host and Port" section of *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle SOA Suite*.

21.8.5 Undeploying a Composite Left in an Inconsistent State Due to a Failed Patch

For troubleshooting information about this issue, see the "Troubleshooting SOA Composite Deployment Failures" section of the *Oracle Fusion Applications Patching Guide*.

21.8.6 Increased Load Observed on RAC Instances when Using Multiple Oracle RAC Servers

Problem

An increased load is seen on Oracle Real Application Clusters (Oracle RAC) instances when using multiple Oracle RAC servers.

Solution

Use a multidata source against Oracle RAC instances using load balancing rather than failover. This ensures that the load is evenly distributed rather than all load staying on one node until it fails over, at which point all the connections are failed over to the other node.

Ensure that you set the **Connection Test Frequency Seconds** property at the multidata source level to a nonzero value. A value that is too low means more load on the listeners/RAC instances when multiple servers attempt to reconnect at the same time.

For more information about configuring JDBC multidata sources, see *Oracle Fusion Middleware Configuring and Managing JDBC Data Sources for Oracle WebLogic Server*.

> **Note:** The Oracle Fusion Applications default configuration for the these settings has been tuned for best performance. Only change these settings when the configuration has *not* been set out-of-the-box by Oracle Fusion Applications.

21.8.7 List of SOA Composite State Values

Problem

Get a list of all composite state values from the *_SOAINFRA.COMPOSITE_INSTANCE state table (for diagnosis).

Solution

Table 21–3 shows the composite state values.

Table 21–3   Composite State Values

State	Description
0	Running
1	Completed
2	Running with faults
3	Completed with faults
4	Running with recovery required
5	Completed with recovery required
6	Running with faults and recovery required
7	Completed with faults and recovery required
8	Running with suspended
9	Completed with suspended
10	Running with faults and suspended
11	Completed with faults and suspended
12	Running with recovery required and suspended
13	Completed with recovery required and suspended
14	Running with faults, recovery required, and suspended
15	Completed with faults, recovery required, and suspended
16	Running with terminated
17	Completed with terminated
18	Running with faults and terminated
19	Completed with faults and terminated
20	Running with recovery required and terminated
21	Completed with recovery required and terminated
22	Running with faults, recovery required, and terminated
23	Completed with faults, recovery required, and terminated
24	Running with suspended and terminated
25	Completed with suspended and terminated
26	Running with faulted, suspended, and terminated
27	Completed with faulted, suspended, and terminated
28	Running with recovery required, suspended, and terminated
29	Completed with recovery required, suspended, and terminated
30	Running with faulted, recovery required, suspended, and terminated
31	Completed with faulted, recovery required, suspended, and terminated
32	Unknown
64	State

21.8.8 Log Files Have Been Deleted or Are Missing

Problem

The SOA log files are getting deleted. This is typically due to the retention policy set in Fusion Applications Control.

Solution

Change the retention policy for the `odl-handler` to the appropriate value in Fusion Applications Control:

1. Right-click **soa-infra**.

2. Choose **Logs** > **Log Configuration**.

21.9 Custom Development (Extensibility)

This section contains the following topics that describe common problems and solutions for Oracle SOA Suite runtime:

- Unable to Use Layered Customization on Specific Artifacts

- New or Deleted Artifacts Do Not Appear as a Layered Customization

- Warnings Regarding xml:id when Building or Deploying a Customized Composite

- Distributed Queue Topic Messages Are Retrieved Multiple Times by Subscribers

- Internal Key is Displayed Instead of a Translated String

- BPEL Activity Errors and the Log Shows Error "com.oracle.bpel.entity.dataprovider.EntityVarMgrException:zero Item"

- BPEL Activity Errors and Log Shows Error "XPath variable or expression <expression> is empty"

- Unavailable Composite Errors Occurring Even Though the Target Service Is Up and Running

- High-Volume Cross Reference (XREF_DATA) Table Impacts Performance and Maintenance

- Access Denied Error While Invoking an Oracle ADF BC Service from BPEL

- Clicking Manage Approvals Page Gives a "SOA server may be down" or "No data to display" Error

- Synchronous Service Invocation Errors Due to WS-Addressing

- Deploying Human Workflow Application Throws "Unable to resolve 'TaskQueryService'" Error

- "Invalid Subject" Error Thrown with Human Workflow API or Notification

- Task is Assigned to the Group/Role when It Is Expected to Go to Every User in the Group/Role Individually

- Task Completes Without Any Assignment Occurring

- Parallel Assignees Must Approve or Reject the Task Even Though the Parallel Completion Criteria is Met

- All Added Adhoc Participants Disappear After a Page Refresh

- Future Approvers Are Not Visible in the History Table

- [Message Appears in the History Table about a Correlation ID Not Being Passed or Any Exception Related to the Correlation ID](#)
- [Edit Toolbar Is Disabled or Not Shown](#)
- [oracle.jrf.UnknownPlatformException Error When Customizing an Oracle Mediator](#)
- [java.lang.NullPointerException When Customizing an Oracle Mediator](#)
- [JDeveloper Compilation Error in a SOA Project with SOA MDS Service Location](#)

21.9.1 Unable to Use Layered Customization on Specific Artifacts

Problem

Only certain artifacts allow layered customization in Oracle JDeveloper because it requires editors that use XMLEF to make delta changes to the underlying DOM. XMLEF is the XML Editing Framework. XMLEF is an Oracle JDeveloper framework that supports flexible editing of XML documents and is used extensively in layered customizations.

Solution

The following SOA artifact types are customizable:

- Composite (`composite.xml`)
- BPEL process (`.bpel`)
- Oracle Mediator (`.mplan`)

The following SOA artifact types are not customizable:

- XSL map (`.xsl`)
- Human task (`.task`)
- Decision service (`.decs`)
- Rules (`.rules`)
- Event definition (`.edl`)
- Domain value map (`.dvm`)

The following common artifact types are not customizable because Oracle XDK/WSDL APIs do not support customization (although their editors do use XMLEF to make delta changes to the underlying DOM):

- XSD schema (`.xsd`)
- WSDL document (`.wsdl`)

Noncustomizable artifacts are not sensitive to a customization context. Any changes to those artifacts from any customization context are visible to other customization contexts.

For more information about customizations in Oracle JDeveloper and runtime tools (such as Oracle SOA Composer and Oracle BPM Worklist), see the "Customizing and Extending SOA Components" chapter of *Oracle Fusion Applications Extensibility Guide*.

21.9.2 New or Deleted Artifacts Do Not Appear as a Layered Customization

Problem

Because MDS customization does not capture artifact creations or deletions as deltas in the customization role, artifact creations and deletions in one customization context are visible to other customization contexts.

Solution

Artifact creations and deletions are not customization context sensitive. Therefore, only existing artifact customizations are supported in layers. Because new artifacts cannot be added as part of a layer, the artifacts are added to the project like any other base artifact. Any component artifacts that are added to the project are not incorporated into the runtime composite model unless the composite.xml file has explicitly referred to it. Therefore, if a customization did not alter composite.xml to include the newly added component, the new artifacts have no impact. Of course, the new artifacts are saved along with the base artifacts.

Additions of any new artifact introduced in the composite definition as a customization (for example, XSLT artifacts) must be named appropriately to prevent collisions with base XSLT artifacts. References to the new XSLT files must be made within the Oracle BPEL Process Manager/Oracle Mediator artifacts (through a layer).

21.9.3 Warnings Regarding xml:id when Building or Deploying a Customized Composite

Problem

When building or deploying the customized composite, you may get warnings similar to this because the xml:id element is not supported:

```
[scac] [WARNING] Line [13] Column [29] Schema validation failed for
 C:\customizationdemo\project1\composite.xml
<Line 13, Column 29>: XML-24535: (Error) Attribute
 'http://www.w3.org/XML/1998/namespace:id' not expected.
```

Solution

This warning can be ignored and should not prevent you from deploying the composite successfully.

21.9.4 Distributed Queue Topic Messages Are Retrieved Multiple Times by Subscribers

Problem

In a clustered environment, a queue topic is being retrieved multiple times by the subscribers. SOA clusters are typically homogenous (that is, each node has the same services running). Therefore, if a queue topic has multiple subscribers (for example, an error queue is subscribed to by a fallout notification service and a trouble ticketing service), then in a clustered environment, each service is running on every node. If it is a four-node cluster, then there are four instances of the fallout notification service and four instances of the trouble ticketing service, each of which retrieve the topic message. This is not desirable.

Solution

Set the consuming service singleton property to ensure that in a multinode environment, only one of the subscriber instances can consume the topic message. In

the example in the "Problem" section, the fallout notification service and the trouble ticketing service each should set their singleton property so that each of them consumes the message only once by any one of the four instances.

```
<service name="ConsumeFaultMessage" ui:wsdlLocation="ConsumeFaultMessage.wsdl">
<interface.wsdl
interface="http://xmlns.oracle.com/pcbpel/adapter/jms/EH/AIAReadJMSNotificationPro
cess/ConsumeFaultMessage/#wsdl.interface(Consume_Message_ptt)"/>
<binding.jca config="ConsumeFaultMessage_jms.jca">
  <property name="singleton">true</property>
</binding.jca>
</service>
```

21.9.5 Internal Key is Displayed Instead of a Translated String

Problem

An internal key is displayed instead of a translated string in the human task title, human task outcome, email subject and body, and human task mapped attributes displayed in Oracle BPM Worklist.

Solution

To resolve this problem:

1. Check the SOA `server`-diagnostic.log for errors for each server in the cluster. If a resource bundle or the key is not found, you see a similar error stack that contains the composite name and resource bundle name.

```
2010-09-29T23:45:46.702-07:00 soa_server1 ERROR \[\]
 oracle.soa.services.workflow.soa tid: orabpel.invoke.pool-4.thread-7 userId:
 weblogic
ecid:59ab739a26595dc1:-38d33932:12b614ddb44,0:1:100000026
APP: soa-infra composite_name: RC3HumanTaskComposite component_name: Humantask1
 component_instance_id: composite_instance_id: 310003
<WorkflowServiceEngine.request> Operation 'initiateTask' failed with exception
    'EJB Exception: : java.util.MissingResourceException:
Can't find bundleforbase name
   oracle.apps.hcm.people.soa.resource.HcmPeopleTopSoaBundleDummy, locale
 en\[\[
at
java.util.ResourceBundle.throwMissingResourceException(ResourceBundle.java:1427
)
at java.util.ResourceBundle.getBundleImpl(ResourceBundle.java:1250)
at java.util.ResourceBundle.getBundle(ResourceBundle.java:952)
```

2. Confirm that the resource bundle has been deployed to the correct location in SOA MDS by first exporting the contents of SOA MDS and confirming that the resource bundle class is defined under /apps/resource. Follow these steps to export the metadata contents from Fusion Applications Control:

   a. From the navigation pane, expand **SOA** and select **soa-infra**.

   b. From the **SOA Infrastructure** menu, choose **Administration** > **MDS Configuration**.

   c. Select **Export metadata documents to an archive on the machine where the web browser is running**.

   d. Click **Export**.

3. If the resource bundle is not defined in SOA MDS or is in the wrong location in SOA MDS, then you may have encountered an error during patching or there is a problem with the resource bundle definition.

If the resource bundle is properly defined in SOA MDS and the string that is not translated is a mapped attribute that appears in the column heading of the worklist, then perform the followings steps to:

- Set the System MBean Browser **WorkflowCustomClasspathURL** property to a value of **""**.

- Apply the changes.

- Set WorkflowCustomClasspathURL to `"oramds:///apps/resource/"`.

- Apply the changes in Fusion Applications Control.

a. Log in to Fusion Applications Control in the domain where the JAR file was deployed.

b. From the navigation pane, expand **SOA** and select **soa-infra**.

c. From the **SOA Infrastructure** menu, choose **SOA Administration** > **System MBean Browser**.

d. Go to **Application Defined MBeans** > **oracle.as.soainfra.config** > **Server:** *SOA_cluster_name* > **WorkflowConfig** > **human-workflow**.

e. Remove the value in the **Value** column for the **WorkflowCustomClasspathURL** attribute.

f. Click **Apply**.

g. Enter `oramds:///apps/resource/` in the **Value** column of the **WorkflowCustomClasspathURL** attribute.

h. Click **Apply**.

If the resource bundle is properly defined in SOA MDS and the string that is not translated is not a mapped attribute that appears in the column heading of the worklist, then stop and restart the composite that references the resource bundle by following these steps:

a. Log in to Fusion Applications Control in the domain where the JAR file was deployed.

b. From the navigation pane, expand **SOA** and select **soa-infra**.

c. Click **Shut Down**.

d. Click **Yes** in the confirmation window.

e. Click **Start Up**.

f. Click **Yes** in the confirmation window.

4. Retest your transaction flow.

21.9.6 BPEL Activity Errors and the Log Shows Error "com.oracle.bpel.entity.dataprovider.EntityVarMgrException:zero Item"

Problem

If entity variables are used in Oracle BPEL Process Manager, this error may be thrown at runtime:

```
com.oracle.bpel.entity.dataprovider.EntityVarMgrException:zero item
```

Solution

When the entity variable tries to bind the key provided to it at runtime in Oracle BPEL Process Manager and if it cannot find the key in the database, it throws this error. This may be due to human error. Either the developer has passed a wrong primary key value to bind to the entity variable or the composite has been deployed to the wrong business component (BC) URL in the deployment plan. To resolve, ensure that the primary key value that gets bound to the entity variable at runtime exists in the database used by the BC service.

21.9.7 BPEL Activity Errors and Log Shows Error "XPath variable or expression <expression> is empty"

Problem

If the payload node used in an activity such as an assign activity in BPEL is empty/null, then it throws this error.

```
The XPath variable or expression
/ns4:findDtsPurchasingDeliverable1FindByDelStatusResponse/ns4:resultns2:Deliverabl
eId=$DeliverableIdVar/ns2:BatchId is empty at line 344. selectionFailure:
 com.oracle.bpel.client.BPELFault: faultName:
 {{http://schemas.xmlsoap.org/ws/2003/03/business-process/}selectionFailure}
```

Solution

This error is thrown when a selection operation performed either in a function or in an assignment encounters an error. If there is a possibility of getting a null/empty value for a payload element at runtime, then you can get around this error by enclosing this element in a switch case with a condition check of (`string-length(payload_ element) > 0`) or a `bpelx:assert` element to test for null conditions in their code. If these errors are encountered at runtime, there is information in the log and the BPEL audit trail flow.

21.9.8 Unavailable Composite Errors Occurring Even Though the Target Service Is Up and Running

Problem

When a BPEL process has partner links targeted to external servers, during start up of the SOA server or the activation of the BPEL composite, the BPEL service engine recognizes that the target server is down and marks the BPEL composite as unavailable. Even after the target server becomes available, there is no way to make the BPEL process available without restarting the complete SOA server or redeploying the BPEL process. Until then, the BPEL composite is not listed in the navigation pane of Fusion Applications Control and is only visible through the **Deployed Composites** tab. Even though the composite is displayed here (sometimes as active), if you click the process, you receive this error:

```
The composite <Process Name>[ <Version> ] is not available. This could happen
because either the composite has been undeployed or EM has not yet discovered
this composite.
```

Solution

The issue is caused by using a concrete WSDL file in the BPEL definition. Using an abstract WSDL prevents the BPEL composite/process from becoming inactive or its state from becoming retired when the target server or service is down.

Oracle Fusion Applications develop composites with `oramds://` references to WSDL files at design time for both the binding and the WSDL location. At deployment, the binding's concrete location (which is only used at runtime) is fixed.

Follow these steps to use an abstract WSDL by making a local copy of the target WSDL file in your BPEL project.

1. When you create a web service partner link in a composite in the Create Web Service dialog, select the **copy WSDL and its dependent artifacts into the project** checkbox.

2. Once you are done with web service creation and the partner link is wired (connected) with the BPEL process service component, you see two WSDL files in the Application Navigator of Oracle JDeveloper:

   - *Webservice*.**WSDL** (the local copy of your target WSDL)

   - *Webservice*Ref.**WSDL**

3. Redeploy the SOA composite.

4. Retest the scenario.

21.9.9 High-Volume Cross Reference (XREF_DATA) Table Impacts Performance and Maintenance

Problem

When the volume of data in the XREF Oracle SOA schemas grows very large, maintaining the database can become difficult. It may have an impact on performance and may be identified in Automatic Workload Repository (AWR) snapshots.

Solution

To address this maintenance challenge, Oracle SOA 11*g* has been instrumented with partition keys that enable database administrators (DBAs) to take advantage of the Oracle RDBMS partitioning features and capabilities. Even thought multiple types of partitioning are possible, the list partitioning strategy is the preferred one for partitioning the cross reference data table.

List partitioning allows greater flexibility in the mapping of rows to partitions than range or hash partitioning. Because the cross reference data (`XREF_DATA`) is distributed based on discrete column values, unordered and unrelated sets of data can be grouped together with no relationship between the partitions. For this table, the list partitioning is based on the column `XREF_TABLE_NAME`. Because you know the list of entities that have cross reference data, this strategy is ideal. The only caveat is that you must repeat this activity if a new entity gets added to the system.

Better understanding of data distribution can lead to further refinement and space management by employment of composite partitioning strategies such as list - list and list - range schemes. For information about partitioning, see the *Oracle Database Performance Tuning Guide*.

In addition to this partitioning strategy, you can also use Oracle Mediator's feature by which custom database tables can be created to store cross reference data for certain

high volume entities. For more information, see *Oracle Fusion Middleware Developer's Guide for Oracle SOA Suite*.

21.9.10 Access Denied Error While Invoking an Oracle ADF BC Service from BPEL

Problem

A common case in Oracle Fusion Applications is Oracle Application Development Framework (Oracle ADF) Business Component (BC) service invocations from various SOA components. If the user context attached with the SOA composite instance or identity switching used in a SOA reference does not have proper permission to access the Oracle ADF BC service, then an `access denied` error is displayed.

Solution

Common causes for this error are as follows. These must be resolved.

- The Oracle ADF BC service does not have required permissions in `jazn-data.xml`.

- The permissions provided in `jazn-data.xml` are not properly reflected in `system-jazn-data.xml`.

- The `ejb-jar.xml` does not have the `<interceptors>` tag added or it has an incorrect application name. An example of `ejb-jar.xml` with the correct tags is as follows:

```
<ejb-jar xsi:schemaLocation="http://java.sun.com/xml/ns/javaee
 [http://java.sun.com/xml/ns/j2ee/ejb-jar_3_0.xsd]"version="3.0">
<enterprise-beans>
<session>
<ejb-name>oracle.apps.financials.payables.invoices.transactions.invoiceApproval
Service.InvoiceApprovalBCServiceBean</ejb-name>
<resource-ref><res-ref-name>jdbc/ApplicationDBDS</res-ref-name><res-type>javax.
sql.DataSource</res-type><res-auth>Container</res-auth></resource-ref>
<resource-ref><res-ref-name>jdbc/ApplicationServiceDBDS</res-ref-name><res-type
>javax.sql.DataSource</res-type><res-auth>Container</res-auth></resource-ref>
</session>
</enterprise-beans>
 
<interceptors> <interceptor>
<interceptor-class>oracle.security.jps.ee.ejb.JpsInterceptor</interceptor-class
>
<env-entry> <env-entry-name>application.name</env-entry-name>
 <env-entry-type>java.lang.String</env-entry-type>
 <env-entry-value>fscm</env-entry-value>
<injection-target>
<injection-target-class>oracle.security.jps.ee.ejb.JpsInterceptor</injection-ta
rget-class>
<injection-target-name>application_name</injection-target-name>
</injection-target>
</env-entry>
</interceptor> </interceptors>
 
<assembly-descriptor>
<interceptor-binding>
<ejb-name>*</ejb-name><interceptor-class>oracle.security.jps.ee.ejb.JpsIntercep
tor</interceptor-class> </interceptor-binding>
</assembly-descriptor>
 
</ejb-jar>
```

21.9.11 Clicking Manage Approvals Page Gives a "SOA server may be down" or "No data to display" Error

Problem

At runtime, when you click the Manage Approvals page (for both requisition and purchase order), you may get a `SOA server may be down` error among other errors. The page may also show a `No data to display` error.

Solution

Check to see if you have `wf_client_config.xml` pointing to the correct SOA server. Oracle Fusion Applications environments are all multidomain and applications that must query their local SOA runtime must do so through the `wf_client_config.xml` configuration.

For more information about the wf_client_config.xml file, see Section 21.5.1.

21.9.12 Synchronous Service Invocation Errors Due to WS-Addressing

Problem

When integrating with external web services, you may have issues related to the WS-Addressing headers in the request and response XML documents used in synchronous request-response scenarios. This may even occur if there is a dynamic endpoint that resolves to an external web service that does not support WS-Addressing.

- A SOA composite acting as a client fails to invoke a service provider that does not support WS-Addressing headers in the request message for synchronous interactions.

- An external service consumer that does not support WS-Addressing headers in the response message fails to invoke a synchronous operation of a SOA composite.

- A service provider throws a `wsa:InvalidAddressingHeader` error when `SOAPAction` is defined in the WSDL `<binding>` element of the service provider interface.

Solution

The WS binding component of Oracle Fusion Middleware 11*g* generates WS-Addressing headers by default, which is as per the specification, but may be incompatible with some web services. This behavior can be overwritten using properties that allow you to control the generation of WS-Addressing header information for interoperability purposes.

- A SOA composite acting as a client fails to invoke a service provider that does not support WS-Addressing headers in the request message for synchronous interactions.

  When your Oracle SOA Suite components act as service clients (consumers) and invoke synchronous operations of external web services, the invocation fails. The invocation failure is due to the WS-Addressing headers included in the request message when invoking the service provider.

  The stack trace displays the following information:

```
NSGetCustomer (faulted)
Nov 10, 2009 4:26:40 PM Faulted while invoking operation "get" on provider
"NetSuite2009_PL".
```

```
<payload>
Nov 10, 2009 4:26:40 PM
"{urn:platform_2009_1.webservices.netsuite.com}UnexpectedErrorFault" has been
thrown.
<payload>
Nov 10, 2009 4:26:40 PM There is a system exception while performing the BPEL
instance, the reason is "Invalid SOAP Header: '<ns1:To
xmlns:ns1="http://www.w3.org/2005/08/addressing">https://webservices.netsuite.c
om/services/NetSuitePort_2009_1</ns1:To>'. Value is
'https://webservices.netsuite.com/services/NetSuitePort_2009_1'."
<payload>
```

In Oracle SOA Suite 11*g*, the WS-Binding component includes the following
WS-Addressing headers:

– `wsa:to`

– `wsa:replyTo`

– `wsa:action`

These headers are included by default regardless of interacting with synchronous
or asynchronous web service operations. Although the SOAP message generated
by 11*g* is fully specification compliant, there are service provider implementations
that do not understand these WS-Addressing headers when included in the
request message for synchronous operations. This is why the invocation fails.

Oracle Fusion Middleware 11*g* uses new standards, WS-Addressing (`wsa`) and
WS-I Basic profile, which demand that WS-Addressing headers be added by
default to the request message.

To increase interoperability, the property `oracle.soa.ws.outbound.omitWSA`
has been introduced for the `<binding.ws>` element in the `composite.xml` file.
This property enables you to optionally suppress generation of WS-Addressing
headers in the request message when set to `true`. The default value is `false`.

1. Manually add the following property to the `<binding.ws>` element in the
   `composite.xml` file.

   ```
   [snippet]
   <binding.ws
   ......
       <property name="oracle.soa.ws.outbound.omitWSA" type="xs:boolean"
   many="false" override="may">true</property>
   </binding.ws>
   [/snippet]
   ```

2. Redeploy the composite.

■ An external service consumer that does not support WS-Addressing headers in the
  response message fails to invoke a synchronous operation of a SOA composite.

  When external services invoke the endpoints of Oracle SOA Suite composites,
  which expose synchronous operation(s) acting as a server (service provider), the
  external web service fails with invocation errors. This is commonly due to the
  WS-Addressing headers included in the response message, which are not
  supported by some web service implementations.

  In SOA Suite 11*g*, the WS-Binding component includes the WS-Addressing
  headers by default in the response messages. Although the SOAP message
  generated by 11*g* is fully specification compliant, there are service consumer
  implementations that do not understand the WS-Addressing headers when

included in the response message in synchronous request-response scenarios. This is why the invocation fails.

To increase interoperability with clients (consumers) that do not support WS-Addressing headers, a property named `oracle.soa.addressing.response.enabled` has been introduced. This property can be used within the `<binding.ws>` element in the `composite.xml` file. This property enables you to optionally suppress generation of WS-Addressing headers in the response message when set to `true`. The default value is `false`.

This means that in a synchronous request-response scenario, in which the Oracle SOA Suite composite acts as a service provider, WS-Addressing headers are included in the response message by default. If the property is set to `false`, no WS-Addressing-related headers are returned in the response message.

1. Set the property in the `composite.xml` file as follows:

```
[snippet]
<binding.ws>
......
<property name="oracle.soa.addressing.response.enabled" type="xs:boolean"
many="false" override="may">false</property>
</binding.ws>
[/snippet]
```

■ A service provider throws a `wsa:InvalidAddressingHeader` error when `SOAPAction` is defined in the WSDL `<binding>` element of the service provider interface.

When invoking a web service provider with a synchronous interface from a composite deployed on 11.1.1.1, a `wsa:InvalidAddressingHeader` message is returned. This happens when both of the following conditions are met:

– The service interface (WSDL) of the service provider does not include the `wsa:Action` element in the WSDL input message for the port operations

and

– The `soapAction` attribute in the WSDL binding section of the service provider interface is specified.

The difference with the first scenario is that the service provider interface may accept WS-Addressing headers for synchronous endpoint operations by using the WS-Addressing extensibility attribute `wsa:UsingAddressing` in the WSDL binding section. A WSDL binding section sample is shown as follows:

```
<wsdl:binding name="mySoapBinding"
type="impl:myService">
  <wsaw:UsingAddressing wsdl:required="false"
 xmlns:wsaw="http://www.w3.org/2006/05/addressing/wsdl"/>
  <wsdlsoap:binding style="document"
 transport="http://schemas.xmlsoap.org/soap/http"/>
  <wsdl:operation name="myOperation">
    <wsdlsoap:operation soapAction="myAction"/>
    <wsdl:input name="myRequestMsg">
      <wsdlsoap:body use="literal"/>
    </wsdl:input>
    <wsdl:output name="myResponseMsg">
      <wsdlsoap:body use="literal"/>
    </wsdl:output>
  </wsdl:operation>
```

```
</wsdl:binding>
```

As the following code sample shows, neither the input message nor the output message (as defined in the service provider WSDL) contain a `<wsa:Action>` element:

```
....
<wsdl:message name="myRequestMsg">
  <wsdl:part name="parameters"
element="impl:myRequestElement"/>
</wsdl:message>
.....
<wsdl:message name="myresponseMsg">
  <wsdl:part name="parameters"
element="impl:myResponseElement"/>
</wsdl:message>
```

The problem has been observed when integrating with synchronous web services deployed on the IBM WebSphere 6.1 Application Server.

The stack trace looks similar to the following code example:

```
......
oracle.fabric.common.FabricInvocationException:
javax.xml.ws.soap.SOAPFaultException: A header representing a Message
Addressing Property is not valid and the message cannot be processed
at
oracle.integration.platform.blocks.soap.WebServiceExternalBindingComponent.thro
wFabricInvocationException(WebServiceExternalBindingComponent.java:415)
at
oracle.integration.platform.blocks.soap.WebServiceExternalBindingComponent.thro
wFabricInvocationExceptionForSoapFault(WebServiceExternalBindingComponent.java:
411)
at
oracle.integration.platform.blocks.soap.WebServiceExternalBindingComponent.proc
essSOAPFault(WebServiceExternalBindingComponent.java:394)
at
oracle.integration.platform.blocks.soap.WebServiceExternalBindingComponent.proc
essOutboundMessage(WebServiceExternalBindingComponent.java:253)
at
oracle.integration.platform.blocks.soap.WebServiceExternalBindingComponent.send
SOAPMessage(WebServiceExternalBindingComponent.java:644)
at
oracle.integration.platform.blocks.soap.WebServiceExternalBindingComponent.requ
est(WebServiceExternalBindingComponent.java:526)
at
oracle.integration.platform.blocks.mesh.SynchronousMessageHandler.doRequest(Syn
chronousMessageHandler.java:139)
at
oracle.integration.platform.blocks.mesh.MessageRouter.request(MessageRouter.jav
a:179)
at
oracle.integration.platform.blocks.mesh.MeshImpl.request(MeshImpl.java:144)
at sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)
at
......
```

If you use HTTP Analyzer or Apache TCP Monitor to trace the message exchange, you observe that the issue is related to the `<wsa:action>` element in the request SOAP header, which is not the same as `SOAPAction` in the HTTP header. The response message looks like the following code sample:

```
<soapenv:Envelope>
.......
  <soapenv:Header>
    <wsa:FaultDetail xmlns:wsa="http://www.w3.org/2005/08/addressing">
        <wsa:ProblemHeaderQName>wsa:Action</wsa:ProblemHeaderQName>
    </wsa:FaultDetail>
  </soapenv:Header>
  <soapenv:Body>
    <soapenv:Fault>
      <faultcode>wsa:InvalidAddressingHeader</faultcode>
      <faultstring>
        <![CDATA[A header representing a Message Addressing Property is not
        valid and the message cannot be processed]]></faultstring>
    </soapenv:Fault>
  </soapenv:Body>
</soapenv:Envelope>
```

As mentioned in the previous sections, the WS Binding component of Oracle Fusion Middleware 11*g* generates WS-Addressing headers by default. According to W3C specification (Web Services Addressing 1.0 - WSDL Binding W3C Candidate Recommendation 29 May 2006) Section 4.4.1 Explicit Association:

"In the absence of a `wsa:Action` attribute on a WSDL input element where a `SOAPAction` value is specified, the value of the [action] property for the input message is the value of the `SOAPAction` specified."

The generated value for the `<wsa:Action>` element of the input message matches the value of the `SOAPAction` element, if that is specified in the WSDL binding section. Otherwise, it generates a default value for the `<wsa:Action>` element.

A sample of the generated request is as follows (provided WSDL `<binding>` and `<messages>` section as shown previously).

```
POST /currencyratesapi/services/myAction HTTP/1.1
Host: <host>:<port>
Connection: TE
TE: trailers, deflate, gzip, compress
User-Agent: Oracle HTTPClient Version 10h
SOAPAction: "myAction"
Accept-Encoding: gzip, x-gzip, compress, x-compress
ECID-Context: 1.0000IUhXOGw3R9BpNWK6ye1BPi6A0002Wk;kYhgp8TLo0hgv0
Content-type: text/xml; charset=UTF-8
Content-Length: 1473

<env:Envelope xmlns:env="http://schemas.xmlsoap.org/soap/envelope/"
xmlns:wsa="http://www.w3.org/2005/08/addressing">
<env:Header>

<wsa:To>http://<host>:<port>/currencyratesapi/services/myAction</wsa:To>
<wsa:Action>myAction</wsa:Action>
<wsa:MessageID>...</wsa:MessageID>
<wsa:RelatesTo>...</wsa:RelatesTo>
<wsa:ReplyTo>
<wsa:Address>http://www.w3.org/2005/08/addressing/anonymous</wsa:Address>
<wsa:ReferenceParameters>
......
</wsa:ReferenceParameters>
</wsa:ReplyTo>
</env:Header>
<env:Body>
```

```
......
</env:Body>
</env:Envelope>
```

21.9.13 Deploying Human Workflow Application Throws "Unable to resolve 'TaskQueryService'" Error

Problem

Almost all Oracle Fusion Applications have the `hwtaskflow.xml` file with the `IntegrateTaskFlowWithTask` entry in the `web.xml` file to register the notification projects. When deploying the J2EE EAR file, it throws the following error if you did not configure the foreign JNDI connection or `wf_client_config.xml` with a valid SOA environment.

```
javax.naming.NameNotFoundException:Unable to resolve 'TaskQueryService'. Resolved
''; remaining name 'TaskQueryService'
```

Solution

If the SOA server is not up and running, then start the SOA server using Fusion Applications Control. If it is up and the problem exists, then check that the `wf_client_config.xml` file has a valid SOA environment. If possible, avoid foreign JNDI references because the JNDI resources get bound into all servers in the domain that causes `BindException` errors for servers in which the actual service exists when source and target JNDI names are the same. Oracle Fusion Applications do not use foreign JNDI references, instead bundling a correctly configured `wf_client_config.xml` file in the EAR file or using APIs at runtime to generate the Java Architecture for XML Binding (JAXB) for the human workflow APIs.

Use `wf_client_config.xml` (bundled at the application level) for all ADF worklist client applications communicating with Oracle SOA Suite. Make sure to specify the correct URL in the file configuration.

```
<?xml version="1.0" encoding="UTF-8" ?>

<workflowServicesClientConfiguration
 xmlns="http://xmlns.oracle.com/bpel/services/client">
 <server name="default" default="true" clientType="REMOTE">
 <localClient>
  <participateInClientTransaction>false</participateInClientTransaction>
  </localClient>
  <remoteClient>
  <serverURL>t3:cluster://CRMSoaCluster</serverURL>
<initialContextFactory>weblogic.jndi.WLInitialContextFactory</initialContextFa
ctory>
  <participateInClientTransaction>false</participateInClientTransaction>
  </remoteClient>
 <soapClient>
<rootEndPointURL>http://adc60048fems.us.oracle.com:6361/soa-infra</rootEndPointURL
>
  </soapClient>
  </server>
</workflowServicesClientConfiguration>
```

21.9.14 "Invalid Subject" Error Thrown with Human Workflow API or Notification

Problem

When Oracle SOA Suite is installed on `Domain1` and ADF on `Domain2` and if any remote calls are performed between them for any human workflow API use or for notifications, errors similar to the following are thrown if the domain trust is not set.

```
javax.servlet.ServletException:
javax.servlet.ServletException:java.lang.SecurityException: Security:090398Invalid
Subject:principals=CVBUYER01, AR_MANAGER_VISION_OPERATIONS_DATA, AR_ACCOUNTS_
RECEIVABLE_MANAGER_JOB, FBI_TRANSACTIONAL_BI_WORKER
```

Solution

Set the domain trust password on both domains (global trust).

1. Log in to Oracle WebLogic Server Administration Console.

2. In the **Domain Structure**, click the domain name.

3. Click the **Security** tab, and then click **Advanced**.

4. Set the password in the **Credential** and **Confirm Credential** fields.

Note: Oracle Fusion Applications do not support RMI across domains. That is, all remote calls between ADF and SOA are always within the same domain. When cross domain communication is required, SOAP is used.

You may encounter the scenario described in this section in your development environment. For example, you may be running the ADF task flow for human tasks in the JDeveloper-integrated Oracle WebLogic Server and running the composite in an independent standalone SOA server. Both are in different domains.

21.9.15 Task is Assigned to the Group/Role when It Is Expected to Go to Every User in the Group/Role Individually

For troubleshooting information about this issue, see the "Task Assignment/Routing/Escalation Issues" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

21.9.16 Task Completes Without Any Assignment Occurring

For troubleshooting information about this issue, see the "Task Assignment/Routing/Escalation Issues" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

21.9.17 Parallel Assignees Must Approve or Reject the Task Even Though the Parallel Completion Criteria is Met

For troubleshooting information about this issue, see the "Task Assignment/Routing/Escalation Issues" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite*.

21.9.18 All Added Adhoc Participants Disappear After a Page Refresh

For troubleshooting information about this issue, see the "Task History Issues" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite.*

21.9.19 Future Approvers Are Not Visible in the History Table

For troubleshooting information about this issue, see the "Task History Issues" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite.*

21.9.20 Message Appears in the History Table about a Correlation ID Not Being Passed or Any Exception Related to the Correlation ID

For troubleshooting information about this issue, see the "Task History Issues" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite.*

21.9.21 Edit Toolbar Is Disabled or Not Shown

For troubleshooting information about this issue, see the "Task History Issues" section of *Oracle Fusion Middleware Administrator's Guide for Oracle SOA Suite and Oracle Business Process Management Suite.*

21.9.22 oracle.jrf.UnknownPlatformException Error When Customizing an Oracle Mediator

Problem

You can receive the following error when making customization modifications to an Oracle Mediator service component in a SOA application.

For example, assume you perform the following steps:

1. Log into JDeveloper with the **Fusion Application Developer** role.

2. Create a SOA application and project with the **Customizable** option selected.

3. Add appropriate JAR files under the **Libraries and Classpath** section of the Project Properties dialog for the project.

4. Under **Application Resources** > **Descriptors** > **ADF META-INF** > **adf-config.xml,** add appropriate customization classes.

5. Add an Oracle Mediator service component to the application.

6. Save the application and restart JDeveloper.

7. Add appropriate routing rule and XSL transformations to the Oracle Mediator service component.

    The following error is displayed in the command prompt window:

    ```
    oracle.jrf.UnknownPlatformException: JRF is unable to determine the current
    application server platform.
    ```

Solution

This error can be ignored. After this error occurs, everything still works correctly.

21.9.23 java.lang.NullPointerException When Customizing an Oracle Mediator

Problem

You can receive the following error when making a customization modification to an Oracle Mediator service component in a SOA application.

For example, assume you perform the following steps:

1. Log into JDeveloper with the **Fusion Application Developer** role.

2. Create a SOA application and project with the **Customizable** option selected.

3. Add appropriate JAR files under the **Libraries and Classpath** section of the Project Properties dialog for the project.

4. Under **Application Resources** > **Descriptors** > **ADF META-INF** > **adf-config.xml,** add customization classes.

5. Add an Oracle Mediator service component to the application.

6. Add a routing rule to the Oracle Mediator service component.

    The following error is displayed in the command prompt window:

    ```
    java.lang.NullPointerException
    ```

Solution

As long you do not see any alerts indicating that a WSDL file cannot be found or parsed, then this exception in the log/console can be ignored.

21.9.24 JDeveloper Compilation Error in a SOA Project with SOA MDS Service Location

Problem

1. Start JDeveloper with the **Fusion Application Developer** role.

2. Create a SOA application.

3. Create and deploy a SOA bundle to the server.

4. Create another SOA application.

5. Create an Oracle Mediator service component.

6. Create a reference web service and access the SOA MDS connection and browser for the service of the SOA application that you created in Step 2.

7. Wire the components.

8. Update the empty `port` and `location` values in the `composite.xml` file. For the `location` value, enter the `oramds` path. For example:

    ```
    <reference name="Service1"

    ui:wsdlLocation="oramds:/apps/CalleeProject/CalleeBPELProcess.wsdl">
        <interface.wsdl
    interface="http://xmlns.oracle.com/Comp2Comp_
    EDGApp/CalleeProject/CalleeBPELProcess#wsdl.interface(CalleeBPELProcess)"/>
        <binding.ws
    port="http://xmlns.oracle.com/Comp2Comp_
    EDGApp/CalleeProject/CalleeBPELProcess#wsdl.endpoint(calleebpelprocess_client_
    ep/CalleeBPELProcess_pt)"
        location="oramds:/apps/CalleeProject/CalleeBPELProcess.WSDL"/>
      </reference>
    ```

9. Compile the application.

You receive the following error message.

```
Warning(26,76): Failed to Find Binding
```

Solution

This is only a warning message thrown by the compiler. Runtime works fine. Having a reference from SOA MDS in a composite makes the `port` and `location` values empty in the bindings. To resolve this warning message:

- Explicitly enter the corresponding details. For example, replace the `binding.ws` `location` value with the concrete WSDL file.

- Deploy the composite using the configuration plan, which replaces the port and `oramds` location with the concrete WSDL file.

21.10 Using My Oracle Support for Additional Troubleshooting Information

You can use My Oracle Support (formerly MetaLink) to help resolve Oracle Fusion Applications problems. My Oracle Support contains several useful troubleshooting resources, such as:

- Knowledge base articles

- Community forums and discussions

- Patches and upgrades

- Certification information

You can access My Oracle Support at `https://support.oracle.com`.

22

Troubleshooting Oracle Identity Management

Use this chapter to troubleshoot runtime Oracle Fusion Applications problems that may have originated in the Oracle Identity Management and security integration layer.

This chapter contains the following topics:

- Introduction to Troubleshooting Oracle Identity Management
- Getting Started with Troubleshooting Oracle Identity Management
- Problems and Solutions
- Additional Information for Troubleshooting Oracle Identity Management
- Using My Oracle Support for Additional Troubleshooting Information

In addition to this chapter, review the *Oracle Fusion Middleware Error Messages Reference* for information about the error messages you may encounter.

22.1 Introduction to Troubleshooting Oracle Identity Management

This section provides guidelines and a process for using the information in this chapter. Using the following guidelines and process will focus and minimize the time you spend resolving problems.

Guidelines

When using the information in this chapter, Oracle recommends:

- Reviewing the various terms and implementations of roles within Oracle Fusion Applications and Oracle Identity Management. For example, a duty role in the context of Oracle Fusion Applications equates to an application role in the context of Oracle Identity Management. Refer to Section 4.5 and Table 4–2, " Equivalent Terminology" for complete information.

- After performing any of the solution procedures in this chapter, immediately retrying the failed task that led you to this troubleshooting information. If the task still fails when you retry it, perform a different solution procedure in this chapter and then try the failed task again. Repeat this process until you resolve the problem.

- Making notes about the solution procedures you perform, symptoms you see, and data you collect while troubleshooting. If you cannot resolve the problem using the information in this chapter and you must log a service request, the notes you make will expedite the process of solving the problem.

Process

Follow the process outlined in Table 22–1 when using the information in this chapter. If the information in a particular section does not resolve your problem, proceed to the next step in this process.

Table 22–1   Process for Using the Information in this Chapter

Step	Section to Use	Purpose
1	Section 22.2	Get started troubleshooting Oracle Identity Management. The procedures in this section quickly address a wide variety of problems.
2	Section 22.3	Perform problem-specific troubleshooting procedures. This section describes:
		■ Symptoms of specific Oracle Fusion Applications runtime problems that may have originated in the Oracle Identity Management and security integration layer
		■ Possible causes of the problems
		■ Solution procedures corresponding to each of the possible causes
3	Section 22.4	Get Oracle Identity Management component-specific troubleshooting information. Use this section if you have isolated your problem to a specific Oracle Identity Management component or want to learn more about a component.
4	Section 22.5	Use My Oracle Support to get additional troubleshooting information about Oracle Fusion Applications or Oracle Identity Management. My Oracle Support provides access to several useful troubleshooting resources, including Knowledge Base articles and Community Forums and Discussions.
5	Section 22.5	Log a service request if the information in this chapter and My Oracle Support does not resolve your problem. You can log a service request using My Oracle Support at `https://support.oracle.com`.

22.2 Getting Started with Troubleshooting Oracle Identity Management

Start troubleshooting by performing the procedures in this section, as they quickly address a wide variety of problems. If the procedures in this section do not resolve your problem, proceed to Section 22.3.

This section contains the following topics:

■ Verifying Oracle Internet Directory Identity Stores Can Perform Look Ups

■ Verifying the Security Providers in the Oracle WebLogic Server Domain

■ Using Selective Tracing to Troubleshoot Inaccessible Functionality

22.2.1 Verifying Oracle Internet Directory Identity Stores Can Perform Look Ups

When using Oracle Internet Directory as the identity store, it must be configured to index the `displayName` attribute. If Oracle Internet Directory is not configured to index the `displayName` attribute, operations that require looking up users and roles in the identity store will fail.

To verify an Oracle Internet Directory identity store is configured to index the `displayName` attribute:

1. Invoke Oracle Directory Services Manager and connect to the Oracle Internet Directory identity store instance. Refer to the "Invoking Oracle Directory Services Manager" and "Connecting to the Server from Oracle Directory Services Manager"

sections in the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory* for more information.

2. Click **Schema** on the Oracle Directory Services Manager task selection bar.

3. Expand the **Attributes** area of the navigation panel, enter displayName in the search field, and click the **Go** (**>**) button to search for the displayName attribute.

4. Click the **displayName** attribute in the search results. The configuration details for the displayName attribute appear in the main screen.

5. Verify the **Indexed** option is selected (checked) in the configuration details.

   If the **Indexed** option is not selected, click **the attribute will be cataloged/decataloged** button below the search field in the navigation tree.

   Refer to the "Adding an Index to an Existing Attribute by Using Oracle Directory Services Manager" section in the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory* for more information.

22.2.2 Verifying the Security Providers in the Oracle WebLogic Server Domain

Small configuration errors in the security providers for the Oracle WebLogic Server domain, such as in the Identity Asserters and Authenticators, frequently are the cause of runtime problems. Use the information in this section to quickly verify a few key security provider settings, including:

- The order of providers, which determines the authentication sequence.
- JAAS Control Flags, which determine how the authentication sequence uses the providers.
- Connection, cache, and user and group lookup settings for the identity store's LDAP Authenticator.

To verify configuration settings for the security providers in the Oracle WebLogic Server domain:

1. Log in to the Oracle WebLogic Server Administration Console by referring to the "Starting the Administration Console" section in the *Oracle Fusion Middleware Introduction to Oracle WebLogic Server*.

2. Click **Security Realms** in the **Domain Structure** area on the left side of the Administration Console Home Page. The Summary of Security Realms screen appears.

3. Click the name of the appropriate security realm in the **Realms** table. The Settings for *REALM_NAME* screen appears.

4. Click the **Providers > Authentication** tabs. The configured providers appear in the **Authentication Providers** table.

Verifying the Order of Providers
The security providers must be configured in the following order, where number 1 in the following list is at the top of the **Authentication Providers** table:

1. Oracle Access Manager Identity Asserter

2. LDAP Authenticator for the identity store: Either the Oracle Internet Directory Authenticator or Oracle Virtual Directory Authenticator, depending on the LDAP server used as the identity store.

3. Default Authenticator

If needed, you can reorder the security providers by performing the following steps from the Settings for *REALM_NAME* screen:

1. Click **Reorder**.

2. Select a provider and use the arrow buttons to move it up or down in the order.

3. Click **OK**.

Verifying JAAS Control Flags

The JAAS Control Flags for the security providers must be set as shown in Table 22–2. Perform the following steps to view, and if needed, edit the JAAS Control Flags.

From the Settings for *REALM_NAME* screen:

1. Click the provider name in the **Authentication Providers** table.

2. Click the **Configuration > Common** tabs.

3. Examine the **Control Flag** setting and adjust it as needed.

4. Click **Save**.

Table 22–2   Required JAAS Control Flags for Security Providers

Security Provider	Required JAAS Control Flag
Oracle Access Manager Identity Asserter	Required
LDAP Authenticator for the identity store: ■ Oracle Internet Directory Authenticator ■ Oracle Virtual Directory Authenticator	Sufficient
Default Authenticator	Sufficient

Verifying Settings for the Identity Store's LDAP Authenticator

Table 22–3 lists settings for the identity store's LDAP Authenticator that you should verify. Perform the following steps on either the Oracle Internet Directory Authenticator or the Oracle Virtual Directory Authenticator, depending on the LDAP server you are using for the identity store.

From the Settings for *REALM_NAME* screen:

1. Click the appropriate authenticator in the **Authentication Providers** table.

2. Click the **Configuration > Provider Specific** tabs.

3. Examine the settings and adjust as needed.

4. Click **Save**.

> **Note:**   You can get more information about each of the settings listed in Table 22–3 by clicking **More Info...** next to each setting in the Oracle WebLogic Server Administration Console.

Table 22–3   Settings to Verify in the Identity Store's LDAP Authenticator

Setting	Verification to Perform
Connection settings	Double-check all to ensure accuracy. Pay particular attention to the **Host** value, which can contain misspelled strings.

Table 22–3 (Cont.) Settings to Verify in the Identity Store's LDAP Authenticator

Setting	Verification to Perform
User Name Attribute	Regardless of which attribute is set, the same attribute must be used to specify the user name in the **All Users Filter** and **User From Name Filter** settings.
All Users Filter and User From Name Filter	The user name attribute used in both of these settings must be the attribute configured for the **User Name Attribute** setting.
Use Retrieved User Name as Principal	Must be enabled (checked).
Static Group Name Attribute	Regardless of which attribute is set, the same attribute must be used to specify the group name in the **All Groups Filter** and **Group From Name Filter** settings.
All Groups Filter and Group From Name Filter	The attribute used to specify the group name in these two settings must be the same attribute configured for the **Static Group Name Attribute** setting.
Cache Enabled	If enabled, examine the value of the **Cache TTL** setting.
Cache TTL	Examine to ensure an appropriate value is set. If you perform an operation that fails, wait for the amount of time specified by the **Cache TTL** to elapse and then retry the failed operation. This will ensure the authenticator's cache has been refreshed and any recent configuration changes have been activated.

22.2.3 Using Selective Tracing to Troubleshoot Inaccessible Functionality

When Oracle Fusion Applications users cannot access a particular functionality, for example, they attempt to log in to an application and are denied access or see an unexpected view of the application, often it is because they are not authorized to access that functionality. In these situations, you can use Fusion Middleware Control's Selective Tracing feature to collect data specific to the user and request, then collaborate with the Security Administrator to compare it against the configured authorizations.

To use Selective Tracing to troubleshoot inaccessible functionality:

1. Update the domain's environment setup script by performing one of the following steps that is appropriate to your environment:

    On Linux/UNIX systems, add the text shown in Example 22–1 to the bottom of the *DOMAIN_HOME*/bin/setDomainEnv.sh file:

Example 22–1 Updating the Domain's Environment Setup Script on Linux/UNIX Systems
```
JAVA_OPTIONS="-Djava.util.logging.manager=oracle.core.ojdl.logging.ODLLogManager ${JAVA_OPTIONS}"
export JAVA_OPTIONS
FMWCONFIG_CLASSPATH=${FMWCONFIG_CLASSPATH}${CLASSPATHSEP}${ORACLE_COMMON_HOME}/modules/oracle.odl_
11.1.1/ojdl.jar
export FMWCONFIG_CLASSPATH
```

On Windows systems, add the text shown in Example 22–2 to the bottom of the *DOMAIN_HOME*\bin\setDomainEnv.cmd file:

Example 22–2 Updating the Domain's Environment Setup Script on Windows Systems
```
set JAVA_OPTIONS=-Djava.util.logging.manager=oracle.core.ojdl.logging.ODLLogManager %JAVA_OPTIONS%
```

```
set FMWCONFIG_CLASSPATH=%FMWCONFIG_CLASSPATH%;%ORACLE_COMMON_HOME%\modules\oracle.odl_
11.1.1\ojdl.jar
```

2. Log in to Fusion Middleware Control by referring to the "Displaying Fusion Middleware Control" section in the *Oracle Fusion Middleware Administrator's Guide* for more information.

3. Navigate to the appropriate domain, then select **Logs > Selective Tracing** from the domain menu. The Selective Tracing page appears.

4. Click the **Tracing Options** tab, configure the following settings, and click **Start Tracing** to generate the selective trace:

   ■ Option Name: Select **User Name** from the list and enter the name of the user that cannot access functionality.

   > **Note:** While this procedure describes troubleshooting inaccessible functionality by selective tracing on a user name, you can also use the other options in the Option Names list for troubleshooting purposes.

   ■ Level: Select **TRACE:32 (FINEST)**.

   ■ Description: Enter a description that will help you identify the trace results, such as: *USER_NAME* cannot access functionality.

   ■ Duration: Enter the number of minutes the selective trace will run.

   ■ Trace ID: Select **Generate a New Unique Trace ID**. Optionally, you can select **Use a Custom Trace ID** and enter an ID, but note that Fusion Middleware Control does not verify the uniqueness of Custom Trace ID strings.

   ■ Loggers: Oracle recommends enabling the following loggers for troubleshooting inaccessible functionality:

   > **Note:** To quickly locate a specific logger, enter the logger name or a string in the logger name in the field above the list of loggers and press return.

   – oracle.jps.authorization

   – oracle.jps.common

   – oracle.security.jps.az.internal.runtime.policy.AbstractPolicyImpl

   – oracle.security.jps.internal.policystore.JavaPolicyProvider

   – oracle.security.jps.internal.policystore.ldap.BulkAuthorizer

   – oracle.security.jps.trace.logger

   – oracle.security.jps.util.JpsAuth

   > **Note:** Refer to the "Debugging the Authorization Process" section in the *Oracle Fusion Middleware Application Security Guide* for information about system properties you can enable for extremely fine grained authorization debugging.

5. Instruct the user that cannot access functionality to try and access it again. Now that you have enabled Selective Tracing for that user, you will collect data specific to that user and the request.

6. Access the results from the selective trace by clicking the **Active Traces And Tracing History** tab and selecting the trace from either the **Active Traces** or **Tracing History** table. If the number of minutes that you specified in the Duration option has elapsed, the trace will be in the **Tracing History** table. If you provided a description for the selective trace, look for it in the **Description** column.

7. Provide the trace results to the Security Administrator.

> **Note:** Typically, the Security Administrator performs the remaining steps in this procedure.

8. Locate the `Failed ProtectionDomain` string and its corresponding `resourceName=`, `resourceType=`, and `Principal=` strings in the trace results. These strings will provide information about the user and the inaccessible resource. As shown in Example 22–3, the user named `user1` was denied access to the resource named `ResourceNameX`:

Example 22–3   Failed ProtectionDomain String in Sample Selective Tracing Results

```
PolicyContext: [JeeScenarioApp]
Resource/Target: [resourceType=TaskFlowResourceType,resourceName=ResourceNameX]
Action:[read]
Permission Class: [oracle.security.jps.ResourcePermission]
          Result:          [FAILED]
          Evaluator:       [ACC]
          FailedProtectionDomain:ClassLoader=weblogic.utils.classloaders.ChangeAwareClassLoader
@c7cee9finder:weblogic.utils.classloaders.CodeGenClassFinder@a05da2 annotation:
JeeScenarioApp@jeescenario
                         CodeSource=file:/somepath/wls-jrfServer/servers/jrfServer_admin/tmp/
_WL_user/JeeScenarioApp/gw8m4w/war/WEB-INF/lib/_wl_cls_gen.jar
                         Principals=total 5 of principals(
                         1. weblogic.security.principal.WLSUserImpl "user1"
                         2. JpsPrincipal:
oracle.security.jps.internal.core.principals.JpsAuthenticatedRoleImpl "authenticated-role"
GUID=null DN=null
                         3. JpsPrincipal:
oracle.security.jps.service.policystore.ApplicationRole "basic_role1"
GUID=734342D04A2811E0AF671B4A95E1598C DN=cn=basic_role1,cn=Roles,cn=JeeScenarioApp,cn=testfarm_
wilu_mlr6,cn=JPSContext,cn=jpsroot
                         4. JpsPrincipal:
oracle.security.jps.service.policystore.ApplicationRole "myrole2"
GUID=738C80D04A2811E0AF671B4A95E1598C DN=cn=myrole2,cn=Roles,cn=JeeScenarioApp,cn=testfarm_wilu_
mlr6,cn=JPSContext,cn=jpsroot
                         5. JpsPrincipal:
oracle.security.jps.internal.core.principals.JpsAnonymousRoleImpl "anonymous-role" GUID=null
DN=null)
                         Permissions=(
                         (oracle.security.jps.service.credstore.CredentialAccessPermission
context=SYSTEM,mapName=default,keyName=* read,write)
                         (oracle.security.jps.service.policystore.PolicyStoreAccessPermission
Context:SYSTEM Context Name:null Actions:getConfiguredApplications)
                         (oracle.security.jps.service.policystore.PolicyStoreAccessPermission
Context:APPLICATION Context Name:* Actions:getApplicationPolicy)
                         (oracle.security.jps.service.policystore.PolicyStoreAccessPermission
```

```
Context:SYSTEM Context Name:null Actions:*)
                        (oracle.security.jps.service.policystore.PolicyStoreAccessPermission
Context:APPLICATION Context Name:* Actions:*)
                    (java.io.FilePermission file2.txt read)
                    (java.io.FilePermission file2.txt write)
                    (java.io.FilePermission file1.txt read)
                    (java.util.PropertyPermission line.separator read)
                    (java.util.PropertyPermission java.vm.specification.version read)
                    (java.util.PropertyPermission java.vm.version read)
                    (java.util.PropertyPermission java.vendor.url read)
                    (java.util.PropertyPermission java.vm.specification.vendor read)
                    (java.util.PropertyPermission java.vm.name read)
                    (java.util.PropertyPermission os.name read)
                    (java.util.PropertyPermission java.vm.vendor read)
                    (java.util.PropertyPermission path.separator read)
                    (java.util.PropertyPermission os.version read)
                    (java.util.PropertyPermission java.specification.name read)
                    (java.util.PropertyPermission os.arch read)
                    (java.util.PropertyPermission java.version read)
                    (java.util.PropertyPermission java.class.version read)
                    (java.util.PropertyPermission java.vendor read)
```

9. Use Oracle Authorization Policy Manager to search for configured security policies that contain the resource and resource type listed in the trace results (look for `resourceName=` and `resourceType=`). In Example 22–3, you would search for configured security polices that contain the resource named `ResourceNameX` that is of the type `TaskFlowResourceType`.

   Refer to the "Finding Application Policies that Match Entitlements or Resources" section in the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)* for more information on how to use Oracle Authorization Policy Manager to search for policies based on resources.

   Note: After identifying the relevant security policies using the "Finding Application Policies that Match Entitlements or Resources" procedure, you will be able to identify the principals and actions granted in each of those configured security policies.

   ---

10. Compare the security policies identified by the search in Step 9 against the relevant `Failed ProtectionDomain` strings in the trace results. Specifically, *for each of the security policies*, compare the granted actions and principals as follows:

    a. Ensure the action granted in the security policies is the same action listed for the `Failed ProtectionDomain` string in the trace results. In Example 22–3, you would ensure the security policy is granting the read action (identified by `Action:[read]` in the trace).

       If the action for the `Failed ProtectionDomain` string is granted in the configured security policy, proceed to Step b.

       If the action for the `Failed ProtectionDomain` string *is not granted* in the security policy, compare the action against all security policies identified by the search in Step 9.

    b. Ensure the principals granted in the security policies are the same principals listed for the `Failed ProtectionDomain` string (identified by `Principals=`).

If the principals configured in the security policy are application roles or external roles and they are not listed in the `Failed ProtectionDomain` string, use Oracle Authorization Policy Manager to determine if the roles are mapped to the relevant user.

> **Note:** Be sure to consult your organization's security policies and the Oracle Fusion Applications security reference manuals before altering any aspect of the configured security policies, as it is possible the user is intentionally unauthorized to access the particular functionality.
>
> You can access the Oracle Fusion Applications security reference manuals in the Oracle Fusion Applications Technology Documentation Library.

If both the actions and principals granted in the security polices are consistent with the authorization request (as identified in the trace), examine Oracle Platform Security Services' cache refresh setting by referring to the problem and solution described in Section 22.3.2.1 of this chapter.

22.3 Problems and Solutions

Use the information in this section if the solution procedures in Section 22.2 did not resolve your problem. This section describes symptoms of specific Oracle Fusion Applications runtime problems that may have originated in the Oracle Identity Management and security integration layer, possible causes of the problems, and solution procedures corresponding to each of the possible causes.

For problems that contain multiple possible causes, the most probable cause and its corresponding solution are listed first. If multiple possible causes are listed, perform the first solution procedure and then retry the failed task. If the problem persists after retrying the failed task, perform the second solution procedure in the topic and then try the failed task again. Repeat this process while proceeding down the list of solution procedures until the problem is resolved.

This section contains the following topics:

- Problems and Solutions for Missing or Incorrect Data

- Problems and Solutions for Accessing Functionality

- Problems and Solutions for Managing Users

- Problems and Solutions for Managing Roles

- Problems and Solutions for Managing Keystores and Certificates

- Problems and Solutions for Identity Propagation Using SAML

> **Notes:** When looking in this section for the problem you encountered, be sure to examine all topics, as many problems fit into multiple topics. For example, while the problem of not being able to see application role hierarchies resides in the topic about managing roles, it could also reside in the topic about missing or incorrect data.

22.3.1 Problems and Solutions for Missing or Incorrect Data

This section describes problems and solutions related to missing or incorrect data. This section contains the following topics:

- LDAP Bulk Changes Not Reconciled in Oracle Identity Manager
- Data is Missing After Migrating or Patching the Policy Store
- Administrator Search for Database Resources Returns No Results
- Data is Missing or Incorrect in a Portlet

22.3.1.1 LDAP Bulk Changes Not Reconciled in Oracle Identity Manager

Bulk changes to an Oracle Internet Directory identity store are not getting reconciled into Oracle Identity Manager.

Problem

The problem may be the Oracle Internet Directory identity store is not configured to generate change logs.

Solution

To verify change log generation is enabled for an Oracle Internet Directory identity store:

Note: If you have multiple Oracle Internet Directory identity store instances, perform this procedure on all of them.

---

1. Log in to Fusion Middleware Control by referring to the "Displaying Fusion Middleware Control" section in the *Oracle Fusion Middleware Administrator's Guide* for more information.

2. Navigate to the appropriate Oracle Internet Directory instance. In the Target Navigation Pane, expand the **Domain > Identity and Access** entries. Alternatively, from the domain home page, expand the **Fusion Middleware > Identity and Access** entries. Oracle Internet Directory instances are listed in both locations. To view the full name of a instance, move the mouse over the instance name.

3. Verify change log generation is enabled. If change log generation is disabled, you must enable it. Refer to the "Enabling or Disabling Change Log Generation by Using Fusion Middleware Control" section of the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*.

22.3.1.2 Data is Missing After Migrating or Patching the Policy Store

After migrating or patching the Oracle Platform Security Services policy store, data that was once available is now missing. This issue may be encountered after the policy store is:

- Migrated from the baseline ("out-of-the-box") jazn-data.xml file policy store to an Oracle Internet Directory instance.

- Migrated from one environment to another, such as moving from a test environment to a production environment.

- Patched using Oracle Authorization Policy Manager.

Problem

The problem may be the application role GUIDs in the Oracle Fusion Data Security repository and their corresponding application role GUIDs in the Oracle Platform Security Services policy store do not match.

Solution

Run the `oracle.apps.fnd.applcore.dataSecurity.util.DSDataMigrator` java program to reconcile the application role GUIDs between the Oracle Fusion Data Security repository and the Oracle Platform Security Services policy store.

Backing Up the fnd_grants Table in the Oracle Fusion Data Security Repository

The `DSDataMigrator` program modifies only the fnd_grants table, which is Virtual Private Database (VPD) enabled. Before running the program, as sys user, back up the existing fnd_grants table in the Oracle Fusion Data Security repository. For example:

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

Running the DSDataMigrator Program

To run the `oracle.apps.fnd.applcore.dataSecurity.util.DSDataMigrator` java program, the following JAR files must be added to the classpath:

```
MW_HOME/ATGPF_HOME/atgpf/modules/oracle.applcore.model_11.1.1/Common-Model.jar
MW_HOME/ATGPF_HOME/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.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/adflogginghandler.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
```

> **Note:** If the classpath is set in the shell, you can run the program from the command line using only the necessary arguments.

The syntax to run the `DSDataMigrator` java program is:

```
java -classpath $CLASSPATH \
-Doracle.security.jps.config=Path_to_jps-config-jse.xml_file \
oracle.apps.fnd.applcore.dataSecurity.util.DSDataMigrator \
-dsdburl URL_to_Oracle_Fusion_Data_Security_repository \
-dsdbuser user_name_for_Oracle_Fusion_Data_Security_repository \
-silentMode [true | false] -forceProcessAllRows [true | false] \
-policyStripe [crm | fscm | hcm]
```

> **Note:** To see usage instructions, execute the following command:
>
> `java oracle.apps.fnd.applcore.dataSecurity.util.DSDataMigrator`

Parameters

The `DSDataMigrator` program supports the following parameters:

- `oracle.security.jps.config`: Identifies the path to the `jps-config-jse.xml` file that the `DSDataMigrator` program will use. For example:

  COMMON_DOMAIN`/config/fmwconfig/jps-config-jse.xml`

  > **Note:** The jps-config-jse.xml file must have credentials for *both* the identity store and policy store—not just the policy store.

- `FND_DS_GUID_RECON_LOG_DIR`: Identifies the output directory for the program's log. For example: `-DFND_DS_GUID_RECON_LOG_DIR=/tmp`

Arguments

The `DSDataMigrator` program supports the following arguments:

- `silentMode`: Set to true if you do not want exceptions to be raised when an entry is not found in the Oracle Platform Security Services policy store.

- `forceProcessAllRows`: Set to true if you want to process all the rows in the policies table. By default, only rows where `compile_flag=Y` are processed.

- `policyStripe`: Identifies the policy stripe to process. Valid values are: crm, fscm and hcm. If the `policyStripe` argument is not specified, all policy stripes and identity store data security role policies are processed.

- `idStoreOnly`: Set to true if you want to process only data security policies made to enterprise roles. If `idStoreOnly` is set to true, the `policyStripe` argument is ignored.

22.3.1.3 Administrator Search for Database Resources Returns No Results

A user with administrator privileges uses Oracle Authorization Policy Manager to search for database resources, but the search does not find any.

Problem

The problem may be data security policies that govern data security administration do not exist in the Oracle Fusion Data Security repository.

Solution

To troubleshoot this issue:

1. Use Oracle Authorization Policy Manager to verify the following application roles are mapped to the external roles of the user performing the search. Refer to the "Mapping Application Roles to an External Role" section in the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)* for more information.

   - `APM_CRM_APPLICATION_OBJECTS_DATA_ADMINISTRATION_DUTY`

- `APM_HCM_APPLICATION_OBJECTS_DATA_ADMINISTRATION_DUTY`

- `APM_FSCM_APPLICATION_OBJECTS_DATA_ADMINISTRATION_DUTY`

- `APM_FND_APPLICATION_OBJECTS_DATA_ADMINISTRATION_DUTY`

Note:

- If the application roles are mapped to the external roles of the user performing the search, go to Step 2.

- If the application roles are not mapped to the external roles of the user performing the search, use Oracle Authorization Policy Manager to map them to the user's external roles and then go to Step 2.

---

2. Determine whether data security policies that govern data security administration exist in the Oracle Fusion Data Security repository. Log in to Oracle Authorization Policy Manager as a user with the Application Developer external role and search for the following roles. Ensure that data security policies for the roles exist on the `FND_OBJECTS` object and that the policies have not expired.

    - `APM_CRM_APPLICATION_OBJECTS_DATA_ADMINISTRATION_DUTY`

    - `APM_HCM_APPLICATION_OBJECTS_DATA_ADMINISTRATION_DUTY`

    - `APM_FSCM_APPLICATION_OBJECTS_DATA_ADMINISTRATION_DUTY`

    - `APM_FND_APPLICATION_OBJECTS_DATA_ADMINISTRATION_DUTY`

    If the policies do not exist in the Oracle Fusion Data Security repository, use Oracle Fusion Functional Setup Manager to upload the Applications Core data security seed data to the Oracle Fusion Data Security repository. Refer to Section 2.8 for more information about using Oracle Fusion Functional Setup Manager.

22.3.1.4 Data is Missing or Incorrect in a Portlet

After logging in to an Oracle Fusion Applications portlet, the data the user expects to see is missing or incorrect.

Problem

The problem may be:

- The application user session was not propagated to the portlet.

- The application user session was not created using the portlet's application stripe and Applications Core did not compute the application roles for the portlet's application stripe.

Solution

To troubleshoot this situation:

1. Log out of the portlet, and then log in again.

2. Execute the following diagnostic tests. Refer to Section 13.10.1 for more information about running diagnostic tests.

    - Data Security Configuration

    - Data Security Configuration with Application User Session Prerequisite

    - Data Security Run Time

■ Data Security Run Time with Application User Session Prerequisite

22.3.2 Problems and Solutions for Accessing Functionality

This section describes problems and solutions related to accessing functionality. This section contains the following topics:

■ Inappropriate User Access After Enterprise Role Membership Removal

■ Newly Created User Does Not Have Correct Access to Oracle Fusion Applications

■ After Logging Out, Access to a Secured Resource is Granted Without Logging in

■ Authenticated User Gets Unexpected Page when Accessing a Different Secured Resource

■ Support Representative Cannot Impersonate an Oracle Fusion Applications User

■ Unauthenticated User Gets Error Page when Accessing a Secured Resource

22.3.2.1 Inappropriate User Access After Enterprise Role Membership Removal

After removing an enterprise role's membership to an application role using Oracle Authorization Policy Manager, access to the application is still being granted.

Problem

Oracle Platform Security Services optimizes the authorization process by caching security artifacts. When an application policy (or some other security artifact) is modified, the change becomes effective depending on where the application and the tool used to modified the artifact (Oracle Authorization Policy Manager in this case) are running.

If the application and the tool (Oracle Authorization Policy Manager) are running on different hosts or in different domains, the change becomes effective after the policy store cache is refreshed. The frequency of the cache refresh is determined by the value of the **Refresh Polling Time (secs)** parameter in Fusion Middleware Control.

Depending on the configuration, access to the application may have been granted (despite the removal of the enterprise role membership to the application role) because the Oracle Platform Security Services cache was not refreshed before the application was accessed.

Refer to the "Caching and Refreshing the Cache" and "An Example" sections in the *Oracle Fusion Middleware Application Security Guide* for more information about authorization behavior relating to the Oracle Platform Security Services cache.

Solution

To examine the refresh interval for Oracle Platform Security Services' cache:

1. Log in to Fusion Middleware Control by referring to the "Displaying Fusion Middleware Control" section in the *Oracle Fusion Middleware Administrator's Guide* for more information.

2. Click the name of the appropriate domain in the target navigation pane on the left side of the screen.

3. Select **Security > Security Provider Configuration** from the domain menu at the top of the screen. The Security Provider Configuration screen appears.

4. Select the **Policy Store Credential Store Keystore** entry in the **Security Stores** table and click **Edit**. The Edit Security Provider Configuration screen appears.

5. Examine the value set for the **Refresh Polling Time (secs)** parameter.

6. Wait for the amount of time specified by the **Refresh Polling Time (secs)** parameter to elapse, then retry the use case. This will ensure that the policy store cache has been refreshed and any recent changes to policies are effective.

22.3.2.2 Newly Created User Does Not Have Correct Access to Oracle Fusion Applications

After creating a new user and external role using Oracle Fusion Human Capital Management, then granting duty roles to that user using Oracle Authorization Policy Manager, the user cannot log in and perform its granted duties.

Problem

The problem may be:

- The user does not exist in the identity store.

- The user to external role membership does not exist in the identity store.

- The Oracle Internet Directory Authenticator's cache or Oracle Platform Security Services' cache has not yet been refreshed.

- Oracle Identity Manager and Oracle Authorization Policy Manager are not configured to use the same identity store or their connection settings to identity store are incorrect.

Solution 1

To verify the user exists in the identity store, use Oracle Directory Services Manager to examine the container in the identity store where users are stored, such as `cn=users,dc=us,dc=oracle,dc=com`.

Refer to the following for more information about examining identity store containers.

If Oracle Internet Directory is the identity store, refer to the following sections in the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory* using this sequence:

1. "Invoking Oracle Directory Services Manager"

2. "Connecting to the Server from Oracle Directory Services Manager"

3. "Displaying Entries by Using Oracle Directory Services Manager"

If Oracle Virtual Directory is the identity store, refer to the following sections in the *Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory* using this sequence:

1. "Invoking Oracle Directory Services Manager"

2. "Connecting to the Server from Oracle Directory Services Manager"

3. "Viewing Oracle Virtual Directory Entries"

Solution 2

To verify the user to external role membership exists in the identity store:

1. Verify the user exists in the identity store. Use Oracle Directory Services Manager to examine the container in the identity store where users are stored, such as `cn=users,dc=us,dc=oracle,dc=com`.

2. Verify the external role exists in the identity store. Use Oracle Directory Services Manager to examine the container where enterprise roles are stored, such as `cn=groups`.

3. Verify the user is a member of the external role. Use Oracle Directory Services Manager to confirm `uniqueMember` is an attribute of the external role.

4. Use Oracle Authorization Policy Manager to verify the external role is mapped to the appropriate application role. Perform a simple search on the application role, open it, and click the **External Role Mapping** tab. Refer to 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)* for more information.

Solution 3

To troubleshoot the Oracle Internet Directory Authenticator's cache and Oracle Platform Security Services' cache:

1. Examine the Oracle Internet Directory Authenticator's cache settings by referring to Section 22.2.2.

2. Examine Oracle Platform Security Services' cache refresh setting by referring to the problem and solution described in Section 22.3.2.1 of this chapter.

> **Note:** Wait for the caches to be refreshed before retrying any failed task or operation.

Solution 4

To verify Oracle Identity Manager and Oracle Authorization Policy Manager are configured to use the same identity store and their connection settings to identity store are correct:

1. Identify the identity store that Oracle Identity Manager is using by performing the following steps:

    a. Log in to the Advanced view of the Oracle Identity Manager Administrative and User Console. You can access the Advanced view by entering a URL similar to the following into a web browser:

        `http://HOST:PORT/oim/admin/`

    b. Click **Manage IT Resource**. The Manage IT Resource screen appears.

    c. Enter `Directory Server` in the IT Resource Name field or select **Directory Server** from the IT Resource Type list and click **Search**. The search results appear at the bottom of the screen.

    d. Click the **Directory Server** link in the search results. The configuration details for the identity store appear. Examine and make a note of the connection settings to the identity store.

2. Identify the identity store that Oracle Authorization Policy Manager is using by examining the connection settings configured for the LDAP Authenticators in the Oracle WebLogic Server domain. To examine the LDAP Authenticators' configuration, refer to Section 22.2.2.

22.3.2.3 After Logging Out, Access to a Secured Resource is Granted Without Logging in

After logging out of a resource secured by Oracle Access Manager and then attempting to access a different secured resource, access is granted without a login page appearing.

> **Note:** Oracle Platform Security Services manages logouts for Oracle Fusion Applications by providing the configured logout URL (typically the Oracle Access Manager logout URL) to Oracle Application Development Framework for redirection. Oracle Access Manager then sets the session status to logged out.

Problem

The problem may be:

- Oracle Access Manager's user session was not removed during logout.

- Oracle Platform Security Services is not configured with the correct Oracle Access Manager logout URL.

Solution 1

Perform either of the following steps to determine whether Oracle Access Manager's user session was removed during logout:

- Examine the cookies in the user's browser. Oracle Access Manager's `OAM_ID` session cookie should *not* be present, as it gets deleted from the browser upon logout.

- Use the Oracle Access Manager Administration Console's Session Management functionality to examine the active sessions. Search on the user to see if any of its sessions are active.

  Refer to the "Logging In to the Oracle Access Manager 11*g* Administration Console" and "Managing Active User Sessions" sections in the *Oracle Fusion Middleware Administrator's Guide for Oracle Access Manager with Oracle Security Token Service*.

Solution 2

To verify Oracle Platform Security Services is configured with the correct Oracle Access Manager logout URL:

1. Log in to Fusion Middleware Control by referring to the "Displaying Fusion Middleware Control" section in the *Oracle Fusion Middleware Administrator's Guide*.

2. Select the appropriate domain from the target navigation pane or the content pane.

3. Select **Security > Security Provider Configuration** from the domain menu. The Security Provider Configuration page appears.

4. Expand the **Single Sign-On Provider** area if it is not already expanded and click the **Configure** button. The Single Sign-On Provider page appears.

5. Select the **Configure Single Sign-on** option. All settings on the Single Sign-On Provider page are invisible until you select the **Configure Single Sign-on** option.

6. Examine the value set in the Logout URL field.

22.3.2.4 Authenticated User Gets Unexpected Page when Accessing a Different Secured Resource

After successfully logging in to and working on a resource secured by Oracle Access Manager and then attempting to access a different secured resource, an unexpected page, such as `Not Authorized`, blank (empty), corrupted, or `500 error`, appears.

Problem

The problem may be Oracle Access Manager's `ObSSOCookie` and `OAM_ID` cookies are not in the user's browser. The `ObSSOCookie` and `OAM_ID` cookies are encrypted, single sign-on, session-based cookies generated by the Oracle Access Manager Access Server when a user authenticates successfully.

Solution

To verify Oracle Access Manager's `ObSSOCookie` and `OAM_ID` cookies are in the user's browser:

1. Display the cookies in the user's browser.

2. Locate Oracle Access Manager's `ObSSOCookie` and `OAM_ID` session cookies.

If the `ObSSOCookie` and `OAM_ID` cookies are *not* in the user's browser:

- Examine the browser's security settings, as they may be too high and preventing cookies from being accepted

- Add the Oracle Fusion application's domain to the browser's exception list

22.3.2.5 Support Representative Cannot Impersonate an Oracle Fusion Applications User

A Support (Help Desk) representative attempts to log in to a resource secured by Oracle Access Manager and impersonate an Oracle Fusion Applications user, but cannot do so.

Problem

The problem may be the user that the Support representative is attempting to impersonate has not granted the privilege to be impersonated or the privilege has expired.

Solution

To verify that the user has granted the privilege to be impersonated and that the privilege is active:

1. Use Oracle Directory Services Manager to locate the account of the user to be impersonated in the identity store. Look in the container where users are stored, such as `cn=users,dc=us,dc=oracle,dc=com`.

   Refer to the following for more information about examining identity store containers.

   If Oracle Internet Directory is the identity store, refer to the following sections in the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory* using this sequence:

   a. "Invoking Oracle Directory Services Manager"

   b. "Connecting to the Server from Oracle Directory Services Manager"

   c. "Displaying Entries by Using Oracle Directory Services Manager"

If Oracle Virtual Directory is the identity store, refer to the following sections in the *Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory* using this sequence:

a. "Invoking Oracle Directory Services Manager"

b. "Connecting to the Server from Oracle Directory Services Manager"

c. "Viewing Oracle Virtual Directory Entries"

2. Verify the user has granted the privilege to be impersonated by examining the user's account for the `orclimpersonationgrants` attribute.

■ If the user's account does not have the `orclimpersonationgrants` attribute, the user has not granted the privilege to be impersonated.

■ If the user's account has the `orclimpersonationgrants` attribute, ensure the privilege has not expired. The `orclimpersonationgrants` attribute will be in a format such as:

```
EEA958988E344BF49740CF00DF9B0421|20110124170000Z|20110124180000Z
```

– `EEA958988E344BF49740CF00DF9B0421` is the GUID of the impersonator.

– `20110124170000Z` is the date on which impersonation can begin

– `20110124180000Z` is the expiration date for the impersonation privilege

Note: The date strings in the `orclimpersonationgrants` attribute use the Coordinated Universal Time (UTC) standard and are of the form: `yyyyMMddHHmmss'Z'`

22.3.2.6 Unauthenticated User Gets Error Page when Accessing a Secured Resource

While attempting to access a resource secured by Oracle Access Manager, an unauthenticated user gets an error page instead of the login page.

Problem

The problem may be:

■ The Oracle HTTP Server Web servers front-ending the Oracle Fusion application are not running.

■ The Managed Servers where Oracle Access Manager is deployed or the requisite Oracle Access Manager services are not running.

Solution 1

To verify the Oracle HTTP Server Web servers front-ending the Oracle Fusion application are running:

1. Connect to a page provided by Oracle Identity Manager. If Oracle Identity Manager is front-ended by Oracle HTTP Server or a load balancer, use the following URL:

```
http(s)://FRONTEND_HOST:FRONTEND_PORT/admin/faces/pages/accountlocked.jspx
```

If Oracle Identity Manager is not front-ended by Oracle HTTP Server or a load balancer, use the following URL:

```
http(s)://OIM_MANAGED_SERVER_HOST:PORT/admin/faces/pages/accountlocked.jspx
```

2. Connect to any public page provided by an Oracle Fusion application through Oracle HTTP Server. For example:

```
http(s)://ORACLE_HTTP_SERVER_FRONTEND_HOST:PORT/fa/app/index.jsp
```

If you cannot access a page in an Oracle HTTP Server front-ending configuration, use Fusion Middleware Control to examine the WebLogic Host and WebLogic Port settings for the Oracle HTTP Server's `mod_wl_ohs` module. Refer to the "Configuring the mod_wl_ohs Module" section in the *Oracle Fusion Middleware Administrator's Guide for Oracle HTTP Server* for more information.

Solution 2

To verify the Managed Servers where Oracle Access Manager is deployed and the requisite Oracle Access Manager services are running:

1. Verify the Managed Servers where Oracle Access Manager is deployed are running by performing the following steps:

   a. Log in to the Oracle WebLogic Server Administration Console by referring to the "Starting the Administration Console" section in the *Oracle Fusion Middleware Introduction to Oracle WebLogic Server* document.

   b. Click **Servers** in the Environment section on the Home page. The Summary of Servers page appears.

   c. Click the **Configuration** tab. A table containing a summary of each server in the domain appears.

   d. Examine the **State and Health** columns for the Managed Servers where Oracle Access Manager is deployed.

2. Verify the HTTP port is open by attempting to connect to it. If Oracle Access Manager is front-ended by Oracle HTTP Server or a load balancer, enter the following URL into a web browser:

```
http://ORACLE_HTTP_SERVER-or-LOAD_BALANCER_HOST:PORT/oam/pages/logout.jsp
```

If Oracle Access Manager is not front-ended, enter the following URL into a web browser:

```
http://MANAGED_SERVER_HOST:PORT/oam/pages/logout.jsp
```

3. Verify Oracle Access Manager authentication is functioning properly by accessing any resource secured by Oracle Access Manager. For example, log in to the Oracle Access Manager Administration Console by referring to the "Logging In to the Oracle Access Manager 11*g* Administration Console" section in the *Oracle Fusion Middleware Administrator's Guide for Oracle Access Manager with Oracle Security Token Service*.

22.3.3 Problems and Solutions for Managing Users

This section describes problems and solutions related to managing users. This section contains the following topics:

- Oracle Fusion Human Capital Management Requests to Assign Roles to Users Fail

- SPML Calls Initiated from an Oracle Fusion Application are Not Processed

- Troubleshooting Oracle Fusion Human Capital Management-Oracle Identity Manager SPML Requests

22.3.3.1 Oracle Fusion Human Capital Management Requests to Assign Roles to Users Fail

Oracle Fusion Human Capital Management makes a request to assign a role to a user, but the role assignment fails.

Problem

The problem may be the user exists in Oracle Identity Manager, but does not exist in the Oracle Internet Directory identity store.

Solution

To troubleshoot this situation:

1. Verify the user does not exist in Oracle Internet Directory by using Oracle Directory Services Manager to examine the container where users are stored, such as `cn=users,dc=us,dc=oracle,dc=com`.

    Refer to the "Displaying Entries by Using Oracle Directory Services Manager" or "Searching for Entries by Using Oracle Directory Services Manager" sections in the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory* for more information.

    > **Note:**
    >
    > - If the user does not exist in Oracle Internet Directory, continue this procedure.
    >
    > - If the user exists in Oracle Internet Directory, perform the steps in Section 22.3.3.3 to get information about why the role assignment failed.

2. Execute the LDAP User Create and Update Reconciliation scheduled job in Oracle Identity Manager. After the job executes, the user account will be removed from Oracle Identity Manager and requests from Oracle HCM for the user will not be created.

    Refer to the "Managing Scheduled Tasks" chapter in the *Oracle Fusion Middleware Administrator's Guide for Oracle Identity Manager* for information about the LDAP User Create and Update Reconciliation scheduled job and how to execute it.

22.3.3.2 SPML Calls Initiated from an Oracle Fusion Application are Not Processed

An Oracle Fusion application invokes an event that initiates an SPML call, for example, Oracle Fusion Human Capital Management sends an SPML request to add a user, but the call fails.

Problem

The problem may be the Oracle Fusion application's composite that invokes the SPML service is using incompatible Oracle Web Services Manager (Oracle WSM) client and server security policies. The client and server security policies must be compatible for calls to succeed.

Solution

To verify Oracle Identity Manager is using the correct Oracle WSM server and client security policies:

1. Verify Oracle Identity Manager is using the correct Oracle WSM server security policy by performing the following steps:

   a. Log in to Fusion Middleware Control on the Oracle WebLogic Server where Oracle Identity Manager is installed. Refer to the "Displaying Fusion Middleware Control" section in the *Oracle Fusion Middleware Administrator's Guide* for more information.

   b. Expand the **Application Deployments** entry in the navigation tree and click **spml-xsd**. The spml-xsd details page appears.

   c. Select **Web Services** from the Application Deployment list. The Web Services tab appears.

   d. Click **SPMLServiceProviderSoap**. The details for the web service appear.

   e. Click the **OWSM Policies** tab.

   f. Verify the following policy is listed as a Directly Attached Policy:

   oracle/wss_saml_or_username_token_service_policy

2. Verify the Oracle Fusion application is using the correct Oracle WSM client security policy by performing the following steps:

   a. Log in to Fusion Middleware Control on the Oracle WebLogic Server where the Oracle Fusion application is running.

   b. Expand the **SOA** entry in the navigation tree and all of its child entries until the list of configured composites appear and then click the name of the appropriate composite. The details of the composite appear.

   c. Click the **Policies** tab.

   d. Verify the following policy is attached to the composite's end points:

   oracle/wss_username_token_client_policy

22.3.3.3 Troubleshooting Oracle Fusion Human Capital Management-Oracle Identity Manager SPML Requests

To collect information about SPML requests between Oracle Fusion Human Capital Management and Oracle Identity Manager:

1. Identify the ID number of the request you want to investigate. After an Oracle Fusion Human Capital Management application performs an operation that sends an SPML request to Oracle Identity Manager, Oracle Identity Manager creates a unique ID for that specific request and returns it to the application. From the application, identify the request ID.

2. Use the Advanced view of the Oracle Identity Manager Administrative and User Console to see general information about the request, such as its status.

   a. Log in to the Advanced view of the Oracle Identity Manager Administrative and User Console by entering a URL similar to the following into a web browser:

   ```
   http://HOST:PORT/oim/admin/
   ```

   b. Click the **Administration** tab, then click **Requests**.

   c. Search for the request by entering the request ID in the search field and clicking **Search**.

d. Click the request ID in the search results. Information about the request appears.

e. Examine the status of the request. If the status is Request Failed, a hyperlink to additional information about the failed request is provided. Click the **Request Failed** link to see more information.

Refer to the "Searching and Tracking Requests" section in the *Oracle Fusion Middleware User's Guide for Oracle Identity Manager* for more information.

> **Note:** To see more detailed information about the request, proceed to the next step in this procedure.

3. Use the Oracle Identity Manager Diagnostic Dashboard's Orchestration Status test to see information such as which Oracle Identity Manager event handlers handled the request and its status at each event handler.

a. Log in to the Oracle Identity Manager Diagnostic Dashboard.

Refer to the "Working with the Diagnostic Dashboard" section in the *Oracle Fusion Middleware Administrator's Guide for Oracle Identity Manager* for more information.

b. Run the **Request Diagnostic Information** test on the request ID. The test will return and Orchestration ID.

c. Run the **Orchestration Status** test on the Orchestration ID. Detailed information about the request appears.

> **Note:** To see the most information available about the request, proceed to the next step in this procedure.

4. Use Fusion Middleware Control to view the Oracle Identity Manager diagnostic logs. Examine the log files to see the most detailed information about the SPML request. Refer to the "Managing Log Files and Diagnostic Data" chapter in the *Oracle Fusion Middleware Administrator's Guide* for more information.

22.3.4 Problems and Solutions for Managing Roles

This section describes problems and solutions related to managing roles. This section contains the following topics:

- Cannot See the Function Security Policies for an External Role
- Cannot See the Data Security Policies for a Data Role
- Problems Mapping an Application Role to an External Role
- Cannot See Application Role Hierarchies
- Attempts to Add an Application Role to a Hierarchy Appear to Have No Effect
- Cannot Create Valid Data Roles Using Data Role Template

22.3.4.1 Cannot See the Function Security Policies for an External Role

The function security polices for a particular external role cannot be seen using Oracle Authorization Policy Manager.

Problem

The problem may be:

- If Oracle Internet Directory is being used as the identity store, it is not configured to index the `displayName` attribute. If Oracle Internet Directory is not indexing the `displayName` attribute, Oracle Authorization Policy Manager cannot retrieve the role during a search.

- The Oracle Internet Directory Authenticator in the Oracle WebLogic Server domain is not configured with the correct connection settings to the Oracle Internet Directory instance.

- The external role has not been provisioned into the identity store.

- If the administrator attempting to identify the function security polices is configured as a Delegated Administrator, the Delegated Administrator role does not have access to the appropriate application stripe.

- The policy store does not have the correct application stripes.

- The external role is not mapped to the correct application roles.

- The external role is mapped to an application role that does not have policy attached to it.

Solution 1

To verify an Oracle Internet Directory identity store is configured to index the `displayName` attribute, refer to Section 22.2.1.

Solution 2

To verify the Oracle Internet Directory Authenticator in the Oracle WebLogic Server domain is configured with the correct connection settings to the Oracle Internet Directory instance, refer to Section 22.2.2.

Solution 3

To verify the external role was provisioned into the identity store, use Oracle Directory Services Manager to examine the container in the identity store where external roles are stored, such as: `cn=groups,dc=mycompany,dc=com`.

> **Note:**
> - If the external role does not exist in the identity store, use Oracle Fusion Human Capital Management to add it to the identity store.
> - If the external role exists in the identity store, verify the security providers in the Oracle WebLogic Server domain are configured in the correct order and with the correct JAAS Control Flags by referring to Section 22.2.2.

Refer to the following for more information about examining identity store containers.

If Oracle Internet Directory is the identity store, refer to the following sections in the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory* using this sequence:

1. "Invoking Oracle Directory Services Manager"

2. "Connecting to the Server from Oracle Directory Services Manager"

3. "Displaying Entries by Using Oracle Directory Services Manager"

If Oracle Virtual Directory is the identity store, refer to the following sections in the *Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory* using this sequence:

1. "Invoking Oracle Directory Services Manager"

2. "Connecting to the Server from Oracle Directory Services Manager"

3. "Viewing Oracle Virtual Directory Entries"

Solution 4

If the administrator attempting to identify the function security polices is configured as a Delegated Administrator, verify the Delegated Administrator role has access to the appropriate application stripe by referring to the "Delegated Administration" chapter in the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)*.

Solution 5

To verify the policy store has the correct application stripes:

1. Identify the application stripes that were loaded into the policy store after the Oracle Fusion Applications environment was provisioned by referring to the Oracle Fusion Applications security reference manuals. You can access the Oracle Fusion Applications security reference manuals in the Oracle Fusion Applications Technology Documentation Library.

2. Verify the application stripes identified in Step 1 exist in the policy store by performing the following steps:

   a. Log in to Oracle Authorization Policy Manager as a security administrator with the APMAdmin application role, which will allow you to see all application stripes in the policy store.

   b. Examine the Browse tab of the Navigation Panel, which lists all policy stripes in the policy store (because you are logged in as a security administrator with the APMAdmin application role).

      Refer to the "The Navigation Panel" section in the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)* for more information about viewing application stripes in the policy store using the Navigation Panel.

Solution 6

To verify the external role is mapped to the correct application roles:

1. Verify the application stripe that the application role is expected to be in exists in the policy store by performing the following steps:

   a. Determine which application stripe the application role is expected to be in by referring to the Oracle Fusion Applications security reference manuals. You can access the Oracle Fusion Applications security reference manuals in the Oracle Fusion Applications Technology Documentation Library.

   b. Log in to Oracle Authorization Policy Manager as a security administrator with the APMAdmin application role, which will allow you to see all application stripes in the policy store.

   c. Examine the **Browse** tab of the Navigation Panel, which lists all policy stripes in the policy store (because you are logged in as a security administrator with

the `APMAdmin` application role). Verify the application stripe identified in Step a exists in the policy store.

Refer to the "The Navigation Panel" section in the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)* for more information about viewing application stripes in the policy store using the Navigation Panel.

2. Use Oracle Authorization Policy Manager to identify the application roles currently mapped to the external role.

Refer to the "Mapping Application Roles to an External Role" section in the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)* for more information.

3. Compare the application roles identified in Step 2 to the application roles listed for the external role in the Oracle Fusion Applications security reference manuals.

If any application roles listed in the Oracle Fusion Applications security reference manuals are not mapped to the external role, use Oracle Authorization Policy Manager to see if they exist in the policy store. Refer to the "Searching Application Roles" section in the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)* for more information.

If the application roles exist in the policy store, use Oracle Authorization Policy Manager to map them to the external role by referring to the "Mapping Application Roles to an External Role" section in the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)*.

If the application roles do not exist in the policy store, use Oracle Authorization Policy Manager to create them by referring to the "Creating a Role" section in the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)*. After creating the application roles, map them to the external role.

Solution 7

To verify the external role is mapped to an application role that has policy attached to it, refer to the "Finding Application Policies that Match Principals" section in the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)*.

22.3.4.2 Cannot See the Data Security Policies for a Data Role

Data security polices for a particular data role cannot be seen in Oracle Authorization Policy Manager.

Problem

The problem may be:

■ The Oracle Internet Directory Authenticator in the Oracle WebLogic Server domain is not configured with the correct connection settings to the Oracle Internet Directory instance.

■ If Oracle Internet Directory is being used as the identity store, it is not configured to index the `displayName` attribute. If Oracle Internet Directory is not indexing the `displayName` attribute, Oracle Authorization Policy Manager cannot retrieve the role during a search.

- The user searching for the data security policies does not have the privileges to do so.

- The data role does not exist in the identity store.

- Data role templates did not create data security policies for the data role.

- The data security role GUIDs in the Oracle Fusion Data Security repository and the Oracle Platform Security Services policy store are not synchronized.

Solution 1

To verify the Oracle Internet Directory Authenticator in the Oracle WebLogic Server domain is configured with the correct connection settings to the Oracle Internet Directory instance, refer to Section 22.2.2.

Solution 2

To verify an Oracle Internet Directory identity store is configured to index the `displayName` attribute, refer to Section 22.2.1.

Solution 3

To verify the user searching for the data security policies has the privileges to do so, perform the solution described in Section 22.3.1.3.

Solution 4

To verify the data role exists in the identity store, use Oracle Directory Services Manager to examine the container in the identity store where data roles are stored, such as `cn=groups,dc=`*mycompany*`,dc=`*com*. If the role does not exist in the identity store, an administrator should add it.

Refer to the following for more information about examining identity store containers.

If Oracle Internet Directory is the identity store, refer to the following sections in the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory* using this sequence:

1. "Invoking Oracle Directory Services Manager"

2. "Connecting to the Server from Oracle Directory Services Manager"

3. "Displaying Entries by Using Oracle Directory Services Manager"

If Oracle Virtual Directory is the identity store, refer to the following sections in the *Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory* using this sequence:

1. "Invoking Oracle Directory Services Manager"

2. "Connecting to the Server from Oracle Directory Services Manager"

3. "Viewing Oracle Virtual Directory Entries"

Solution 5

To verify data role templates created data security policies for the data role:

1. Use Oracle Authorization Policy Manager to perform a simple search for the data role using **External Role** as the object type. Refer to the "Finding Artifacts with a Simple Search" section in the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)* for more information.

2. Select the data role in the search results and click the **View** button. Details about the data role appear.

3. Click the **Find Global Policies** button. The **Data Security Policies** table appears and lists the data security policies attached to the data role. Examine the entries in the table to ensure the data role template created the appropriate data security policies.

Solution 6

To reconcile the data security role GUIDs in the Oracle Fusion Data Security repository and the Oracle Platform Security Services policy store, run the `oracle.apps.fnd.applcore.dataSecurity.util.DSDataMigrator` java program to reconcile the GUIDs. Refer to the solution in Section 22.3.1.2 for information about using this program.

22.3.4.3 Problems Mapping an Application Role to an External Role

While attempting to map an application role to an external role using Oracle Authorization Policy Manager, issues such as the following are encountered:

- Either the external role or application role cannot be seen in Oracle Authorization Policy Manager.

- The mapping succeeds in Oracle Authorization Policy Manager, but is activated after a delay.

Problem

The problem may be:

- If Oracle Internet Directory is being used as the identity store, it is not configured to index the `displayName` attribute. If Oracle Internet Directory is not indexing the `displayName` attribute, Oracle Authorization Policy Manager cannot retrieve the roles during a search.

- The security providers for the Oracle WebLogic Server domain are configured incorrectly. Specifically, the order of providers, JAAS Control Flags, or connection settings to the Oracle Internet Directory instance may be incorrect.

- If the mapping succeeds in Oracle Authorization Policy Manager, but is activated after a delay, the cache refresh settings for the Oracle Internet Directory Authenticator or for Oracle Platform Security Services may need to be adjusted.

Solution 1

To verify an Oracle Internet Directory identity store is configured to index the `displayName` attribute, refer to Section 22.2.1.

Solution 2

To troubleshoot the configuration of the security providers for the Oracle WebLogic Server domain, perform the steps in Section 22.2.2 and examine the:

- Order of providers

- JAAS Control Flags

- Connection settings to the Oracle Internet Directory instance

Solution 3

If the mapping is activated after a delay, to troubleshoot the cache refresh settings for the Oracle Internet Directory Authenticator and for Oracle Platform Security Services:

1. Examine the Oracle Internet Directory Authenticator's cache settings by referring to Section 22.2.2.

2. Examine Oracle Platform Security Services' cache refresh setting by referring to the problem and solution described in Section 22.3.2.1 of this chapter.

> **Note:** Wait for the caches to be refreshed before reattempting a failed task or operation.

22.3.4.4 Cannot See Application Role Hierarchies

Attempts to view application role hierarchies using Oracle Authorization Policy Manager fail.

Problem

The problem may be:

- The identity store's LDAP Authenticator in the Oracle WebLogic Server domain is configured to use the wrong identity store.

- The administrator is attempting to view the application role hierarchy from the incorrect application role in the Oracle Authorization Policy Manager interface.

- Role hierarchies are not defined.

Solution 1

To verify the identity store's LDAP Authenticator in the Oracle WebLogic Server domain is configured to use the correct identity store, refer to Section 22.2.2 and examine the connection settings configured for the identity store's LDAP Authenticator.

Solution 2

To verify the correct application role is being used to display the application role hierarchy, in the Oracle Authorization Policy Manager interface, ensure attempts to display the role hierarchy are based on the correct application role. Application roles frequently have similar names, such as roles that are qualified by region. Double-check that the intended application role is being used to display the role hierarchy.

Solution 3

To verify role hierarchies are defined, perform the procedures in the "Managing the Application Role Hierarchy" section in the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)*.

22.3.4.5 Attempts to Add an Application Role to a Hierarchy Appear to Have No Effect

After using Oracle Authorization Policy Manager to add an application role to a hierarchy, no changes can be seen in the hierarchy.

Problem

The problem may be:

- The application role already exists as a member of the hierarchy.

- The incorrect application role was added to the hierarchy, or the correct application role was added to the incorrect hierarchy.

Solution

To verify the application role hierarchy:

1. Display the application role hierarchy the role was intended for. Refer to the "Managing the Application Role Hierarchy" section in the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)* for information about viewing the application role hierarchy.

2. Ensure that the application role does not already exist in the hierarchy.

3. Ensure that when the application role was added to the hierarchy, the intended application role and the intended hierarchy were used. It is possible the intended application role was added to the incorrect hierarchy, or the incorrect application role was added to the intended hierarchy.

Refer to the "Permission Inheritance and the Role Hierarchy" section in the *Oracle Fusion Middleware Application Security Guide* for information about rules for application role hierarchies.

22.3.4.6 Cannot Create Valid Data Roles Using Data Role Template

While attempting to create a data role using a data role template in Oracle Authorization Policy Manager, issues such as the following are encountered:

- The data role is not created

- The data role is created with a null displayName and description

Problem

The problem may be:

- The SQL query used in the Dimension tab of the template is invalid or returns no records.

- The Oracle Authorization Policy Manager application ID used by the data role template does not have sufficient privileges to create the data role in the intended identity store container.

- A general issue in the identity store, such as the instance is not running.

Solution 1

To troubleshoot the SQL query used in the Dimension tab of the template:

1. Review the SQL query and ensure the intended string was entered correctly.

2. Review the SQL query and ensure it does not contain special characters such as "," (comma) that are unsupported by the identity store. Role names must be comprised of only alphanumeric characters.

3. Verify the database table referenced in the SQL query contains data (is not empty).

Solution 2

To troubleshoot the privileges of the Oracle Authorization Policy Manager application ID used by the data role template, perform the following steps on the identity store:

1. Verify the `cn=fusion_apps_apm_rgx_appid` user exists in the `cn=appidusers container`.

   Refer to the following for more information about examining identity store containers.

   If Oracle Internet Directory is the identity store, refer to the following sections in the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory* using this sequence:

   a. "Invoking Oracle Directory Services Manager"

   b. "Connecting to the Server from Oracle Directory Services Manager"

   c. "Displaying Entries by Using Oracle Directory Services Manager"

   If Oracle Virtual Directory is the identity store, refer to the following sections in the *Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory* using this sequence:

   a. "Invoking Oracle Directory Services Manager"

   b. "Connecting to the Server from Oracle Directory Services Manager"

   c. "Viewing Oracle Virtual Directory Entries"

2. Verify the `cn=fusion_apps_apm_rgx_appid` group exists in the `cn=appidgroups` container.

3. Identify all groups that the `fusion_apps_apm_rgx_appid` group is a member of, and then verify those groups have write permission to the container where enterprise roles are stored, such as `cn=groups`.

   > **Note:** If using Oracle Virtual Directory as the identity store, you must verify the groups' permissions *in both* Oracle Virtual Directory and the back-end (source) repositories.

4. If using Oracle Virtual Directory as the identity store, verify that the ACLs for Oracle Virtual Directory and its back-end (source) data repositories are configured correctly.

   To focus the ACL verification, perform the following steps:

   a. Temporarily disable access control checking in Oracle Virtual Directory using Fusion Middleware Control. To disable access control checking, deselect (disable) the **Enable Access Control Check** option on Oracle Virtual Directory's Server Properties page.

      Refer to the "Configuring Oracle Virtual Directory Server Properties Using Fusion Middleware Control" section in *Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory* for more information.

   b. Perform the steps to create the data role using a data role template.

      – If you can create the data role when Oracle Virtual Directory access control checking is disabled, the Oracle Virtual Directory ACLs are configured incorrectly.

        To isolate the error in the Oracle Virtual Directory ACLs, re-enable access control checking in Oracle Virtual Directory, set its logging to TRACE message type at level 32, try creating the data role using a data role template,

and then examine Oracle Virtual Directory's log, which will now contain the result of each ACL test.

Refer to the "Setting the Level of Information Written to Log Files" section and the "Managing Log Files and Diagnostic Data" chapter in the *Oracle Fusion Middleware Administrator's Guide* for more information about Oracle Virtual Directory logging.

– If you cannot create the data role when Oracle Virtual Directory access control checking is disabled, the error is not in the Oracle Virtual Directory ACLs and you should examine the ACLs in the back-end (source) data repositories by referring to their documentation.

Solution 3

To troubleshoot the identity store:

- If using Oracle Internet Directory as the identity store:

  1. Verify Oracle Internet Directory is running.

     You can view the status of Oracle Internet Directory using Fusion Middleware Control. After logging in to Fusion Middleware Control, navigate to the Farm home page and examine the Identity and Access components within the Fusion Middleware section of the content pane.

     Refer to the "Displaying Fusion Middleware Control" section in the *Oracle Fusion Middleware Administrator's Guide* for more information.

  2. Verify Oracle Internet Directory is configured to index the `displayName` attribute by referring to .

- If using Oracle Virtual Directory as the identity store:

  1. Verify Oracle Virtual Directory is running.

     You can view the status of Oracle Virtual Directory using Fusion Middleware Control. After logging in to Fusion Middleware Control, navigate to the Farm home page and view the Identity and Access components within the Fusion Middleware section of the content pane.

     Refer to the "Displaying Fusion Middleware Control" section in the *Oracle Fusion Middleware Administrator's Guide* for more information.

  2. Verify the connectivity between Oracle Virtual Directory and its back-end (source) data repositories. Use Oracle Directory Services Manager's Client View Data Browser to view the directory tree. If Oracle Virtual Directory is not connected to a back-end repository, a message will appear when the Data Browser attempts to connect it.

     Refer to the following sections (in the listed sequence) in the *Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory* for more information about using Oracle Directory Services Manager's Client View Data Browser:

     a. "Invoking Oracle Directory Services Manager"

     b. "Connecting to the Server from Oracle Directory Services Manager"

     c. "Viewing Oracle Virtual Directory Entries"

22.3.5 Problems and Solutions for Managing Keystores and Certificates

This section describes problems and solutions for managing keystores and certificates. This section contains the following topics:

- Key or Credential Store Error After an Application Invokes Web Service
- Trust Certificate Error After Application Invokes Web Service

22.3.5.1 Key or Credential Store Error After an Application Invokes Web Service

After an Oracle Fusion application invokes a web service, a key store or credential store error such as the following appears:

- `WSM-00056: The key orakey is not retrieved`
- `WSM-00256: The property "Keystore Sign Alias" is not set`

Problem

The problem may be:

- The alias for the signature key or encryption key in the Oracle WSM keystore configuration does not exist in the Oracle WSM keystore file.

- The signature key, encryption key, or Oracle WSM keystore file password is not synchronized in the keystore file and the keystore configuration for Oracle WSM. That is, at least one of the passwords does not have identical values in both locations.

Solution 1

To verify the alias for the signature key and encryption key in the Oracle WSM keystore configuration exist in the Oracle WSM keystore file:

1. Use Fusion Middleware Control to identify the alias for the signature key and encryption key in the Oracle WSM keystore configuration. Perform the procedure in the "Configuring Keystores for Message Protection" section in the *Oracle Fusion Middleware Security and Administrator's Guide for Web Services*.

2. Verify the aliases identified in Step 1 exist in the Oracle WSM keystore file. Use the `keytool -list` command on the Oracle WSM keystore file to view its aliases. Refer to the *keytool - Key and Certificate Management Tool* document on the Java SE Technical Documentation Web site for more information about using `keytool`. You can access this document by searching for it on the Search Java SE Technical Documentation Web page at:

   `http://download.oracle.com/javase/search.html`

   - Ensure each alias is synchronized in both locations. If they are not, you can edit the alias in the Oracle WSM keystore configuration by performing the procedure in the "Configuring Keystores for Message Protection" section in the *Oracle Fusion Middleware Security and Administrator's Guide for Web Services*. You can edit the alias in the Oracle WSM keystore file using the `keytool -changealias` command.

     Note: Before you edit an alias, be sure that doing so will not affect any other web service.

   - If the alias for the signature key or encryption key does not exist in the Oracle WSM keystore file, add it by referring to the "Generating Private Keys and

Creating the Java Keystore" section in the *Oracle Fusion Middleware Security and Administrator's Guide for Web Services*.

> **Note:** If you make any changes to the Oracle WSM keystore file using `keytool`, you must restart all Managed Servers where Oracle WSM is deployed to activate the changes.

Solution 2

To ensure that the signature key, encryption key, and Oracle WSM keystore file passwords are each synchronized in the keystore file and the keystore configuration for Oracle WSM:

1. Use `keytool` to reset the passwords in the Oracle WSM keystore file. Because the passwords are not visible, resetting them is the only method to ensure that they have identical respective values in both locations.

   - Use the `keytool -storepasswd` command to reset the Oracle WSM keystore file password.

   - Use the `keytool -keypasswd` command to reset the signature key password and encryption key password.

   > **Note:** After resetting passwords in the Oracle WSM keystore file using `keytool`, you must restart all Managed Servers where Oracle WSM is deployed to activate the changes.

2. Use Fusion Middleware Control to reset the passwords in the Oracle WSM keystore configuration to the same respective values you set in Step 1. Refer to the "Configuring Keystores for Message Protection" section in the *Oracle Fusion Middleware Security and Administrator's Guide for Web Services* for more information.

22.3.5.2 Trust Certificate Error After Application Invokes Web Service

After an Oracle Fusion application invokes a web service, a trust certificate error such as the following appears:

```
WSM-00138: The path to the certificate is invalid due to
exception
```

Problem

The problem may be, if the web service is advertising its certificate in the Web Services Description Language (WSDL), the client is not configured correctly to trust that certificate or its issuer.

Solution

To verify the client is configured to trust the web service's certificate advertised in the WSDL or its issuer:

1. Verify the client keystore has either the public certificate of the web service or the public certificate of its issuer. Use the `keytool -list` command to identify the certificates in the client keystore. If either of the public certificates are missing from the client keystore, use the `keytool -importcert` command to add them.

   Refer to the *keytool - Key and Certificate Management Tool* document on the Java SE Technical Documentation Web site for more information about using `keytool`.

You can access this document by searching for it on the Search Java SE Technical Documentation Web page at:

http://download.oracle.com/javase/search.html

2. Verify the value for the keystore.recipient.alias override property of the client Oracle WSM policy is identical to the alias of the trusted public certificate in the Oracle WSM keystore file. Refer to the "Attaching Web Service Policies Permitting Overrides" section of the *Oracle Fusion Middleware Security and Administrator's Guide for Web Services* for more information.

22.3.6 Problems and Solutions for Identity Propagation Using SAML

After an Oracle Fusion application attempts to propagate a user's identity by calling a different Oracle Fusion application using Oracle SOA, `InvalidSecurityToken-`, `FailedAuthentication-`, or `SAML assertion issuer`-related errors appear.

Problem

The problem may be:

- The SAML issuer name for the SAML token is not configured or is configured incorrectly.

- The subject.precendence configuration override is set incorrectly.

- A user is not logged in on the client. The following error is a symptom of this problem:

  ```
  WSM-00263: Failed to create SAML token as 'anonymous' user
  principal found in Subject
  ```

Solution 1

To troubleshoot the SAML issuer name configuration, verify the SAML Issuer Name the client is using is among the issuers configured on the Oracle WebLogic Server domain by performing the steps in the "Adding an Additional SAML Assertion Issuer Name" section of the *Oracle Fusion Middleware Security and Administrator's Guide for Web Services*.

If the SAML Issuer Name that the client is using is not configured as an issuer in the Oracle WebLogic Server domain, Oracle recommends changing the issuer name on the client by updating its saml.issuer.name override to one of the issuers configured on the Oracle WebLogic Server domain.

If you cannot change the issuer name on the client, you can add its issuer name to the Oracle WebLogic Server domain by performing the steps in the "Adding an Additional SAML Assertion Issuer Name" section of the *Oracle Fusion Middleware Security and Administrator's Guide for Web Services*.

> **Note:** If you make any changes to the issuers configured in the Oracle WebLogic Server domain, you must restart the Managed Server where Oracle WSM is deployed.

Solution 2

To troubleshoot the subject.precendence configuration override:

1. Set the subject.precendence override value in your current client policy to false to change the identity to a different user. By default, the subject.precendence override is set to true.

2. Set the appropriate Credential Store Framework key override on the client policy that contains the user name and password of the user you want to send to the service. If an entry for this user does not exist in the Credential Store Framework, you must add it. Refer to the "Adding Keys and User Credentials to the Credential Store" section in the *Oracle Fusion Middleware Security and Administrator's Guide for Web Services* for more information.

3. Ensure the appropriate Web Services Identity Permission is set for the client application by performing the steps in the "Configuring SAML Web Service Clients for Identity Switching" section of the *Oracle Fusion Middleware Security and Administrator's Guide for Web Services*.

Solution 3

If you encountered the `WSM-00263: Failed to create SAML token as 'anonymous' user principal found in Subject` error, ensure the client has been authenticated before invoking the web service. You can ensure portlet and Oracle Application Development Framework application clients have been authenticated by logging out of them and logging back in.

22.3.7 Problems and Solutions for Logging in to Secured Resources

This section describes problems and solutions for logging in to secured resources. This section contains the following topics:

- Incorrect Language Appears After Logging in to a Secured Resource
- Login Page Unexpectedly Reappears (No Single Sign-On)
- Cannot Access Forgotten Password Functionality from Login Page

22.3.7.1 Incorrect Language Appears After Logging in to a Secured Resource

While attempting to access a resource secured by Oracle Access Manager, a user changes the language preference on the login page. After logging in successfully, the secured resource appears in a language different from what the user selected on the login page.

Problem

The problem may be Oracle Access Manager's `ORA_FUSION_PREFS` cookie is not in the user's browser. The `ORA_FUSION_PREFS` cookie determines which language the secured resource appears in. After the user chooses a language preference on the login page and gets authenticated, Oracle Access Manager sends the `ORA_FUSION_PREFS` cookie to the user's browser.

Solution

Examine the cookies in the user's browser and try to locate the `ORA_FUSION_PREFS` cookie. If the `ORA_FUSION_PREFS` cookie is not in the user's browser:

- Examine the browser's security settings, as they may be too high and preventing cookies from being accepted.
- Add the Oracle Fusion application's domain to the browser's exception list.

22.3.7.2 Login Page Unexpectedly Reappears (No Single Sign-On)

After successfully logging in to a resource secured by Oracle Access Manager, a login page unexpectedly reappears. Regardless if the reappearing login page is for Oracle

Access Manager or Oracle Fusion Applications, a user may not expect to see it in a single sign-on environment.

Problem

The problem may be:

- If the login page reappeared after attempting to access a different secured resource, the authentication level of the authentication scheme securing the subsequently accessed resource is greater (higher) than the authentication level of the authentication scheme securing the resource that was accessed first. In this situation, the reappearing login page is expected behavior.

- The Oracle Access Manager server's Idle Timeout or Session Lifetime configuration parameters are set to a value that is too small. The Idle Timeout parameter specifies the amount of time, in minutes, that a user's authentication session remains valid without accessing a resource secured by Oracle Access Manager. The Session Lifetime parameter specifies the amount of time, in minutes, that a user's authentication session remains valid. For both parameters, the smaller the value, the more frequently users must re-authenticate.

- Oracle Access Manager's `ObSSOCookie` and `OAM_ID` cookies are not in the user's browser. The `ObSSOCookie` and `OAM_ID` cookies are encrypted, single sign-on, session-based cookies generated by the Oracle Access Manager Access Server when a user authenticates successfully.

Solution 1

To examine the authentication levels of the authentication schemes securing the resources:

1. Log in to the Oracle Access Manager Administration Console by referring to the "Logging In to the Oracle Access Manager 11*g* Administration Console" section in the *Oracle Fusion Middleware Administrator's Guide for Oracle Access Manager with Oracle Security Token Service*.

2. Identify the authentication policies securing the resources and the authentication schemes configured for those policies. You can reduce the number of policies to examine by first looking at the policies for the Host Identifier that the Webgate is using.

   Refer to the "Searching for an Authentication Policy" and "Viewing or Editing an Authentication Policy" sections in the *Oracle Fusion Middleware Administrator's Guide for Oracle Access Manager with Oracle Security Token Service* for more information.

3. Identify the authentication levels for each authentication scheme. Refer to the "Viewing or Editing a Authentication Scheme" section in the *Oracle Fusion Middleware Administrator's Guide for Oracle Access Manager with Oracle Security Token Service*.

   > **Note:** If the authentication level for the subsequently accessed resource is greater than that of the first accessed resource, the reappearing login page is the expected behavior.

Solution 2

To verify the settings for the Idle Timeout and Session Lifetime configuration parameters:

1. Log in to the Oracle Access Manager Administration Console by referring to the "Logging In to the Oracle Access Manager 11*g* Administration Console" section in the *Oracle Fusion Middleware Administrator's Guide for Oracle Access Manager with Oracle Security Token Service*.

2. Verify the values configured for the **Idle Timeout** and **Session Lifetime** configuration parameters by referring to the "Configuring User Session Lifecycle Settings" section in the *Oracle Fusion Middleware Administrator's Guide for Oracle Access Manager with Oracle Security Token Service*.

Solution 3

To verify Oracle Access Manager's `ObSSOCookie` and `OAM_ID` cookies are in the user's browser:

1. Display the cookies in the user's browser.

2. Locate Oracle Access Manager's `ObSSOCookie` and `OAM_ID` session cookies.

If the `ObSSOCookie` and `OAM_ID` cookies are not in the user's browser:

- Examine the browser's security settings, as they may be too high and preventing cookies from being accepted.

- Add the Oracle Fusion application's domain to the browser's exception list.

22.3.7.3 Cannot Access Forgotten Password Functionality from Login Page

While attempting to access a resource secured by Oracle Access Manager, the Forgotten Password feature is inaccessible from the login page.

Problem

The problem may be:

- Network issues are preventing a connection to Oracle Identity Manager.

- Oracle Access Manager's configuration to Oracle Identity Manager's lost password functionality is incorrect.

Solution 1

To test connectivity to Oracle Identity Manager, from the system hosting the Administration Server where Oracle Access Manager is deployed, `ping` the system hosting the Managed Server where Oracle Identity Manager is deployed.

Solution 2

To verify Oracle Access Manager's configuration to Oracle Identity Manager's lost password functionality:

1. Use a text editor to open the following file on the Administration Server for the domain where Oracle Access Manager is deployed:

   DOMAIN_HOME/config/fmwconfig/oam-config.xml

2. Locate the `<Setting Name="IdentityManagement" Type="htf:map">` entry.

3. Examine the `ServerConfiguration` settings similar to those shown in Example 22–4 and verify the following values:

Example 22–4 ServerConfiguration Settings Within IdentityManagement Entry

```
<Setting Name="ServerConfiguration" Type="htf:map">
  <Setting Name="OIM-SERVER-1" Type="htf:map">
    <Setting Name="Host" Type="xsd:string">OIM_HOST</Setting>
    <Setting Name="Port" Type="xsd:integer">OIM_PORT</Setting>
    <Setting Name="SecureMode" Type="xsd:boolean">true|false</Setting>
  </Setting>
</Setting>
```

- **OIM-SERVER-1**: Must be identical value of the same setting in the `IdentityServiceProviderConfiguration` entry described in Step 4.

- If Oracle Identity Manager is front-ended by Oracle HTTP Server or a load balancer:

  - **OIM_HOST**: Fully-qualified host name of Oracle HTTP Server or load balancer.

  - **OIM_PORT**: The port for the Oracle HTTP Server or load balancer.

  - **SecureMode**: Set to true for connecting to Oracle Identity Manager over HTTPS, set to false for connecting over HTTP.

- If Oracle Identity Manager is not front-ended:

  - **OIM_HOST**: Fully-qualified host name of the Managed Server where Oracle Identity Manager is deployed.

  - **OIM_PORT**: The port for the Managed Server where Oracle Identity Manager is deployed.

  - **SecureMode**: Set to true for connecting to Oracle Identity Manager over HTTPS, set to false for connecting over HTTP.

4. Examine the `IdentityServiceProviderConfiguration` settings similar to those shown in Example 22–5 and verify the following values:

Example 22–5 IdentityServiceProviderConfiguration Settings Within IdentityManagement Entry

```
<Setting Name="IdentityServiceProviderConfiguration" Type="htf:map">
  <Setting Name="IdentityManagementServer" Type="xsd:string">OIM-SERVER-1</Setting>
  <Setting Name="DateFormatPattern" Type="xsd:string">yyyy-MM-dd'T'HH:mm:ss'Z'</Setting>
  <Setting Name="PasswordExpiredURL" Type="xsd:string">/admin/faces/pages/pwdmgmt.jspx</Setting>
  <Setting Name="ChallengeSetupNotDoneURL" Type="xsd:string">/admin/faces/pages/pwdmgmt.jspx</Setting>
  <Setting Name="ForcedPasswordChangeURL" Type="xsd:string">/admin/faces/pages/pwdmgmt.jspx</Setting>
  <Setting Name="AccountLockedURL" Type="xsd:string">/admin/faces/pages/accountlocked.jspx</Setting>
</Setting>
```

- **OIM-SERVER-1**: Must be identical value of the same setting in the `ServerConfiguration` entry described in Step 3.

- Confirm the following URL Settings are configured with the values shown in Example 22–5:

  - PasswordExpiredURL

  - ChallengeSetupNotDoneURL

  - ForcedPasswordChangeURL

  - AccountLockedURL

22.4 Additional Information for Troubleshooting Oracle Identity Management

The following is a list of Oracle Identity Management documents that provide additional information and will help you troubleshoot. Use these documents if you have isolated your problem to a specific Oracle Identity Management component or to learn more about a specific component.

> **Note:** A few of the documents in the following list do not contain explicit troubleshooting information, but are a source of information that will help you during troubleshooting.

- "Troubleshooting Oracle Fusion Middleware" appendix of the *Oracle Fusion Middleware Administrator's Guide*

- "Troubleshooting Security in Oracle Fusion Middleware" appendix in the *Oracle Fusion Middleware Application Security Guide*

- *Oracle Fusion Middleware Enterprise Deployment Guide for Oracle Identity Management (Oracle Fusion Applications Edition)*

- "Troubleshooting Oracle Internet Directory" appendix of the *Oracle Fusion Middleware Administrator's Guide for Oracle Internet Directory*

- "Troubleshooting Oracle Virtual Directory" appendix of the *Oracle Fusion Middleware Administrator's Guide for Oracle Virtual Directory*

- "Troubleshooting Oracle Authorization Policy Manager" appendix in the *Oracle Fusion Middleware Oracle Authorization Policy Manager Administrator's Guide (Oracle Fusion Applications Edition)*

- "Troubleshooting" appendix in the *Oracle Fusion Middleware Administrator's Guide for Oracle Access Manager with Oracle Security Token Service*

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

- *Oracle Fusion Middleware Security and Administrator's Guide for Web Services*

22.5 Using My Oracle Support for Additional Troubleshooting Information

You can use My Oracle Support (formerly MetaLink) to help resolve Oracle Fusion Applications problems. My Oracle Support contains several useful troubleshooting resources, such as:

- Knowledge Base articles

- Community Forums and Discussions

- Patches and Upgrades

- Certification information

> **Note:** You can also use My Oracle Support to log a service request.

You can access My Oracle Support at `https://support.oracle.com`.

23

Troubleshooting Oracle WebCenter Portlets

This chapter describes common problems that you might encounter when using Oracle WebCenter portlets and explains how to solve them.

This chapter contains the following topics:

- Introduction to Troubleshooting Oracle WebCenter Portlets
- Problems and Solutions
- Diagnosing Oracle WebCenter Portlet Problems
- Using My Oracle Support for Additional Troubleshooting Information

This chapter focuses on troubleshooting Oracle WebCenter portlets. For information about troubleshooting general Oracle WebCenter configuration issues, see the "Troubleshooting WebCenter Application Configuration Issues" section in the *Oracle Fusion Middleware Administrator's Guide for Oracle WebCenter*. Also, review the *Oracle Fusion Middleware Error Messages Reference* for information about the error messages you may encounter.

23.1 Introduction to Troubleshooting Oracle WebCenter Portlets

Oracle Fusion Applications utilizes portlet technology in various places, typically to remotely invoke a business view that is implemented as an Oracle Application Development Framework (ADF) task flow. This enables functionality that is implemented on one Oracle Fusion application, for example, HCM, to be incorporated into another Oracle Fusion application, for example, CRM, as if it was embedded in the same application. The CRM application is in fact embedding a portlet that obtains its markup from the remote HCM application, which is running on another server. The task flow that is implemented on the HCM server is made available as a portlet through a component called the Oracle JSF Portlet Bridge. This wrapper makes the task flow available as a portlet producer that can be consumed by another application.

Useful Terminology

The following list defines some common terms for Oracle WebCenter portlets:

- Portlet

  A portlet is a region of the screen that is displayed from a remote source. In Oracle Fusion Applications, portlets conform to the Web Services for Remote Portlets (WSRP) standard, and implement the JSR 286 portlet specification.

- Oracle JSF Portlet Bridge

The Oracle JSF Portlet Bridge is a component that enables an ADF application to be exposed as a WSRP portlet producer application. Oracle Fusion applications are all implemented as ADF applications.

- Producer application

  A producer application is an ADF application with pages or task flows that have been enabled to run as portlets. This type of application can run in dual modes: as a servlet (like a regular web application), or as a portlet (when consumed by a consumer application through the Oracle JSF Portlet Bridge).

- Consumer application

  A consumer application is an application that consumes the portlets exposed by a producer application. Before consuming a portlet (by dropping it onto a page), application developers must first register the portlet producer application with the consumer application.

Useful Resources

The following list provides some useful resources to use when diagnosing problems with Oracle WebCenter portlets:

- Portlet Consumer Test Page

  A page that provides diagnostic information about the consumer application. You can access the Portlet Consumer Test Page using the following URL:

  ```
  http://host:port/context-root/faces/oracle/portlet/client/adf/diagnostic/pages/
  ConsumerTestPage.jspx
  ```

  where:

  - *host* is the server to which the consumer application is deployed
  - *port* is the port to which the server is listening for HTTP requests
  - *context-root* is the consumer web application's context root

  For example:

  ```
  http://mymanagedserver.example.com:8888/myapp/faces/oracle/portlet/client/adf/d
  iagnostic/pages/ConsumerTestPage.jspx
  ```

  For more information, see Section 23.3.1.2.

- Producer Test Page

  A page that provides diagnostic information about the portlet producer application. You can access the Producer Test Page using the following URL:

  ```
  http://host:port/context-root/info
  ```

  where:

  - *host* is the server to which the portlet producer is deployed
  - *port* is the port to which the server is listening for HTTP requests
  - *context-root* is the producer web application's context root

  For example:

  ```
  http://portlets.example.com:9999/sample/info
  ```

  The Producer Test Page includes a link to the Web Service Definition Language (WSDL) document to use for registration, for example:

```
http://portlets.example.com:9999/sample/portlets/wsrp2?WSDL
```

For more information, see Section 23.3.1.3.

- Running a producer as a servlet application through Faces

  This is also known as running the application as a servlet. Before an application can act as a portlet provider, it must be able to run correctly through standard HTTP requests.

  To run an application as a servlet, use the following URL:

  ```
  http://host:port/context-root/faces/path-to-page/page.jspx
  ```

  where:

  - *host* is the server to which the portlet producer is deployed
  - *port* is the port to which the server is listening for HTTP requests
  - *context-root* is the producer web application's context root
  - *path-to-page* is the path to the page you want to run
  - *page* is the name of the page you want to run

  For example:

  ```
  http://portlets.example:9999/sample/faces/index.jspx
  ```

  The Producer Test Page provides links to run such pages or task flows as servlets. For more information, see Task 2, "Run the JSF Portlet as a Servlet".

- Logging configuration file

  The logging configuration file, logging.xml, is located in:

  ```
  DOMAIN_HOME/config/fmwconfig/servers/server/logging.xml
  ```

- Diagnostic log file

  The default location of the diagnostic log file is:

  ```
  DOMAIN_HOME/servers/server/logs/server-diagnostic.log
  ```

Process

Follow the process outlined in Table 23–1 when using the information in this chapter. If the information in a particular section does not resolve your problem, proceed to the next step in this process.

Table 23–1 Process for Using the Information in this Chapter

Step	Section to Use	Purpose
1	Section 23.2	Perform problem-specific troubleshooting procedures. These section describes: - Possible causes of the problems - Solution procedures corresponding to each of the possible causes
1	Section 23.3	Perform general diagnostics steps with when you encounter problems with Oracle WebCenter portlets.

Table 23–1   (Cont.) Process for Using the Information in this Chapter

Step	Section to Use	Purpose
2	Section 23.4	Use My Oracle Support to get additional troubleshooting information about Oracle Fusion Applications or Oracle WebCenter portlets. My Oracle Support provides access to several useful troubleshooting resources, including Knowledge Base articles and Community Forums and Discussions.
		When you encounter problems with Oracle WebCenter portlets, there are some general diagnostics steps that you can follow.
3	Section 23.4	Log a service request if the information in this chapter and My Oracle Support does not resolve your problem. You can log a service request using My Oracle Support at `https://support.oracle.com`.

23.2 Problems and Solutions

When running Oracle Fusion Applications, it may not be readily apparent which portions of the user interface are implemented as portlets. The only time this may actually be evident is when there is a problem.

This section describes common problems and solutions. It contains the following topics:

- Portlet Displays a Portlet Consumer Error
- Portlet Displays a Portlet Timeout
- Portlet Displays a Remote Portlet Communication Error
- Portlet Displays a Remote Portlet Error
- Portlet Displays a Server Connection Failed Dialog

23.2.1 Portlet Displays a Portlet Consumer Error

The message **Portlet Consumer Error** (shown in Figure 23–1) typically indicates that an error occurred within the operation of the portlet parts of the portlet consumer application. Occasionally, the error may indicate that the remote portlet producer is refusing connection.

Figure 23–1   Portlet Displaying a Portlet Consumer Error
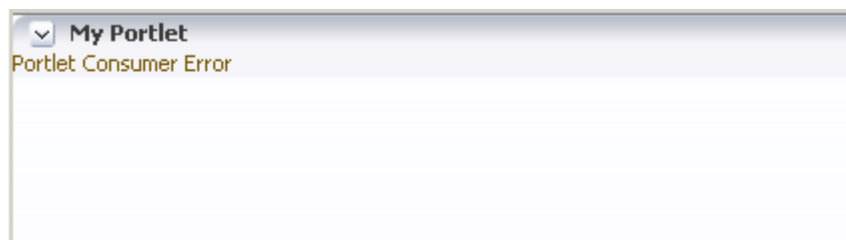
Problem 1

An error has occurred within the operation of the portlet parts of the portlet consumer application. In other words, the error is unrelated to the remote portlet producer application.

Solution 1

Consult the diagnostic log file to determine the cause of the exception.

If the `DebugErrorRenderer` is enabled, the cause exception is displayed in the portlet along with links to the log file. If running in production mode, then consult the consumer-side logs.

The exception that caused the error message to be displayed is logged. Wherever possible, a message is included in the log at the start of the exception stack to indicate for which portlet binding the exception occurred. Example 23–1 shows a message logged for a portlet error.

Example 23–1   Example Message Logged for a Portlet Error

```
<PortletRenderer> <setErrorState> An error has occured for Portlet Binding
portlet1.
oracle.portlet.client.container.PortletContentTypeException: Unexpected content
type "null" in WSRPGetMarkup response.
...
```

Pay particular attention to the cause exceptions in the stack as this is likely to indicate what the real underlying problem is.

The cause is likely to be an internal error and the appropriate course of action is to contact Oracle Support.

Problem 2

The remote portlet producer application server is refusing connections completely. This usually occurs when the server is not running. This situation is characterized by the an exception stack on the consumer-side logs similar to that shown in Example 23–2.

Example 23–2   Exception Stack When Portlet Producer Refuses Connection

```
oracle.portlet.client.container.PortletContentTypeException: Unexpected content
type "null" in WSRPGetMarkup response.
 at oracle.portlet.client.techimpl.wsrp.WSRPBaseTerminalPipe.processException
  (WSRPBaseTerminalPipe.java:109)
 at oracle.portlet.client.techimpl.wsrp.WSRPGetMarkupPipe.execute
  (WSRPGetMarkupPipe.java:248)
...
Caused by: HTTP transport error: java.net.ConnectException: Connection refused
 at oracle.portlet.client.connection.wsrp.HTTPClientTransport.invokeOneWayInternal
  (HTTPClientTransport.java:546)
 at oracle.portlet.client.connection.wsrp.HTTPClientTransport.invoke
  (HTTPClientTransport.java:197)
 at oracle.j2ee.ws.client.StreamingSender._sendImpl(StreamingSender.java:232)
 at oracle.j2ee.ws.client.StreamingSender._send(StreamingSender.java:148)
 at oracle.portlet.wsrp.v2.soap.runtime.WSRP_v2_Markup_Binding_SOAP_
Stub.getMarkup(WSRP_v2_Markup_Binding_SOAP_Stub.java:80)
```

Solution 2

If the cause exception indicates that the remote producer is down, then the resolution is to bring the producer up. You need an understanding of which Oracle Fusion applications are associated with which portlet producers in order to determine which applications need to be started up.

23.2.2 Portlet Displays a Portlet Timeout

If a **Portlet Timeout** is displayed in the area of the page that you would expect to contain a portlet (as shown in Figure 23–2), this means that the consumer waited for a

configured period of time for the producer to respond and did not get a response during that time, or the response did not complete during that time. There are a number of possible causes.

Figure 23–2   Portlet Displaying a Portlet Timeout Error

Problem 1

The producer machine is overloaded.

Solution 1

Check the load on the producer managed server (the tools used to do this vary depending on the operating system that is running on the producer). If the load is high, check whether a particular process is causing this high load, and whether such a process could be run on another machine, or at a less busy time. If no single process is causing the high load, or if the Oracle WebLogic Server is causing the high load, and if the load is consistently high, consider whether the producer hardware is adequate, or whether it is necessary to upgrade it (or add further nodes to the cluster). Also consider adjusting the Oracle WebLogic Server configuration to increase the size of the request thread pool. For more information, see the Oracle WebLogic Server documentation.

Problem 2

The network is overloaded, or there are problems with the network affecting communication between the consumer and producer.

Solution 2

Check that you can ping the producer machine from the consumer machine. Check that you can access the producer's WSRP Producer Test Page in your local browser. If this works, check that you can access this same page from a browser running on the consumer machine. If any of these steps cause problems, and the machine is not overloaded, this could be a network problem, which should be investigated by a system administrator.

Problem 3

There is a deadlock (or a stuck thread) on the producer machine causing the request thread to hang.

Solution 3

This should not happen during normal operation. If it does occur, there will generally be an error in the producer's log files indicating the point at which the deadlock occurred. This may help diagnose the problem. In some cases, it may be possible to alleviate this by modifying the configuration of Oracle WebLogic Server. For more information, see the Oracle WebLogic Server documentation.

Problem 4

The producer application is running slowly (for example, due to processing large quantities of data).

Solution 4

In this case, the producer application may be processing large quantities of data, causing it to spend too long building the response. If the application will regularly deal with large quantities of information, it may be necessary to either add or improve producer hardware, or to increase the portlet timeout duration. The portlet timeout can be configured on the producer connection in the consumer application using Enterprise Manager or the WLST `setWSRPProducerConnection` command. Additionally, minimum and maximum timeouts for all producer connections within the application can be configured within the portlet section of the `adf-config.xml` file.

Problem 5

The producer application is waiting for a response from another resource, such as a database, that is taking too long (for example, if the database is overloaded).

Solution 5

Check that the resource in question is functioning correctly. If it is, the solution the same as Solution 4.

Problem 6

The producer application is being debugged and has hit a breakpoint.

Solution 6

This only happens if the producer application Oracle WebLogic Server is running in debug mode, and a debugger has been connected. In this case, if an incoming request causes a breakpoint to be triggered, the execution of the producer application stops and the consumer times out when the portlet timeout interval is reached. This scenario is unlikely to occur in a production environment.

Problem 7

The portlet timeout values in `adf-config.xml` have been misconfigured such that the timeout period is too short.

Solution 7

The portlet section of the `adf-config.xml` file allows minimum, maximum, and default values for portlet timeouts to be configured across the whole application. The maximum timeout imposes an upper limit on timeouts specified by portlet producers, so if the maximum timeout is too short, this could cause unwanted portlet timeout errors even if the timeout specified on the producer connection is longer.

23.2.3 Portlet Displays a Remote Portlet Communication Error

When a section of the screen shows the **Remote Portlet Communication Error** message (as shown in Figure 23–3), and there is an otherwise blank region surrounding it, this area is expected to be filled with a portlet, which the application is not able to contact.

Figure 23–3   Portlet Displaying a Remote Portlet Communication Error
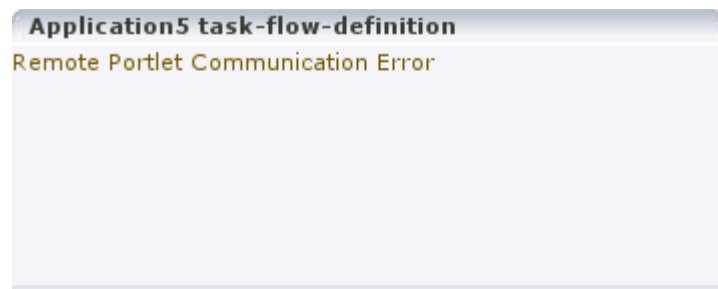
Problem 1

The producer is down.

Solution 1

It could be that the producer application is not running, or the managed server on which it is deployed is not started. In this case, it will need to be started. Identify the application that needs to be started based on the task being attempted at the time of the portlet failure.

Problem 2

The web services security is incorrectly configured.

Solution 2

In Oracle Fusion Applications deployments, web services security (WS-Security) is managed with global web services security policies.

Troubleshooting steps for web services security depend on the type of security profile being used, for example AuthN, SSL, or Message Security.

For more information about troubleshooting web service security, see:

- The "Diagnosing Problems" chapter in the *Oracle Fusion Middleware Security and Administrator's Guide for Web Services*

- Section 22.3.5

- Section 22.3.6

For information about the different types of security profile, see the "Locking Down Web Services: Points to Consider" section in the "Hardening Backchannel Network and Services" chapter of the *Oracle Fusion Applications Security Hardening Guide*.

The security policies set on the portlet producer's `WSRP_v2_Service` web service ports are as follows:

- `WSRP_v2_ServiceDescription_Service` port: `oracle/no_authentication_service_policy`

- `WSRP_v2_PortletManagement_Service` port: `oracle/no_authentication_service_policy`

- `WSRP_v2_Markup_Service` port: no policy specified, so that it picks up the globally attached policy

- `WSRP_v2_Registration_Service` port: `oracle/no_authentication_service_policy`

If the producer ports are configured in any other way, then it may be the cause of the problem. In particular, if a local policy is applied to the `WSRP_v2_Markup_Service` port, and the policy does not match the corresponding policy on the producer connection, then the port or the connection will need to be updated to specify matching policies, or be removed, so that the globally attached policies can take effect.

Problem 3

The producer managed server cannot be reached.

Solution 3

The producer may be in a location that cannot be reached by the consumer application, due to intervening firewalls or incorrect routing rules. In an environment that is installed by Oracle's provisioning software, this should not be the case, but it is worth checking that you are able to access the WSDL endpoint for the producer from the machine hosting the consumer, by going to:

```
http://host:port/context-root/portlets/wsrp2?WSDL
```

Where:

- *host* is the server to which the portlet producer is deployed

- *port* is the port to which the server is listening for HTTP requests

- *context-root* is the producer web application's context root

For example:

```
http://portlets.example.com:9999/sample/portlets/wsrp2?WSDL
```

23.2.4 Portlet Displays a Remote Portlet Error

If the portlet displays a **Remote Portlet Error** (as shown in Figure 23–4), this indicates that the producer responded with an error message. The error message is returned in the form of a SOAP fault message inside the response document. There are a number of reasons the producer might return an error. The best strategy to diagnose these errors is to first find the corresponding exception stack trace in the consumer diagnostic logs. This stack trace shows what kind of fault was returned by the producer, plus any further information required in the response. Some faults you may encounter are listed in the following sections.

Figure 23–4   Portlet Displaying a Remote Portlet Error
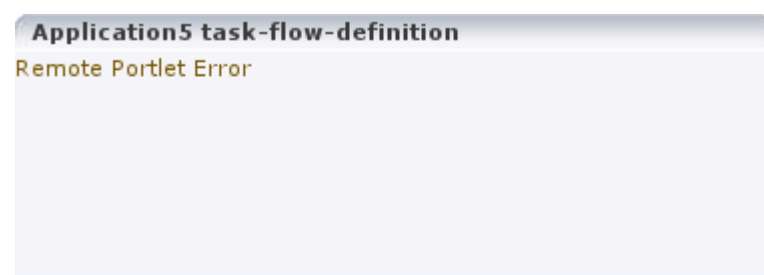
Problem 1

`OperationFailedException`. This is the most common type of Remote Portlet Error and it is a catch-all for most unhandled exceptions raised in the producer application.

Solution 1

To resolve an `OperationFailedException`, examine the exception in the consumer diagnostic logs. This generally shows any exception that was raised in the producer application to trigger the fault response as the final `Caused by` exception.

If required, you can then examine the diagnostic logs on the producer application for more detail, or for any related exceptions that occurred prior to the fault being triggered. In some cases, the exception in the producer log indicates a problem that can be simply resolved, such as a database connection failure or configuration problem. In other cases, the exception might indicate a product bug.

Problem 2

`InvalidRegistrationException`. This error indicates that the producer has not been properly registered with the consumer before the consumer attempted to communicate with it. This could also occur if the producer's preference store has been moved or deleted since the consumer registered it.

Solution 2

This error should not occur in the Oracle Fusion Applications environment. If it is observed, the most likely cause is a problem during provisioning. It is also worth checking the producer application's `web.xml` file setting to ensure that the entry shown in Example 23–3 is present.

Example 23–3 Persistent Store Setting in web.xml

```
<env-entry>
  <env-entry-name>oracle/portal/wsrp/server/persistentStore</env-entry-name>
  <env-entry-type>java.lang.String</env-entry-type>
  <env-entry-value>Consumer</env-entry-value>
</env-entry>
```

Problem 3

`InvalidHandleException`. This indicates that the consumer has asked the producer to render, or otherwise interact with, a portlet instance that the producer does not know about. This could occur if the producer's preference store has been corrupted in some way since the portlet was added to the page.

Solution 3

This error should not occur in the Oracle Fusion Applications environment. It is most likely caused by a problem during provisioning, or a missing `persistentStore` setting in the `web.xml` file, as described in Solution 2.

Problem 4

`AccessDeniedException`. This indicates that the producer application decided that the current user did not have access to the portlet or task flow in question.

Solution 4

This error should not occur in the Oracle Fusion Applications environment. This could either be a legitimate error message or an indication of a configuration problem. In most cases, Oracle Fusion Applications should deal with authorization errors gracefully, without a Portlet Remote Error being displayed.

23.2.5 Portlet Displays a Server Connection Failed Dialog

If the portlet displays a dialog containing the message "A connection to the server has failed" (as shown in Figure 23–5), this indicates that the producer responded to a portlet resource proxy request with an error message.

Figure 23–5   The Server Connection Failed Dialog

Portlet resource proxy requests are made from the browser for any partial page render (PPR) request from the portlet. There are a number of reasons the producer might return an error.

The best strategy to diagnose these errors is first to find the corresponding exception stack trace in the consumer diagnostic logs. You can use the ECID in the error dialog to find the relevant exceptions in the log files. The stack trace shows what kind of fault was returned by the producer, plus any further information. You may also use other tools, such as Firebug, to look at the response of the request.

Problem

`PortletTimeoutException`. This exception means that the consumer has waited longer than the configured time out period for a response from the producer. This may happen if the system is under excessive load, or it may indicate problems in the producer application, for example, a slow database connection. Note that you will typically not see any errors logged in the producer application in this scenario, this is an indication that the producer is slow, not that anything is causing errors.

Solution

To resolve a `PortletTimeoutException`, use the Oracle Enterprise Manager Fusion Middleware Control Console or the Oracle WebLogic Server Administration Console console to look at the health of the producer server. If possible, render the same task flow on the producer server as a servlet, rather than as a portlet. If it is slow to respond check whether there are slow running database queries in the application.

23.3 Diagnosing Oracle WebCenter Portlet Problems

When you encounter problems with Oracle WebCenter portlets, there are some general diagnostics steps that you can follow.

This section contains the following topics:

- Using Diagnostic Tools
- Configuring the Portlet Logging File

23.3.1 Using Diagnostic Tools

There is a set of tools available for both the consumer and producer to help identify and resolve issues when running Oracle JSF Portlet Bridge portlets.

If you encounter a portlet error message when a portlet is rendered, or if the portlet displays but you cannot interact correctly with it, there are some general steps using these tools that you should follow to diagnose the issue.

This section contains the following topics:

- Identify the Portlet Instance
- Examine the Portlet Consumer Test Page
- Examine the Producer Test Page

23.3.1.1 Identify the Portlet Instance

The first step when you encounter a portlet error, is to identify which portlet producer and portlet instance is being invoked. Execute the `portletDebugShow()` JavaScript from your browser to display this information in the main portlet content area.

To identify the portlet instance:

1. Enter the following command in the Location field of your browser:

   ```
   javascript:portletDebugShow()
   ```

2. After running the script, every portlet now displays the following information:

   - Producer name
   - Portlet name
   - Portlet instance ID
   - Execution Context IDs (ECIDs)

     The ECIDs are unique IDs used to identify a portlet request. Use the ECIDs to correlate the messages across different consumer and producer log files using Fusion Middleware Control. The same ECID is propagated from the consumer to the producer. For more information, see the "Correlating Messages Across Log Files and Components" section in the *Oracle Fusion Middleware Administrator's Guide*.

     Note: Broken portlets show two ECIDs: one for the request in which the error occurred and one for request in which the error was reported. For inline portlets (that is, portlets that are not displayed in an IFRAME), these two ECIDs are the same.

     For IFRAME portlets, for example Oracle JSF Portlet Bridge portlets, the ECIDs are different. This is because the error is reported in a later request than the one in which the original exception occurred. When checking the logs, you should look for both ECIDs, as either may contain relevant information.
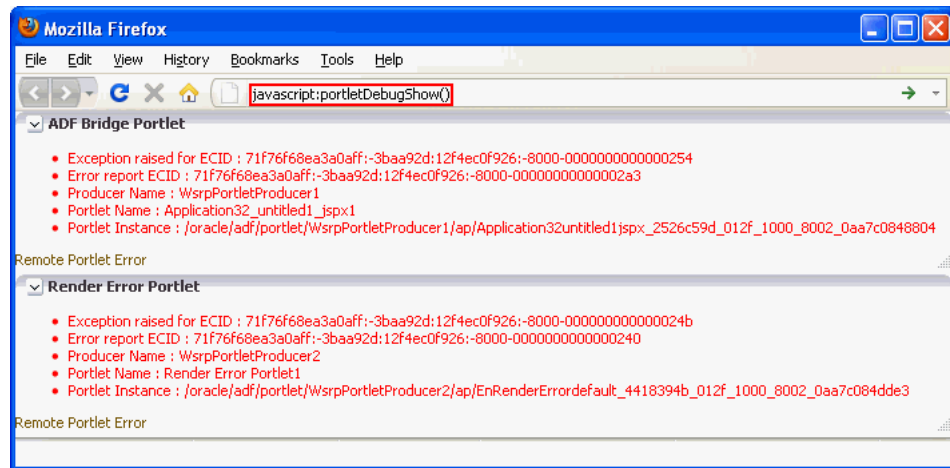
You can use this information in the subsequent diagnostic steps to help locate the issue.

> **Note:** The ECIDs shown in the portlet diagnostic information do not reflect partial page rendering requests that have been made to the portlet producer (using the portlet consumer resource proxy). These requests may update the portlet, but the ECIDs are not recorded in the portlet diagnostic information. Errors that occur during these requests are logged on the producer and by the portlet resource proxy on the consumer but you cannot use the ECID information reported in the portlet diagnostic information to help you determine the ECIDs for the relevant log entries.

3. When you have finished debugging the portlets, enter the following command to hide the portlet debugging information:

```
javascript:portletDebugHide()
```

23.3.1.2 Examine the Portlet Consumer Test Page

The next step in diagnosing a portlet error is to access the Portlet Consumer Test Page (shown in Figure 23–6) to locate the portlet producer and, if necessary, test the portlet in isolation.

Figure 23–6   The Portlet Consumer Test Page
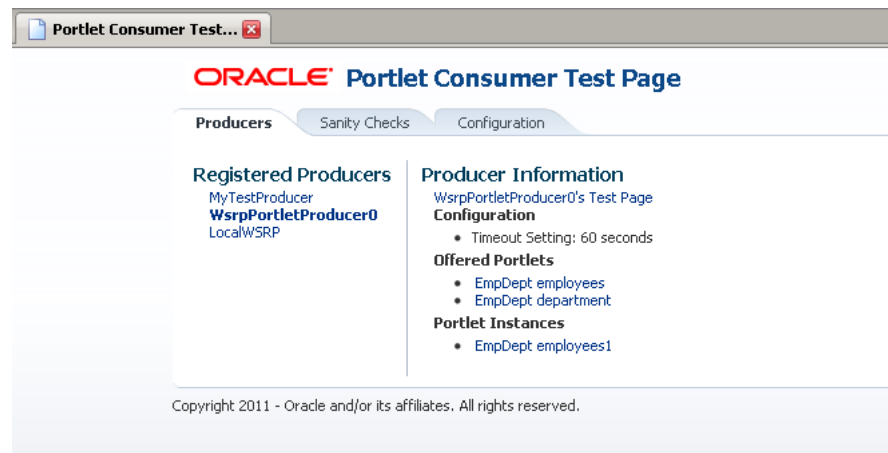
The Portlet Consumer Test Page contains three tabs:

- **Producers:** This tab lists all the producers registered with the consumer application. Selecting a producer provides specific information about that producer.

- **Sanity Checks:** This tab may contain a predefined set of portlet instances and required parameters that can be run in the consumer application, as configured by the consumer application developer. Any failures within these portlets indicate a problem with the corresponding producer and/or portlet.

- **Configuration:** This tab enables you to identify the consumer configuration entries for portlet consumption. You cannot change these values as they are stored within the application; they are displayed for reference information only.

After accessing the Portlet Consumer Test Page, you can perform further diagnostic steps.

This section contains the following topics for using the Portlet Consumer Test Page to diagnose portlet issues:

- Task 1, "Access the Portlet Consumer Test Page"

- Task 2, "Locate the Portlet Producer"

- Task 3, "Locate and Run the Portlet Instance"

- Task 4, "Perform Sanity Checks"

- Task 5, "Check Consumer Configuration Entries"

Task 1 Access the Portlet Consumer Test Page

The Portlet Consumer Test Page provides diagnostic information about the portlet consumer.

To access the Portlet Consumer Test Page:

1. In your browser, enter the URL for the Portlet Consumer Test Page:

   ```
   http://host:port/context-root/faces/oracle/portlet/client/adf/diagnostic/pages/
   ConsumerTestPage.jspx
   ```

> **Note:** If the consumer application is secured, the Portlet Consumer Test Page can be accessed only by users granted permission to view those pages.

2. In the Portlet Consumer Test Page, you can perform further diagnostic steps as described in the following sections.

Task 2 Locate the Portlet Producer

The Producers tab of the Portlet Consumer Test Page lists all the producers that have been registered with the consumer application. If a portlet instance in your application displays an error message, you can view information about the producer that owns the portlet by selecting it on this tab.

To locate the portlet producer:

1. In the Portlet Consumer Test Page, select the portlet producer that owns the portlet instance that is reporting the error.

   You noted this information in Section 23.3.1.1.

2. The following information is provided for the selected producer:

   – Producer Test Page: A link to the Producer Test Page.

   – Configuration: Details of potential issues surrounding skins, security, and timeouts associated with the using producer.

   – Offered Portlets: A list of all portlets offered by the producer. If there are no offered portlets listed, this indicates that there is a problem with the registration metadata for the producer.

   – Portlet Instances: A list of all portlet instances for the selected producer in the consumer application. This list may be empty.

   You can use this information to identify potential issues with the producer.

Task 3 Locate and Run the Portlet Instance

If you have still not been able to identify the cause of the portlet error, the issue may lie with the portlet instance itself.

To locate and run the portlet instance:

1. In the Portlet Consumer Test Page, select the portlet producer that owns the portlet instance that is reporting the error.

   You noted this information in Section 23.3.1.1.

2. Under Portlet Instances, select the portlet instance to display the Consumer Test Page: Portlet page.

   You noted this information in Section 23.3.1.1.
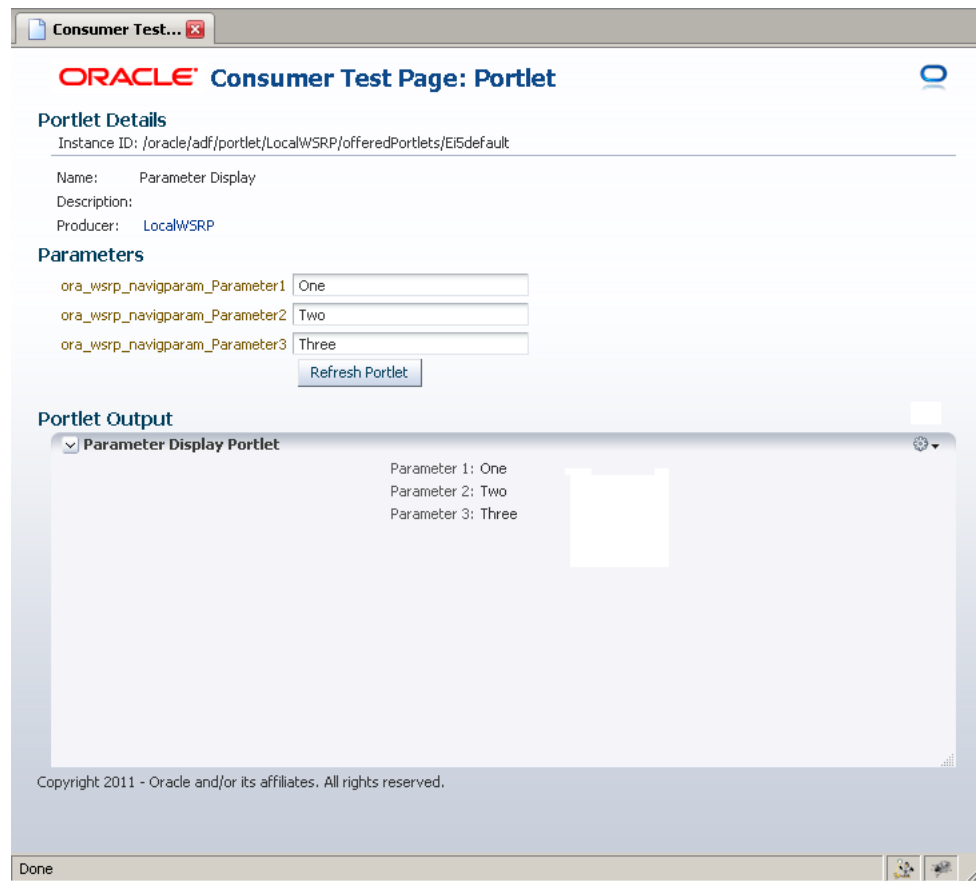
3. The Portlet Consumer Test Page: Portlet page renders the portlet in a standalone page. If the portlet runs correctly on this page, the problem is most likely caused by other components on the page containing the broken portlet.

4. If the portlet contains parameters, a parameter section is displayed that lists all the public parameters for the portlet. Enter values for any parameters to test that the portlet is receiving parameters correctly.

5. To navigate back to the Portlet Consumer Test Page, click the producer name link at the top of the page.

Task 4 Perform Sanity Checks

The Sanity Checks tab of the Portlet Consumer Test Page (shown in Figure 23–7) provides a quick overview of the state of portlet communication in your application across all producers.

Figure 23–7   The Sanity Checks Tab
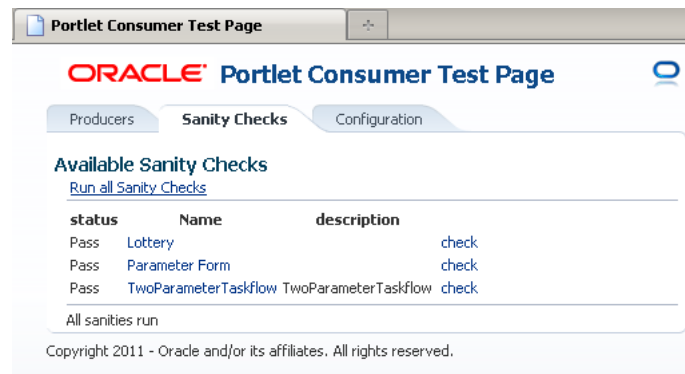
The Sanity Checks tab references portlet instances used within the consumer application. This list is configured by the application developer who chose the portlets to include and the parameters to pass to these portlets.

The checks on this page do not render the output in the UI, but simply create a runnable instance of the portlet under the covers and report any failures if any exception is returned by the portlet.

To perform sanity checks:

1. In the Portlet Consumer Test Page, click the **Sanity Checks** tab.

2. Click the **check** link next to the portlet that you want to test.

    The results of the sanity tests are displayed in the **Status** column.

3. To run sanity checks on all listed portlets, click the **Run all Sanity Checks** link.

Task 5   Check Consumer Configuration Entries

The Configuration tab of the Portlet Consumer Test Page (shown in Figure 23–8) enables you to identify the consumer configuration entries for portlet consumption. This tab displays settings defined in the `adf-config.xml` file, for example, the minimum and maximum timeout values and the consumer version number. You cannot change these values as they are stored within the application; they are displayed for reference information only.

Figure 23–8   The Configuration Tab
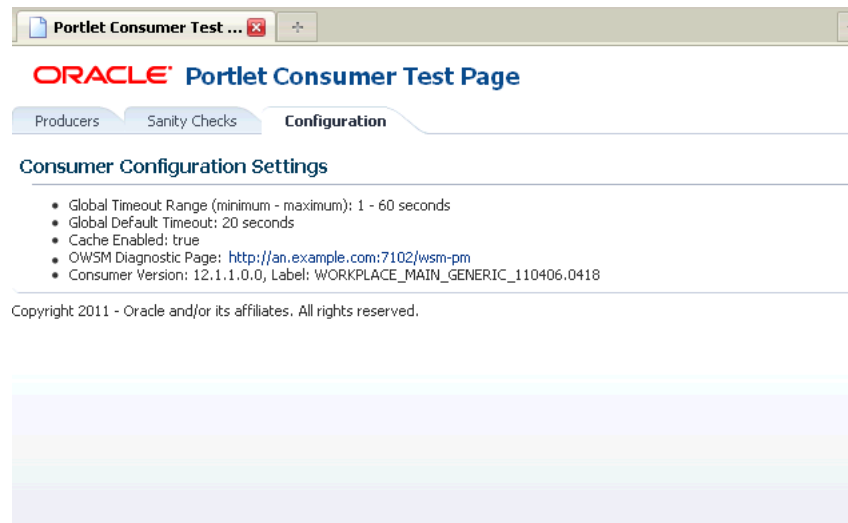
23.3.1.3 Examine the Producer Test Page

If you cannot identify the cause of the error in the consumer application, the next step is to use the Producer Test Page (shown in Figure 23–9) to identify potential issues with the portlet producer application.

Figure 23–9   The Producer Test Page

Access to the main Producer Test Page is public, but links to the test pages for each portlet are accessible only to users granted permission on the underlying pages and task flows.

The Producer Test Page contains five sections:

- Portlets

  A list of all the portlets within the producer. For Oracle JSF Portlet Bridge portlets, each portlet also provides a separate link to run the portlet as a servlet (this is a prerequisite to running them as portlets: if a portlet does not run as a servlet, it cannot run as a portlet).

- Container Configuration

  Information on where the consumer preference information is stored.

- Container Version

  The version number of the Portlet Producer Container.

- WSDL URLs

  Links to the Web Service Definition Language (WSDL) documents to use for registration.

- SOAP Monitor

  A link to the WSRP SOAP monitor where users with the `Monitors` or `Administrators` role can track the SOAP messages between the consumer and producer.

After accessing the Producer Test Page, you can perform further diagnostic steps.

This section contains the following topics:

- Task 1, "Access the Producer Test Page"
- Task 2, "Run the JSF Portlet as a Servlet"
- Task 3, "Check the Persistent Store Type"
- Task 4, "Examine the SOAP Monitor"

Task 1 Access the Producer Test Page

The Producer Test Page provides diagnostic information about the portlet producer.

To access the Producer Test Page:

1. In your browser, enter the URL for the Producer Test Page:

   ```
   http://host:port/context-root/info
   ```

2. In the Producer Test Page, you can perform further diagnostic steps as described in the following sections.

Task 2 Run the JSF Portlet as a Servlet

To verify that an Oracle JSF Portlet Bridge portlet producer is running correctly, you must first verify that the producer application runs correctly through standard HTTP requests. If the artifacts the producer exposes as portlets do not run as servlets, they will not run as portlets.

To run a JSF portlet as a servlet:

1. In the Producer Test Page, click the **run as servlet** link next to the portlet.

2. The portlet is called using standard HTTP to request the underlying page or task flow. The results of the request are displayed in a new browser window.

   If the resulting page or task flow does not render correctly, then there is a problem with the producer application that must be resolved before you can run the page or task flow as a portlet.

3. If the application developer configured the producer application with seeded parameters, there is also a **run as servlet with parameters** link next to the portlet in the Producer Test Page. Click this link to run the underlying page or task flow with those seeded parameters.

Task 3 Check the Persistent Store Type

Oracle Fusion Applications has adopted a standard to use a consumer preference store as the persistent store. Therefore, for Oracle Fusion applications producers, the **Persistent Store Type** displayed on the Producer Test Page should always be set to **Consumer**.

Although other configurations are acceptable for applications that are built to assume such a configuration, having a non-consumer setting in Oracle Fusion applications indicates an issue in the producer. For Oracle Fusion applications to work correctly, they require a consumer preference store.

Task 4 Examine the SOAP Monitor

The SOAP monitor provides access to the SOAP requests between the consumer and producer when rendering a portlet. This is very useful in diagnosing problems at the communication level.

To examine the SOAP monitor:

1. In the Producer Test Page, click the SOAP Monitor link at the bottom of the page.

2. When prompted, enter your user name and password.

> **Note:** To access the SOAP monitor you must be a member of the `Monitors` or `Administrators` role in the Identity Management System.

3. By default, the SOAP monitor is disabled, so the page is empty. You must first enable the monitor by clicking the **Enable** link at the top of the page.

4. The page does not automatically refresh, so to display SOAP messages, you must click the **Refresh** link.

5. To force a request to the failing portlet, go to the Portlet Consumer Test Page: Portlet page for the portlet and select **Refresh Portlet**.

6. When the portlet has rendered, or failed, click the Refresh link in the SOAP monitor to display the captured request.
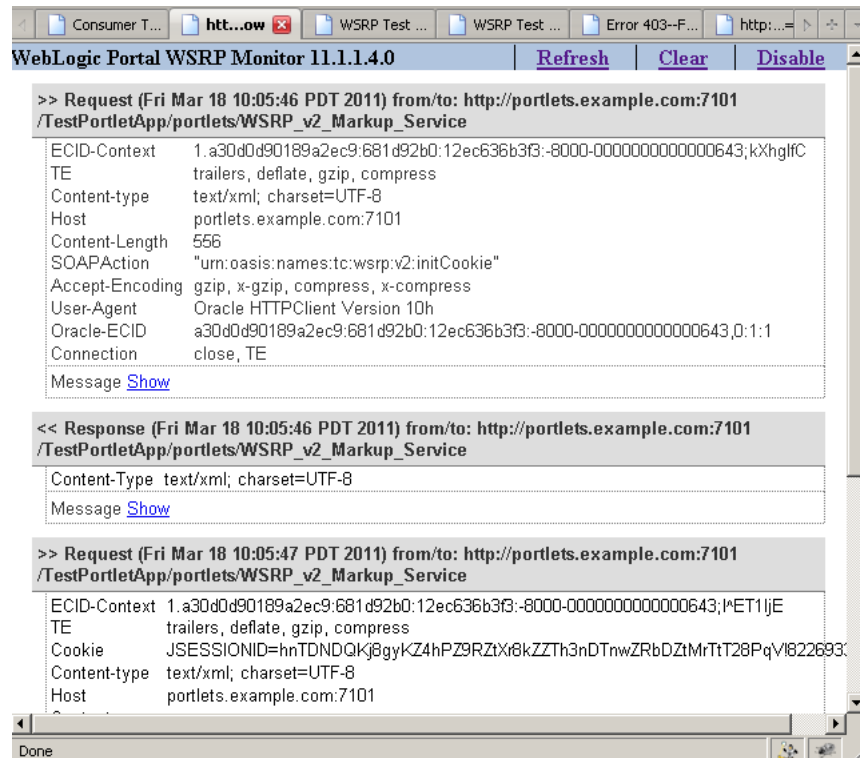
7. Now, you can investigate the SOAP messages that were sent and the responses to try to narrow down the cause of the problem.

> **Note:** If, after rerunning the portlet and refreshing the SOAP monitor, you see no messages displayed, this indicates that there may be a security issue between the producer and the consumer. You must verify that the correct WS-Security settings are set up for the producer and consumer to communicate.

23.3.2 Configuring the Portlet Logging File

To troubleshoot portlet issues, it is useful to add portlet log-handlers and loggers to the logging configuration file, `logging.xml`.

Example 23–4 shows how to add the portlet log-handlers and loggers. The example assumes that you are running the consumer and producer applications on the same WebLogic Server instance. If you are running the consumer and producer applications on different instances, you must split them up appropriately.

> **Note:** Add the log entries at the end of the file to ensure that they override any seeded settings.

Example 23–4   Configuring Log Files for Troubleshooting Portlet Issues

```
<!-- NOTE: You need to change the path where the logfile is located -->
<log_handlers>
...
   <!-- Portlet Consumer -->
   <log_handler name="portlet-consumer-handler" class="oracle.core.ojdl.logging.ODLHandlerFactory">
      <property name="format" value="ODL-Text"/>
```

```
            <property name="path" value="/scratch/logs/portlet-consumer.log"/>
    </log_handler>

    <!-- Portlet Producer -->
    <log_handler name="portlet-producer-handler" class="oracle.core.ojdl.logging.ODLHandlerFactory">
        <property name="format" value="ODL-Text"/>
        <property name="path" value="/scratch/logs/portlet-producer.log"/>
    </log_handler>

    <!-- Portlet Bridge -->
    <log_handler name="portlet-bridge-handler" class="oracle.core.ojdl.logging.ODLHandlerFactory">
        <property name="format" value="ODL-Text"/>
        <property name="path" value="/scratch/logs/portlet-bridge.log"/>
    </log_handler>
...
</log_handlers>

<loggers>
...
    <!-- Portlet Consumer -->
    <logger name="oracle.portlet.client" level="FINEST" useParentHandlers="false">
        <handler name="portlet-consumer-handler"/>
    </logger>

    <!-- Portlet Servers -->
    <logger name="com.bea.portlets" level="FINEST" useParentHandlers="false">
        <handler name="portlet-producer-handler"/>
    </logger>
    <logger name="com.bea.netuix" level="FINEST" useParentHandlers="false">
        <handler name="portlet-producer-handler"/>
    </logger>
    <logger name="com.bea.wsrp" level="FINEST" useParentHandlers="false">
        <handler name="portlet-producer-handler"/>
    </logger>
    <logger name="oracle.portlet.producer" level="FINEST" useParentHandlers="false">
        <handler name="portlet-producer-handler"/>
    </logger>

    <!-- Portlet Bridge -->
    <logger name="oracle.portlet.bridge" level="FINEST" useParentHandlers="false">
        <handler name="portlet-bridge-handler"/>
    </logger>
    <logger name="oracle.portlet.server.bridge" level="FINEST" useParentHandlers="false">
        <handler name="portlet-bridge-handler"/>
    </logger>
...
</loggers>
```

The logging configuration file is located in:

DOMAIN_HOME/config/fmwconfig/servers/*server*/logging.xml

The log file name is also defined in `logging.xml`. By default the log file name is:

DOMAIN_HOME/servers/*server*/logs/*server*-diagnostic.log

23.4 Using My Oracle Support for Additional Troubleshooting Information

You can use My Oracle Support (formerly MetaLink) to help resolve Oracle Fusion Applications problems. My Oracle Support contains several useful troubleshooting resources, such as:

- Knowledge base articles

- Community forums and discussions

- Patches and upgrades

- Certification information

You can access My Oracle Support at `https://support.oracle.com`.

24

Troubleshooting Performance, Scalability, Reliability

This chapter describes performance, scalability, and reliability issues that you might encounter and explains how to solve them.

This chapter includes the following topics:

- Introduction to Troubleshooting Performance, Reliability, and Scalability

- Overall Slowness

- Stuck Threads

- View Detailed Timing Of a Request Using a JRockit Flight Recorder (JFR) File

- Memory Leaks and Heap Usage Pressure

- Connection Usage

- Connection Leaks

- Slow Requests Using SQL trace

- Excessive Activation

- Slow Oracle Enterprise Scheduler Jobs of SQL Type

- Using My Oracle Support for Additional Troubleshooting Information

24.1 Introduction to Troubleshooting Performance, Reliability, and Scalability

This section provides guidelines and a process for using the information in this chapter. Using the following guidelines and process will focus and minimize the time you spend resolving problems.

Guidelines

When using the information in this chapter, Oracle recommends:

- After performing any of the solution procedures in this chapter, immediately retrying the failed task that led you to this troubleshooting information. If the task still fails when you retry it, perform a different solution procedure in this chapter and then try the failed task again. Repeat this process until you resolve the problem.

- Making notes about the solution procedures you perform, symptoms you see, and data you collect while troubleshooting. If you cannot resolve the problem using

the information in this chapter and you must log a service request, the notes you make will expedite the process of solving the problem.

Process

Follow the process outlined in Table 24–1 when using the information in this guide. If the information in a particular section does not resolve your problem, proceed to the next step in this process.

Table 24–1 Process for Resolving Performance Issues

Step	Section to Use	Purpose
1	Chapter 9 and Chapter 10 for locating key metrics Chapter 11 for diagnosing Java applications in the middle tier Chapter 12 for monitoring and tuning the Oracle database	Collect symptoms about the performance problem to determine if the problem is related the following: ■ Response time or throughput ■ Widespread or limited to specific users and flows Determine what changed since the system was last performing well.
2	Section 24.2 through Section 24.10	Use Section 24.2 if the problem is widespread. Otherwise, review the problem description in Section 24.3 through Section 24.10 to see if there is a match These section describes: ■ Possible causes of the problems ■ Solution procedures corresponding to each of the possible causes
4	Section 24.11	Use My Oracle Support to get additional troubleshooting information about Oracle Fusion Applications or performance, scalability, and reliability. My Oracle Support provides access to several useful troubleshooting resources, including Knowledge Base articles and Community Forums and Discussions.
5	Section 24.11	Log a service request if the information in this chapter and My Oracle Support does not resolve your problem. You can log a service request using My Oracle Support at `https://support.oracle.com`.

In addition to this process, for more information about determining if database cache sizes need to be increased, see "Automatic Database Performance Monitoring" chapter in the *Oracle Database 2 Day + Performance Tuning Guide* to use Automatic Database Diagnostic Monitor (ADDM) reports

24.2 Overall Slowness

Problem

When many users report slowness across many business flows, you need to determine if the cause is the host, components, heap usage, web sessions, or too many users accessing the application at one time.

Solution

To resolve this problem, perform the following steps:

1. From Grid Control, check host health.

   For example, you can view the host CPU usage from Performance Summary page. To access this page from Grid Control:

   a. From the home page, click the **Targets** tab.

   b. From the **Targets** tab, click the **Hosts** secondary tab.

      The Hosts page displays the overall status of all the computers in the environment.

   c. From the **Search** list, search for specific host.

   d. Click a specific host name to monitor the performance.

      The Host home page displays.

   e. Click the **Performance** tab.

2. Check if any key components are down:

   ■ For Fusion Applications Control, see Section 10.2.1 and Section 10.2.2.

   ■ For Grid Control, see Section 2.7.

3. Check heap usage in EM Performance Summary page.

   ■ For Fusion Applications Control, see "Monitoring the Oracle Fusion Applications Middle Tier" chapter in the *Oracle Fusion Middleware Administrator's Guide* and Table 10–6.

   ■ For Grid Control, see Table 10–9.

4. If the heap is constantly close to 100 percent, then search with the string `OutOfMemoryErrors` in the Oracle WebLogic Server *server_name*.out in the following directories:

   ```
   (UNIX) DOMAIN_HOME/servers/server_name/logs
   (Windows) DOMAIN_HOME\servers\server_name\logs
   ```

5. If there are `OutOfMemoryErrors`, a heap dump would have been generated in the directory specified by the `-DHeapDumpPath` parameter from the Oracle WebLogic Server startup JVM option. Submit the heap dump to Oracle Support for further analysis of what is retaining memory.

6. Use Fusion Applications Control to view log messages, and see if there are a lot of messages getting logged and incidents getting raised. See Chapter 13 for and Chapter 17 for information about Oracle Fusion Applications log file, and see the "Managing Log Files and Diagnostic Data" chapter in *Oracle Fusion Middleware Administrator's Guide* for information about Oracle Fusion Middleware log files.

7. Address the source of the errors and incidents, and verify the log level is set to `SEVERE`.

8. Determine the number of web sessions in Fusion Applications Control or Grid Control. This number can fluctuate depending on the flow and heap size. Therefore, monitoring the trend can help you find spikes.

   For Fusion Applications Control:

   ■ View the number of active sessions for a specific Oracle Fusion application product:

      a. From the navigation pane, expand the product family, then **Products**, and then select the product.

    b. In the Product home page, locate the **Servers** sections.

    c. View the **Active Sessions** metric.

- View the number of active sessions for an Oracle WebLogic Server domain:

    a. From the navigation pane, expand the farm and then **WebLogic Domain**.

    You can similarly check the active session for an Oracle WebLogic Server cluster or Managed Server from the navigation pane.

    b. Select a domain.

    c. In the WebLogic Domain home page, in the table on the left-hand side of the page, view the **Active Sessions** column.

For Grid Control:

a. Click the **Targets** tab.

b. Click the **Middleware** secondary tab.

c. From the **Search** list, select **Oracle WebLogic Server Domain**, and then click **Go**.

You can similarly check the active sessions for an Oracle WebLogic Server cluster or Managed Server by selecting **Oracle WebLogic Server Cluster** or **Oracle WebLogic Managed Server**.

d. Click on a domain.

The WebLogic Server Domain home page displays.

e. In the table on the left-hand side of the page, view the **Active Sessions** column.

f. In the Product home page, locate the **Servers** sections.

g. View the **Active Sessions** metric.

9. If there is a spike in the number of web sessions, find out what is generating the additional load and what tests are being run around that time. By default, the session timeout is set to 15 minutes.

10. Check data source health and see if it is running out of connection.

- For Fusion Applications Control, see Section 10.4.1.

- For Grid Control, check the **Server Datasource** metrics, as described in Section 10.5.2.

11. Review current execution stacks for the Oracle WebLogic Server threads. There are several ways to perform this step:

Oracle WebLogic Server Administration Console:

a. In the **Domain Structure**, expand **Environment** and then **Servers**.

b. In the Summary of Servers page, click the server from the table.

c. In the Settings for *server_name* page, click on the **Monitoring** tab, then click on **Threads** subtab.

d. Click Dump Thread Stacks.

e. Click **Save**.

For Grid Control:

a. Click the **Targets** tab.

    b. Click the **Middleware** secondary tab.

    c. From the **Search** list, select **Oracle WebLogic Managed Server**, and then click **Go**.

       You can similarly check the active sessions for an Oracle WebLogic Server cluster or Managed Server by selecting **Oracle WebLogic Server Cluster** or **Oracle WebLogic Managed Server**.

    d. Click on a specific server having problems.

       The WebLogic Server home page displays.

    e. From the **WebLogic Server** menu, choose **JVM Diagnostics** > **Threads** > **Real-Time Analysis**.

    f. In the **JVMs** section, click on a thread in the upper section show details in the **Threads** section.

12. If the threads are not blocked, follow the instructions in Section 11.3 to review the top Java methods and Section 11.4 to review the top SQL using JVM diagnostics.

13. Extract JFR recording and review timing breakdown of slow requests. See Section 24.4.

24.3 Stuck Threads

Problem

Stuck threads may result if the server is nearing out of memory. If the server is close to out of memory, all requests should slow down. To resolve an out-of-memory issue, see Section 24.4.

If a request is taking longer than 10 minutes, the stuck thread is reported to Oracle WebLogic Server *server_name*.out in the following directories:

```
(UNIX) DOMAIN_HOME/servers/server_name/logs
(Windows) DOMAIN_HOME\servers\server_name\logs
```

For example:

```
<Mar 4, 2011 7:44:08 AM PST> <Error> <WebLogicServer> <BEA-000337> <[STUCK]
ExecuteThread: '19' for queue: 'weblogic.kernel.Default (self-tuning)' has been
busy for "600" seconds working on the request
"weblogic.servlet.internal.ServletRequestImpl@18986012[
GET /productManagement/faces/PimDashboardUiShellPage?_afrLoop=1398820150000&_
afrWindowMode=0&_adf.ctrl-state=a44e7uxcc_13 HTTP/1.1
Accept: image/gif, image/x-xbitmap, image/jpeg, image/pjpeg,
application/x-shockwave-flash, application/x-ms-application,
application/x-ms-xbap, application/vnd.ms-xpsdocument, application/xaml+xml,
application/vnd.ms-excel, application/vnd.ms-powerpoint, application/msword, */*
Accept-Language: fr
UA-CPU: x86
...
]", which is more than the configured time (StuckThreadMaxTime) of "600" seconds
. Stack trace:
Thread-164 "[STUCK] ExecuteThread: '19' for queue: 'weblogic.kernel.Default
(self-tuning)'" <alive, in native, suspended, priority=1, DAEMON> {
    jrockit.net.SocketNativeIO.readBytesPinned(SocketNativeIO.java:???)
    jrockit.net.SocketNativeIO.socketRead(SocketNativeIO.java:24)
    java.net.SocketInputStream.socketRead0(SocketInputStream.java:???)
    java.net.SocketInputStream.read(SocketInputStream.java:107)
```

...

In this example, the request has been running longer than the configured 600 seconds. Here is the associated stack trace showing the thread is stuck:

```
Thread-164 "[STUCK] ExecuteThread: '19' for queue: 'weblogic.kernel.Default
(self-tuning)'" <alive, in native, suspended, priority=1, DAEMON> {
jrockit.net.SocketNativeIO.readBytesPinned(SocketNativeIO.java:???)
jrockit.net.SocketNativeIO.socketRead(SocketNativeIO.java:24)
java.net.SocketInputStream.socketRead0(SocketInputStream.java:???)
java.net.SocketInputStream.read(SocketInputStream.java:107)
...
```

Solution

If the stack shows the thread is waiting for a response from another server, check the status of the other server and see it has performance problems before proceeding with this solution.

To determine what the stuck thread was doing prior to becoming stuck, perform the following steps:

1. Look at the next few log messages in `server_name.out` for a message indicating an incident has been created. For example:

   ```
   <Mar 4, 2011 7:44:10 AM PST> <Alert> <Diagnostics> <BEA-320016>
   <Creating diagnostic image in DOMAIN_HOME/servers
   /ProductManagementServer_1/adr/diag/ofm/SCMDomain/
   ProductManagementServer_1/incident/incdir_394 with a lockout minute
   period of 1.>
   ```

   The above message may not always appear after each stuck thread reported. It is printed at most four times an hour. If the message does not appear, manually look for the `incident` directory by checking the `readme` file in the subdirectories under the following directories:

   ```
   (UNIX) DOMAIN_HOME/servers/server_name/adr/diag/ofm/domain_name/server_
   name/incident
   (Windows) DOMAIN_HOME\servers\server_name\adr\diag\ofm\domain_name\server_
   name\incident
   ```

   The incident directory contains a WLDF diagnostic image which contains the JFR recording, and a file containing the thread dump

   For more information about diagnosing incidents, see the "Diagnosing Problems" chapter in the *Oracle Fusion Middleware Administrator's Guide*.

2. Review thread dump to see call stack of the thread. If thread is blocked waiting for lock, check what the thread holding the lock is doing.

3. If call stack involves executing JDBC calls, you can go to Grid Control and check the top activity around that time window, and see if there is a session with a matching module and action. See Section 11.4.

4. Review the JRockit flight recording file `JRockitFlightRecorder.jfr` for more details. You will also need the ECID of the request which is recorded in the `readme.txt` file of the incident directory, and also the Oracle WebLogic Server log.

5. Perform the tasks in Section 24.4.

24.4 View Detailed Timing Of a Request Using a JRockit Flight Recorder (JFR) File

Problem

Certain requests are slow and there is a need to find out where time is spent

Solution

The JRockit Flight Recorder (JFR) file contains a record of various events that consume time, and can be used to help understand why a request is taking time

To resolve this problem, create a JFR file:

1. Extract a JFR file from an Oracle WebLogic Server server by running the following command:

   ```
   (UNIX) JROCKIT_HOME/bin/jrcmd jrockit_pid dump_flightrecording recording=1
   copy_to_file=path compress_copy=true
   (Windows) JROCKIT_HOME\bin\jrcmd.exe jrockit_pid dump_flightrecording
   recording=1 copy_to_file=path compress_copy=true
   ```

   See the "Running Diagnostic Commands" chapter in the *Oracle JRockit JDK Tools Guide* for more information about the `jrcmd` command-line tool.

2. To view the file, start the JRockit Mission Control Client from the following directories:

   ```
   (UNIX) JAVA_HOME/bin/bin/jrmc
   (Windows) JAVA_HOME\bin\jrmc.exe
   ```

3. Choose **File** > **Open File** to select the JFR file.

4. Locate the slowest requests or investigate a specific request:

To locate the slowest requests:	To investigate a specific request:
1. In the JRockitFlightRecorder.jfr page, click the **Events** icon. 2. Click the **Log** tab at the bottom of the page. 3. In the **Event Type** navigation pane on the left, locate **Dynamic Monitoring System** and then **HttpRequest**. 4. Click **HTTP request**; de-select all the other event types. 5. In the **Log** tab, in the **Event Log** section, click the **Duration** column to sort the duration in descending order. Each row corresponds to a HTTP Request and the duration column shows the response time for that request. 6. Click the row in the table to view the attributes of the requests. 7. In the **Event Attributes** sections, note the start time and the thread that serviced the request.	1. Find the Execution Context Identifier (ECID) of that request. If the request is related to an incident triggered by a `STUCK` thread, the incident `readme.txt` file will contain the ECID. Alternatively, you can search the Oracle WebLogic Server HTTP `access.log` for requests from specific users. See the "Viewing and Searching Log Files" section in the *Oracle Fusion Middleware Administrator's Guide.* 2. In the JRockit Mission Control Client, in the JRockitFlightRecorder.jfr page, choose the **WebLogic** icon, and then If the **Weblogic** icon is not available, choose **Help** > **Install Plugins** to download the Oracle WebLogic Server plug-in. 3. Click the **ECIDs** tab at the bottom of the age. 4. In the **ECIDs** section, from **Filter Column** list, select **ECID**. 5. Enter the ECID in the search box and choose <Enter>. 6. In the results table, highlight the row with the matching ECID and right-click to bring up the menu. 7. Choose **Operative Set** > **Clear**, and then **Operative Set** > **Add matching ECID** > *ECID* to add the ECID to the operative set. This enables users to view only events associated with the operative set. 8. Click the **Events** icon. 9. In the Event Type navigation pane on the left, locate **Dynamic Monitoring System** and then **HttpRequest**. 10. Click **HTTP request**; de-select all the other event types. ** 11. In the **Event Log** section, click **Show Only Operative Set**. Each row corresponds to the request with the matching ECID 12. Click the row in the table to view the attributes of the requests. 13. Note the start time and the thread that serviced the request

5. Once the start time and the thread that serviced the request are identified, in the Logs tab, drag the time selector at the top of the screen to include only the time window for the duration of the request.

6. In the **Event Log** section, perform the following search:

   a. Deselect **Show Only Operative Set**.

   b. Enter the thread name in the search box.

   c. From the **Filter Column** list, select **Thread**.

    d. Choose <Enter>.

7. In the **Event Type** navigation pane on the left, click the events of interest. Typically, these events are located under nodes **Dynamic Monitoring System**, **Java Application**, and **WebLogic** > **JDBC**.

The selected events appear in the table in the **Event Log** section.

8. Click the **Start Time** column to sort y the time when these events occur, or click the **Duration** column to view the events that took longest

The **JDBC Statement Execute** events corresponds to SQL execution. If there are slow SQLs, the event details give the SQL text. These events do not have callstacks.

9. To see to callstack for slow SQLs, view the **Socket Read** event that happens right after the **JDBC Statement Execute** event.

This event corresponds to Oracle WebLogic Server waiting for the SQL results to return, and it has callstack in the event details.

10. Review the callstacks for long **Java Blocked** and **Java Wait** events to see if the cause can be identified. See the "Analyzing Flight Recorder Data in JRockit Mission Control" section in the *Oracle Fusion Middleware Configuring and Using the Diagnostics Framework for Oracle WebLogic Server*.

11. If more details are needed to compare with what is captured in the default recording, and the user can reproduce the slowness, start an explicit recording. See the "Starting an Explicit Recording" section in the *Oracle JRockit Flight Recorder Run Time Guide*.

24.5 Memory Leaks and Heap Usage Pressure

Problem

Application performance degrades over time, heap usage and garbage collection activity increases overtime, sometimes `OutOfMemoryErrors` are seen.There could be memory leaks in the application, which causes the amount of free memory in the JVM to continuously decrease.

Solution

To solve this problem, perform the following:

1. Review the `server_name.out` file for `OutOfMemoryErrors` errors, which indicate a heap dump file has been written. The `server_name.out` file is located in the following directories:

```
(UNIX) DOMAIN_HOME/servers/server_name/logs
(Windows) DOMAIN_HOME\servers\server_name\logs
```

2. Restart the Managed Server.

See the following documentation resources to learn more about other methods for starting and stopping the Managed Servers:

- "Starting Managed Servers with a Startup Script" section in *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server*

- "Starting Managed Servers with the java weblogic.Server Command" section in the *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server*

- "Starting and Stopping Managed Servers Using Fusion Middleware Control" section in *Oracle Fusion Middleware Managing Server Startup and Shutdown for Oracle WebLogic Server*

- "Start and Stop Servers" and various startup and shutdown procedures in the Cluster section of the *Administration Console Online Help*.

If the problem persists, proceed to Step 3.

3. Open the file with a heap-dump analysis tool that can handle binary HPROF format, such as Eclipse Memory Analyzer.

4. Review what objects and classes are retaining most memory. Send the heap dump file to Oracle Support for further analysis.

5. Sometimes it may be necessary to take several heap dumps to see what objects or classes are consuming and increasing the amount of memory.

   To take heap dumps on demand, use the `jrcmd` command-line tool. See the "Running Diagnostic Commands" chapter in the *Oracle JRockit JDK Tools Guide*. Many heap dump analysis tools, such as Eclipse Memory Analyzer, enable you to compare two heap dumps to identify memory growth areas.

   Heap dumps provide information on why memory is retained. Sometimes it is necessary to know how memory is allocated to further resolve the issue. For these cases, proceed to Step 6.

6. Use the JRockit Memory Leak Detector tool that is part of JRockit Mission Control Client to understand how memory is allocated.

   For more information, see the JRockit Mission Control online help.

24.6 Connection Usage

Problem

The connection usage on the Oracle Database is high, or there is an Oracle process on the database host consuming high amount of CPU.

Solution

To find out the source of the connection causing the high CPU on To adjust the reference pool size from Fusion Applications Control:

1. Oracle Fusion Applications set values on a number of `v$session` attributes to indicate how the connection is being used. When looking at a connection consuming high CPU on the database, or when trying to understand what connections are used for what processes, inspect the value of these attributes as follows:

Attribute in v$session	Value Being Set
Process	Data Source Name (for example, ApplicationDB)
Program	Oracle WebLogic Server Domain plus the Managed Server name, prefixed by DS (for example, DS/FinancialDomain/AccountsReceivableServer_1)

Attribute in v$session	Value Being Set
Module	Oracle Application Development Framework: ADF BC application module name
	Oracle Enterprise Scheduler:
	■ Java job type: Class name, except `oracle.apps`
	■ PLSQL: the package and procedure name (for example, `mypkg.myproc`)
	■ Other jobs: Static: Executable name
	Oracle BI Publisher: Name of the report
Action	Oracle Application Development Framework: jspx name
	Oracle Enterprise Scheduler: Job definition name
	Oracle BI Publisher, if request is submitted:
	■ Oracle Enterprise Scheduler: Oracle Enterprise Scheduler job definition name
	■ Oracle BI Publisher Scheduler Job: Oracle BI Publisher job name submitted by the user
	■ Oracle BI Publisher online: Static string `BIP:Online`
	■ Oracle BI Publisher Web services: Name of the web services
Client_Identifier	Application User Name

2. If the error messages related to connection pool capacity being reached are also seen in Oracle WebLogic Server logs, use the solution for connection leaks described in Section 24.7.

24.7 Connection Leaks

Problem

When there are errors in the log, and the error message indicates connection pool size has been reached

Solution

To resolve this problem:

1. When data source is at maximum capacity and there are errors during connection reservation requests, then there may be connection leaks in the code

2. Enable JDBC profiling from the Oracle WebLogic Server Administration Console:

   a. In the **Domain Structure**, expand **Services** and then **Data Sources**.

   b. Click on the data source that needs to profiled, for example, **ApplicationDB**.

   c. In the Settings page, click on the **Configuration** tab, then click on **Diagnostics** subtab.

   d. Check the profiles that need to be collected (**PROFILE_TYPE_CONN_ USAGE_STR**).

   e. Click **Save**.

3. Configure the diagnostic archive where the profiling data is saved from the Oracle WebLogic Server Administration Console:

a. In the **Domain Structure**, expand **Services**, **Diagnostics**, and then **Archives**.

b. Click on the server where you want to make changes (archives are stored for each server)

c. In the Settings page, you can change archive location, size and how to retire data.

d. Check the profiles that need to be collected (**PROFILE_TYPE_CONN_ USAGE_STR**).

e. Click **Save**.

4. To retrieve profiling data, use the sample code (`http://download.oracle.com/docs/cd/E15051_ 01/wls/docs103/wldf_configuring/access_diag_ data.html#wp1100898`), with changes to the URL, username and password in the initialize method.

5. Run the sample code as a standalone program.

6. The program will capture the stack trace for each request for a connection from that data source. Inspect the callers to see the suspicious stack. This sample program requires connecting to a live Oracle WebLogic Server instance.

The diagnostic archive file under the archive location can also be provided to Oracle Support for further analysis.

Oracle WebLogic Server will not report a leak unless inactive connection timeout connection pool setting is set to a positive value. This cannot be done for Oracle Fusion Applications, as it will break functionality.

24.8 Slow Requests Using SQL trace

Solution

When a user reports that a specific operation is slow, and the slowness is reproducible and that slow database operations are suspected but the top activity reports did not provide sufficient information for resolving the problem.

Solution

To resolve this problem:

1. Enable SQL trace for the user session. See Section 12.2.4.6.

2. Ask user to re-run the problematic flow and collect the SQL trace files and review

24.9 Excessive Activation

Problem

When response time suddenly increases with rising user count, even though there is no memory pressure, it is possible that the reference pool size for key application modules needs to be increased. If there is a JFR recording to review, and you observe many events containing callstacks containing the `activateState` method, you should also try adjusting the reference pool size.

Solution

To adjust the reference pool size from Fusion Applications Control:

1. Review the number of web sessions from Performance Summary pages:

   For Fusion Applications Control:

   a. From the navigation pane, expand the farm, **Application Deployments**.

   b. From the Applications Deployments page, select the application.

   c. From the **Application Deployments** menu, choose **ADF > ADF Performance**.

      The ADF Performance page displays.

   d. Click the **Application Module Pools** tab.

   e. Sort the request by descending order.

   f. For the top 10 or so application modules, click the application module name to view the **Activations** count.

   For Grid Control:

   a. Click the **Targets** tab.

   b. Click the **Middleware** secondary tab.

   c. From the **Search** list, select **Oracle WebLogic Server Domain**, and then click **Go**.

   d. Click on a domain.

      The WebLogic Server Domain home page displays.

   e. In the table on the right-hand side of the page, expand the **Application Deployments** node.

   f. Click the target application.

      The Application Deployment page displays.

   g. From the **Application Deployments** menu, choose **ADF > ADF Performance**.

      The ADF Performance page displays.

   h. In the **Application Module Pools** table, from the **View** list, select **Total Requests**, and once selected, from the **Total Requests** column, click **Sort Descending**.

   i. For the top 10 or so application modules, click to see the details of each one.

   j. After selecting an application module, on the **Requests** graph, from the **Select metric to display in chart** list, select **Passivation** and **Activation** to add to the graph.

      If activation count is close to passivation and is constantly above 0, then following Step 2 to adjust

2. If the activation count constantly increases, increase the application module reference pool size from Fusion Applications Control:

   a. From the **Application Deployments** menu, choose **ADF > Configure ADF Business Components**.

      The ADF Configuration BC Configurations page displays.

   b. From the **Application Modules** section, click the application module of interest. From the left hand side, select the local by selecting a name that ends in **Local**.

    c. Click the **Pooling and Scalability** tab, and change the **Reference Pool Size** parameter.

24.10 Slow Oracle Enterprise Scheduler Jobs of SQL Type

Problem

When the user submits a SQL job type, the job remains in a `RUNNING` state for too long.

Solution

To resolve this problem, perform the following steps:

1. Use Fusion Applications Control to find the database session ID that was used to process the job:

   1. Search for the request, as described in Section 5.7.2.

   2. On the Request Details page, in the **Request Properties** section, next to the **Execution Type** field, click the eye glasses icon.

      The Spawned Process Details dialog displays. This will bring up a pop-up with the database session id that was used to process this job

   3. Take note of the value in the **Session Id** field, and then click **OK**.

2. Use Grid Control to ensure the request processor and request dispatcher are running:

   a. Run an Active Session History (ASH) report for the session within the relevant time window to inspect top SQLs and top wait events. See the "Resolving Transient Performance Problems" section in the *Oracle Database 2 Day + Performance Tuning Guide*.

   b. Identify time consuming SQLs and tune following normal SQL tuning procedures. See Section 11.4.

24.11 Using My Oracle Support for Additional Troubleshooting Information

You can use My Oracle Support (formerly MetaLink) to help resolve Oracle Fusion Applications problems. My Oracle Support contains several useful troubleshooting resources, such as:

- Knowledge base articles

- Community forums and discussions

- Patches and upgrades

- Certification information

> **Note:** You can also use My Oracle Support to log a service request.

You can access My Oracle Support at `https://support.oracle.com`.

Part VI

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"

A

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 15. 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 `ApplSession` using the `HttpSession`, but it is marked as non-serializable so that it does not get persisted in the event of failover/replication. The `ApplSession` will be re-established from the session cookie from the client that is set when the session is first created.

B

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 Chapter 20.

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.

This section includes the following topics:

- Oracle Enterprise Scheduler Architecture
- Oracle Enterprise Scheduler Components
- Oracle Enterprise Scheduler Life Cycle
- Oracle Enterprise Scheduler Life Cycle Tools

B.2.1 Oracle Enterprise Scheduler Architecture

Oracle Enterprise Scheduler is installed to an Oracle WebLogic Server instance, on which it runs. The Oracle Enterprise Scheduler service component sits on top of Oracle JRF and is secured by Oracle Web Services Manager. Oracle Enterprise Scheduler manages scheduled job submissions and job definitions.

Figure B–1 shows the Oracle Enterprise Scheduler runtime architecture in the context of Oracle Fusion Middleware components.

Figure B–1   Oracle Enterprise Scheduler Runtime Architecture
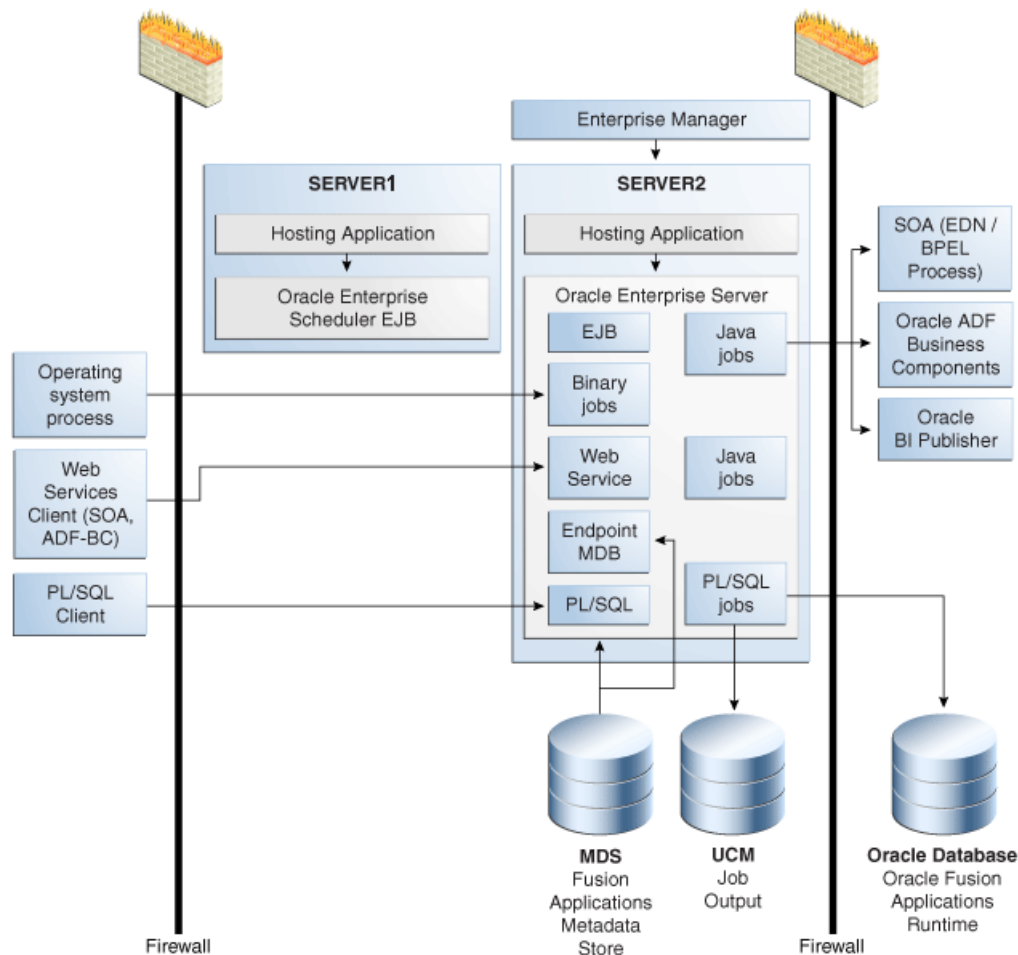
The components of the Oracle Enterprise Scheduler runtime architecture are as follows:

- Oracle Enterprise Scheduler client applications: Various applications can request the execution of a scheduled job. Applications include Oracle Fusion applications, Web service clients such as SOA or Oracle ADF Business Components and PL/SQL applications.

- Oracle Enterprise Scheduler: Fusion Applications Control enables you to manage Oracle Enterprise Scheduler clusters, services and jobs. Oracle Enterprise Scheduler accesses Oracle Fusion Applications metadata via MDS. Scheduled job output is saved to Oracle Universal Content Management. Oracle Enterprise Scheduler includes interfaces and APIs that enable interaction with applications and external components. For example, a PL/SQL client uses the Oracle Enterprise Scheduler PL/SQL API to request a scheduled job.

- Components accessed by Oracle Enterprise Scheduler: Components that may be accessed by Oracle Enterprise Scheduler include SOA components, Oracle ADF Business Components services, and Oracle Business Intelligence Publisher.

Client applications accessing EJBs connect to Oracle Enterprise Scheduler over RMI, whereas client applications using Oracle Enterprise Scheduler web services use HTTP. Connections from client applications to the server are persistent, short-lived asynchronous interactions that sometimes use callback functions.

B.2.2 Oracle Enterprise Scheduler Components

Oracle Enterprise Scheduler components are as follows:

- Oracle ADF Server (Oracle Enterprise Scheduler client): `RuntimeServiceEJB` and `MetadataServiceEJB` are deployed as shared libraries. These libraries are imported in the ADF client applications (ears).

- Oracle Enterprise Scheduler Server (Oracle Enterprise Scheduler Runtime): The Oracle Enterprise Scheduler service component manages all scheduled jobs.

- Core runtime: This is an Oracle Enterprise Scheduler application EAR file which contains a JCA resource adapter, multiple EJB components and JRF web service modules (WAR files).

- Hosting applications: A hosting application is an EAR file that imports the `ESSEndpointMDB`, `RuntimeServiceEJB` and `MetadataServiceEJB` shared libraries. An Oracle Enterprise Scheduler hosting application submits job requests using the Oracle Enterprise Scheduler libraries or an integrated job request submission interface.

- Oracle Database Scheduler: The standard Oracle Database Scheduler is used to execute Oracle Enterprise Scheduler PL/SQL jobs.

- Oracle Enterprise Scheduler uses Java process APIs to spawn native binary jobs.

Oracle Enterprise Scheduler relies on the following data sources:

- Oracle Enterprise Scheduler runtime (XA)

- Oracle Enterprise Scheduler runtime (non-XA)

- Oracle Enterprise Scheduler Metadata Store (non-XA)

An XA transaction, in the most general terms, is a global transaction that may span multiple resources. A non-XA transaction always involves just one resource, and generally cannot participate in a global transaction.

External dependencies include a runtime database, an MDS repository, as well as Oracle SOA Suite, Oracle ADF Business Components, Oracle BI Presentation Services, Oracle Universal Content Management and so on, depending on the components involved in implementing a given job.

B.2.3 Oracle Enterprise Scheduler Life Cycle

The Oracle Enterprise Scheduler engine starts up as part of the standard J2EE application initialization by Oracle WebLogic Server. The Oracle Enterprise Scheduler JCA adapter connects to the runtime schema and polls for scheduled work items.

The following is the sequence of the execution of a client request in Oracle Enterprise Scheduler.

1. A client application submits a job request.

2. The client application sends the job request to Oracle Enterprise Scheduler.

3. Oracle Enterprise Scheduler reads the metadata for the job request.

4. Oracle Enterprise Scheduler places the job request and the job metadata in a queue in the Oracle Enterprise Scheduler runtime data store.

5. Based on schedule and request processor availability, Oracle Enterprise Scheduler sends a message to the hosted application, which includes all the job request parameters and metadata captured at the time of submission.

6. The hosting application executes the job and returns a status. Job output and logs are written to Oracle Universal Content Management.

7. Oracle Enterprise Scheduler updates the history with the job request status.

Figure B–2 shows the changes in job state during the life cycle of an executed job request.

Figure B–2   Job Request Changes in State During Runtime
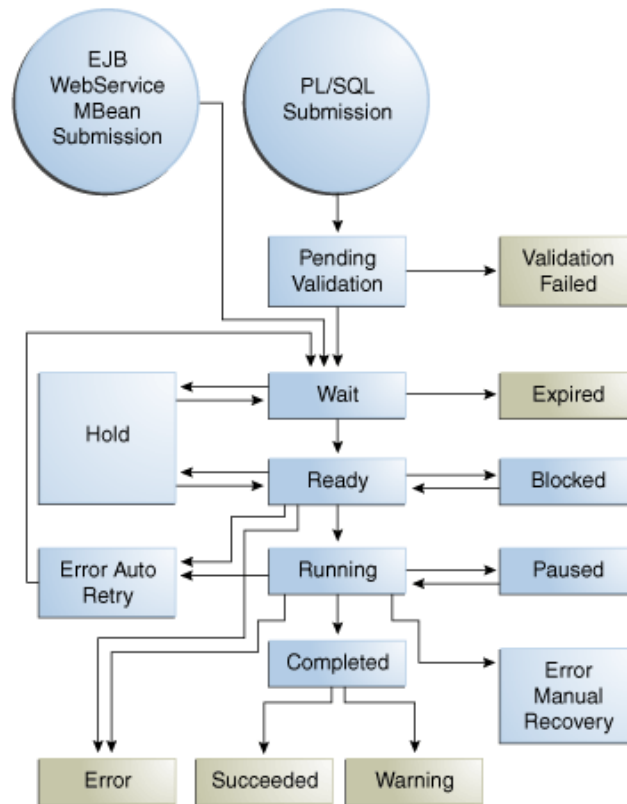
Figure B–3 displays the changes in state for an executable job request for which the executing user has cancelled the request.

Figure B–3   Changes in Job State Following Cancellation
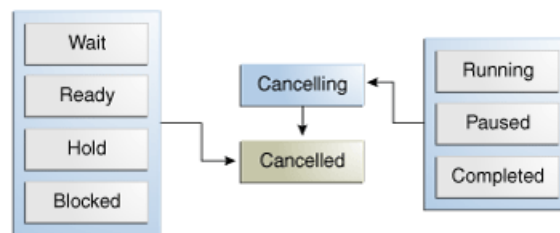
Figure B–4 displays state transitions for a job request submitted with a schedule.

Figure B–4   State Transitions for a Job Request Submitted with a Schedule
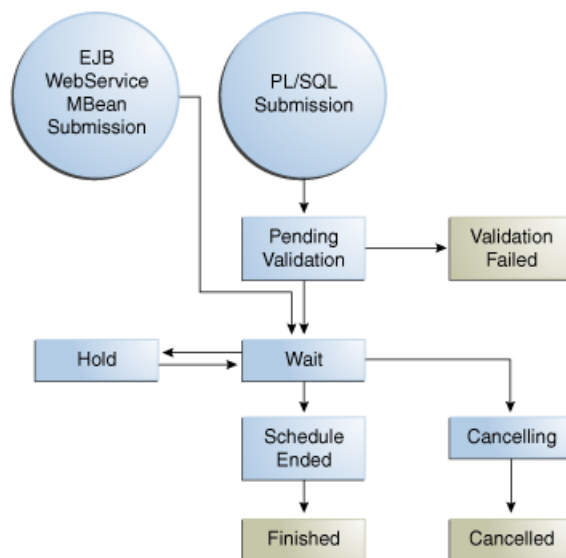
B.2.4 Oracle Enterprise Scheduler Life Cycle Tools

As Oracle Enterprise Scheduler runs on an Oracle WebLogic Server instance, you can manage Oracle Enterprise Scheduler using Oracle Fusion Middleware Node Manager for SOA.

Oracle Enterprise Scheduler jobs can be hosted on the same Oracle WebLogic Server instance (or a remote Oracle WebLogic Server instance), database and binary processes. Oracle Enterprise Scheduler controls the life cycle of Oracle Enterprise Scheduler jobs. Use Oracle Enterprise Manager to monitor and manage Oracle Enterprise Scheduler jobs. For more information about managing Oracle Enterprise Scheduler jobs, see Chapter 5.

For more information about Oracle Fusion Middleware Node Manager, see the chapter "Using Node Manager" in *Oracle Fusion Middleware Node Manager Administrator's Guide for Oracle WebLogic Server*.

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
- Oracle Enterprise Scheduler Cluster Architecture
- Failover Requirements
- Scalability
- Backup and Recovery
- Load Balancing

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.

- 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 Universal Content Management 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 Oracle Enterprise Scheduler Cluster Architecture

is an architectural diagram of a two node Oracle Enterprise Scheduler cluster.

Figure B–5   Oracle Enterprise Scheduler Two Node Cluster
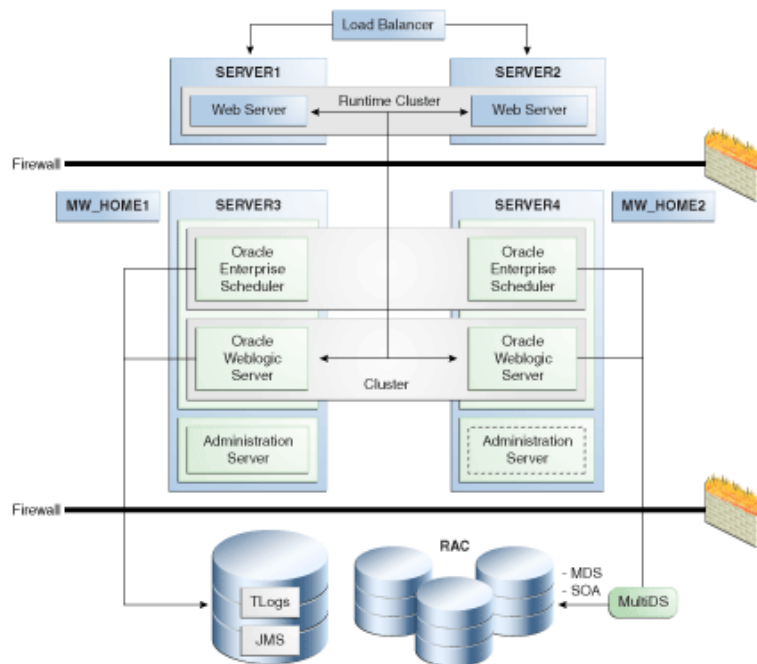
This configuration includes the following components:

- Hardware load balancer.

- A cluster of Oracle WebLogic Servers running on two servers.

- Two Oracle Fusion Middleware homes running on two servers.

- A two node cluster of Oracle Enterprise Scheduler instances, each running on a different instance of Oracle Fusion Middleware.

- A cluster of Oracle RAC databases with multi-DS configuration is required for availability of runtime and MDS schemas. Multi-DS and Oracle RAC provide for database failure.

- Shared persistence storage

- A Metadata Store

- An HTTP load balancer is required for web service interactions.

For more information regarding high availability architectures, see "Configuring High Availability for Oracle Fusion Middleware SOA Suite" in the *Oracle Fusion Middleware High Availability Guide*.

B.3.4 Failover Requirements

An HTTP load balancer provides load balancing so as to re-route requests in the event of node failure. There are no time out requirements for the load balancer or firewalls, as long as components use persistent connections. Likewise, session state replication and failover are not required.

Load balancing is used for actions such as submitting a job and querying its status using the Oracle Enterprise Scheduler web service interface. This load balancing occurs independently of where the job is scheduled to execute.

As Oracle Enterprise Scheduler does not use JMS, no JMS failover is needed. However, JTA failover is not required in the default Fusion Applications context. Oracle Enterprise Scheduler uses stateless session beans and JCA RAR files rather than stateful session beans. Therefore, no EJB state failover is required.

This section contains the following topics:

- Request Processor Failover
- External Component Failover

B.3.4.1 Request Processor Failover

Oracle Enterprise Scheduler includes a request processor component, which represents a single Managed Server in the Oracle Enterprise Scheduler cluster. Request processors process job requests, such that job execution is connected to one or more request processors.

If all jobs are targeted at a number of request processors, jobs are not dependent on a particular request processor. If a job is targeted at a particular request processor, any jobs tied to that request processor execute only when the Managed Server is available and an active workshift exists for the job.

B.3.4.2 External Component Failover

Oracle Enterprise Scheduler interacts with Oracle Fusion Middleware and other components such as Oracle SOA Suite, Oracle ADF Business Components, and so on. If one of these external components fail, it is possible that any running jobs may fail.

You can prevent external component failure from affecting jobs using proper configuration. Table B–1 lists the external components that may fail, along with the steps to take to prevent failed Oracle Enterprise Scheduler jobs.

Table B–1   Oracle Enterprise Scheduler External Component Failover

External Component	Steps to Prevent Failure
Oracle WebCenter	Integrate with a cluster of Oracle WebCenter service through a load balancer.
RAC Database	Use multi-DS for Oracle RAC database integration.
Oracle SOA Suite, Oracle ADF, and so on	Configure retries for jobs that depend on these components.

B.3.5 Scalability

Horizontal scalability—adding Managed Servers on different machines—tends to enable better performance than vertical scalability—(adding Managed Servers on the same machine.

Use standard Oracle WebLogic Server cluster scaling methodologies for horizontal scaling. For more information about Oracle WebLogic Server clustering, see the "High Availability for WebLogic Server" chapter in the *Oracle Fusion Middleware High Availability Guide*. You can increase the concurrent processing of jobs within a work assignment by increasing the thread allocation of the request processor (by editing the workshift for the request processor) or by binding the work assignment to more than one request processor. For more information, see Section 5.11.2.1 and Section 5.11.1.1.

B.3.6 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 Universal Content Management upon job completion. Only one database is used, and two phase commit is not used.

B.3.7 Load Balancing

Per Oracle best practice, use any hardware load balancer to load balance two or more Oracle HTTP servers, which, in turn, load balance the Oracle WebLogic Server Managed Servers in the cluster.

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.

This section contains the following topics:

- Starting and Stopping the Cluster

- Propagating Configuration Changes to the Cluster

- Deploying Applications to the Cluster

- Failures and Expected Behavior

B.4.1 Starting and Stopping the Cluster

Oracle Enterprise Scheduler uses standard J2EE components. As such, Oracle WebLogic Server determines the startup sequence. Oracle Enterprise Scheduler also allows implementing throttling to prevent surges in load.

When stopping a cluster, Oracle Enterprise Scheduler and all local Java jobs terminate. Oracle Enterprise Scheduler also attempts to stop all local binary jobs. However, asynchronous jobs such as SOA or PL/SQL jobs will continue. The asynchronous callback from SOA or Oracle ADF Business Component services cannot be delivered if the entire Oracle Enterprise Scheduler cluster is down.

In the case of an abrupt shutdown, the server attempts to recover any on-going transactions.

B.4.2 Propagating Configuration Changes to the Cluster

There are two types of configuration: container or server, and job metadata configuration. Job metadata is stored in Oracle Metadata Repository. The process of configuring the server is the same as that of configuring a standard Oracle WebLogic Server, as part of platform configuration maintained by the Oracle WebLogic Server

Configuration Framework. Job metadata configuration changes are deployed from Oracle JDeveloper.

At the cluster level, any configuration changes are propagated by deploying EAR files or metadata. Configuration data is stored either in the database or an EAR file. You can modify the configuration files in the EAR file, such as `connections.xml`, using the MBean browser in Fusion Applications Control.

Cluster members are independent, sharing only the database. There is no communication among members of a cluster.

B.4.3 Deploying Applications to the Cluster

Applications are deployed to the cluster using standard Oracle WebLogic Server EAR files, using standard J2EE defined Oracle WebLogic Server mechanisms. EAR files can be deployed without restarting the server.

An application deployment includes an EAR file, which contains a JAZN and MAR file. The JAZN file, which contains access privileges for the scheduled jobs, is stored in LDAP under the control of Oracle Authorization Policy Manager. The MAR file contains metadata, and is stored in MDS. Oracle WebLogic Server deploys the application EAR file to all Managed Servers in the cluster.

B.4.4 Failures and Expected Behavior

In the event of failure, the main way to ensure the continuation of job processing is to configure a cluster of Oracle Enterprise Schedulers. If a server fails, another node in the cluster transitions all jobs running on the failed server to the relevant state. Synchronous jobs for example, end in an error state and may be retried depending upon whether retries are configured.

In order to enable high availability for the data tier, use Oracle RAC.

This section contains the following topics:

- Retries
- Death Detection and Restart
- Oracle Java Transaction API Migration and Oracle Java Message Service

B.4.4.1 Retries

You can configure retries for jobs and job time outs for asynchronous jobs. For more information about configuring retries and time-outs for a job, see Section 5.7.1.1.

B.4.4.2 Death Detection and Restart

In order to enable death detection and recovery, each Oracle Enterprise Scheduler cluster node updates its record in the database every minute. This is called a heartbeat. Other nodes monitor the heartbeat, and when the record does not change for a period of time, the server is assumed to be dead. When a server death is detected, each job running on that server is handled. Synchronous jobs are marked as completed with errors, whereas asynchronous jobs continue to run remotely. If retry has been configured, the job marked as completed with errors restarts. Death detection tends to take about ten minutes.

On death detection, the output and log files of a node must be accessible from another node. As such, the file directory containing the job output and log files must be located on a shared file system. This directory is listed in the file `connections.xml`.

B.4.4.3 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.

C

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

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.infrastructure 360.publicViewEcsf.CustomerAccountPVO	Both	
	Campaigns - oracle.apps.marketing.coreMarketing.campaigns.publicViewEcsf.CampaignsPVO	Global	

Table C–1   (Cont.) Searchable Objects by Product Family

Product Family	Searchable Object Name	Scope	Notes
	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	
	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	

Table C–1   (Cont.) Searchable Objects by Product Family

Product Family	Searchable Object Name	Scope	Notes
	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.
	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.

Table C–1 (Cont.) Searchable Objects by Product Family

Product Family	Searchable Object Name	Scope	Notes
	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	
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	

Table C–1   (Cont.) Searchable Objects by Product Family

Product Family	Searchable Object Name	Scope	Notes
	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	

Table C–2   (Cont.) Search Categories by Product Family

Product Family	Search Category Name	Business Object	Scope
		Sourcing Contracts	
	Campaigns	Campaigns	GLOBAL
	Competitors	Competitors	GLOBAL
	Contacts	Contacts	GLOBAL
	Contract Documents	Sourcing Contract Documents	LOCAL
		Purchasing Contract Documents	
		Contract Documents	
	Customer	Customer Interactions	LOCAL
		Contacts	
		Billing Accounts	
		Customer Assessments	
		Customer Contracts	
		Customer Profiles	
		Customer References	
		Customer Opportunities	
		Customer Notes	
		Customer Leads	
	Customers	Customer Profiles	GLOBAL
		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
Oracle Fusion Financials	Payables	Payables Standard Invoice	GLOBAL
		Payables Payment	
		Payables Credit Memo	
		Payables Prepayment	
		Payables Refund	

Table C–2   (Cont.) Search Categories by Product Family

Product Family	Search Category Name	Business Object	Scope
	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	GLOBAL
Oracle Fusion Human Capital Management	Worker	Worker	GLOBAL
		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
		Material Transactions	GLOBAL
		COGS Transactions	GLOBAL
		Item Costs	GLOBAL
	Receipt Accounting	Receiving Transactions	GLOBAL
		Accrual Adjustments	GLOBAL

Table C–2 (Cont.) Search Categories by Product Family

Product Family	Search Category Name	Business Object	Scope
	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

Table C–3   (Cont.) Index Schedules for Oracle Fusion Customer Relationship Management

Schedule Name	Searchable Object	Business Object	Repeat Interval
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
Leads	Leads	oracle.apps.marketing.leadMgmt.leads.publicViewEcsf.MklLeadPVO	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

Table C–3   (Cont.) Index Schedules for Oracle Fusion Customer Relationship Management

Schedule Name	Searchable Object	Business Object	Repeat Interval
Purchasing Contracts	Purchasing Contracts	oracle.apps.contracts.dash Board.textSearch.publicMo del.view.PoHeaderSearchV O	Manual
References	References	oracle.apps.sales.baseSales. references.publicViewEcsf. ReferencePVO	Manual
Sourcing Contracts	Sourcing Contracts	oracle.apps.contracts.dash Board.textSearch.publicMo del.view.AuctionHeaderSe archVO	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.TopicSearchPV O	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.prod uctModel.publicView Ecsf.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	ProjectBudgetSearch PVO	1	Daily		12:00 AM
Project Contract Invoices	ProjectContractInvoic e	ProjectContractInvoic ePVO	2	Weekly	Saturday	12:00 PM
Project Events	ProjectEvent	ProjectEventPVO	2	Weekly	Saturday	12:00 PM
Project Expenditure Items	ProjectExpenditureIt em	ProjectExpenditureIt emPVO	1	Daily		12:00 AM
Project Forecasts	ProjectForecastSearch	ProjectForecastSearch PVO	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.recei ving.publicViewEcsf. InboundShipmentPV O	Manual
Receipt Schedule	Receiving	Receiving	oracle.apps.scm.recei ving.publicViewEcsf. ReceiptPVO	Manual
Shipping Schedule	Shipping	Shipping	oracle.apps.scm.ship ping.shipConfirm.del iveries.publicViewEc sf.ShipmentPVO	Manual
Orchestration Orders Schedule	DOO	Orchestration Orders	oracle.apps.scm.doo. workbench.publicVie wEcsf.view.Orchestra tionOrderPVO	Manual
Cost Adjustments Schedule	Cost Management	Cost Accounting	oracle.apps.scm.costi ng.distProcessors.rev iewDistributions.pub licViewEcsf.CostAdju stmentsPVO	Manual
Cogs Transactions Schedule	Cost Management	Cost Accounting	oracle.apps.scm.costi ng.distProcessors.rev iewDistributions.pub licViewEcsf.CogsTran sactionPVO	Manual
Material Transactions Schedule	Cost Management	Cost Accounting	oracle.apps.scm.costi ng.distProcessors.rev iewDistributions.pub licViewEcsf.Material TransactionsPVO	Manual

Table C–9 (Cont.) Index Schedule for Oracle Fusion Supply Chain Management

Schedule Name	Product	Search Category	Business Object	Repeat Interval
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.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

Table C–10   (Cont.) Facets by Product Family

Product Family	Search Category	Facet
		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