

Oracle® E-Business Suite

Maintenance Guide

Release 12.2

Part No. E22954-13

July 2013

Oracle E-Business Suite Maintenance Guide, Release 12.2

Part No. E22954-13

Copyright © 1994, 2013, Oracle and/or its affiliates. All rights reserved.

Primary Author: Robert Farrington, Mildred Wang, Clara Jaeckel, Melody Yang

Contributing Author: Santiago Bastidas, Jason Brincat, George Buzsaki, Anne Carlson, Steve Carter, Steven Chan, Siu Chang, Ada Constanzo-Muller, Ivo Dujmovic, Mark Fisher, Paul Ferguson, Rajesh Ghosh, Kevin Hudson, Kunal Kapur, Ryan Landowski, Ruth Mamo, Ravi Mohan, Muhannad Obeidat, Gursat Olgun, Richard Ou, Venu Palakurthy, Justin Pfenning, Elke Phelps, Pranab Pradhan, Traci Short, Jan Smith, Vikas Soolapani, Seth Stafford, Susan Stratton, Leslie Studdard, Vani Subramanian, Venkat Vengala, Mark Warren, Aaron Weisberg, Sara Woodhull, Maxine Zasowski

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

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

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

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

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

U.S. GOVERNMENT END USERS: Oracle programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. No other rights are granted to the U.S. Government.

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

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

Contents

Send Us Your Comments

Preface

Part 1 Patching

1 Patching Overview

Release Maintenance: Patching Concepts.....	1-1
Patching Scope and Strategy.....	1-6
Introduction to Patching Utilities.....	1-7

2 Patching Utilities

Oracle Patch Application Assistant.....	2-1
The adop Utility.....	2-2
AD Merge Patch.....	2-43

3 Patching Procedures

Preparing for Patching.....	3-1
The Online Patching Cycle.....	3-6
Performing Interactive Patching.....	3-33
Performing Non-Interactive Patching.....	3-38
Patching HRMS Legislative Data.....	3-41
Patching NLS Systems.....	3-41
Keeping Your System Current.....	3-43
Analyzing Applied Patches.....	3-54

4 Patch Tracking Utilities

Patch Wizard.....	4-1
Register Flagged Files.....	4-22

5 Patch Reporting Utilities

Applied Patches.....	5-1
Timing Reports.....	5-15
Software Updates.....	5-31

Part 2 General Maintenance

6 Basic DBA Tasks

Overview of Oracle E-Business Suite DBA Duties.....	6-1
Resource Consumer Groups in Oracle E-Business Suite.....	6-3
Oracle E-Business Suite Password Management.....	6-4
ORACLE Users Window.....	6-15
Applications Window.....	6-18
Network Test Window.....	6-21
Administering Folders.....	6-23

7 Applications DBA System Maintenance Tasks and Tools

Overview.....	7-1
Managing Files.....	7-2
Adding NLS Languages.....	7-12
Maintaining the Database.....	7-14
Performing Maintenance Tasks Non-Interactively.....	7-19
Distribute Processing With Distributed AD.....	7-21
Managing Application Tier Services.....	7-22
Oracle E-Business Suite Maintenance Utilities.....	7-29
Command Line Utilities.....	7-37
Running AD Utilities.....	7-52
Using Parallel Processing.....	7-55
About System Maintenance.....	7-61
AD Administration Overview.....	7-62
Running AD Administration Interactively.....	7-65
Generating Applications Files.....	7-66
Maintaining Applications Files.....	7-69
Managing Database Entities.....	7-75

Using AD Relink.....	7-78
Relinking AD Executables.....	7-78
8 Applications DBA Reporting and Tracking Tasks	
Timing Information.....	8-1
General System Reporting.....	8-2
Oracle E-Business Suite Reporting Tools.....	8-3
AD Job Timing Report.....	8-4
AD Configuration Report.....	8-5
AD File Identification Report.....	8-6
AD Check Digest.....	8-6
9 Troubleshooting Applications DBA Operations	
Managing Worker Processes.....	9-1
Restarting Processes.....	9-7
Shutting Down and Restarting Managers.....	9-9
10 Technology Inventory Utility	
Technology Inventory Utility.....	10-1
11 Managing Oracle Fusion Middleware Log Files	
Collecting and Managing Log Files.....	11-1
12 Logging Features in Oracle E-Business Suite	
Overview.....	12-1
The Logging Framework.....	12-1
Guidelines for the Logging Framework.....	12-5
Log Files in Applied Technology Products.....	12-10
13 Running Diagnostics	
Diagnostics in Oracle E-Business Suite.....	13-1
Oracle Diagnostics Framework.....	13-1
Controlling Access to the Oracle Forms-based Applications Diagnostics Menu.....	13-1
14 Monitoring Oracle E-Business Suite with Oracle Applications Manager	
Overview of Monitoring.....	14-1
The Applications Dashboard.....	14-1
Additional Monitoring Features and Options.....	14-12

System Alerts, Metrics, and Logs.....	14-27
Diagnostics in Oracle Applications Manager.....	14-37
Support Cart.....	14-46
Oracle Applications Manager Log.....	14-47
Purging in Oracle Applications Manager.....	14-48

15 Oracle Workflow Manager

Oracle Workflow Manager Overview.....	15-1
Service Components.....	15-6
Notification Mailers.....	15-19
Agent Listeners.....	15-64
Java Agent Listeners.....	15-72
Web Services Outbound.....	15-80
Background Engines.....	15-86
Purging Workflow Data.....	15-90
Workflow Control Queue Cleanup.....	15-96
Active Work Items.....	15-98
Deferred Work Items.....	15-101
Suspended Work Items.....	15-104
Errored Work Items.....	15-107
Agents.....	15-110
Queue Propagation.....	15-114

16 License Manager

License Manager.....	16-1
----------------------	------

17 Functional Administrator and Functional Developer Tasks

Overview of Functional Administrator and Functional Developer Responsibilities.....	17-1
---	------

18 Using Alerts

Overview of Oracle Alert.....	18-1
Basic Business Needs.....	18-1
Oracle Alert Runtime Features.....	18-2
Alert Definitions.....	18-2
Predefined Alerts	18-3
Using Predefined Alerts.....	18-3
Customizing Predefined Alerts.....	18-4
Oracle Alert Precoded Alerts	18-6
Terms	18-7

Oracle Alert DBA Alerts	18-7
Applications DBA Alerts Descriptions	18-8
Oracle Alert Purging Alerts.....	18-10
Oracle Alert Purging Alerts Descriptions.....	18-11

19 Managing Query Optimization Statistics

Oracle E-Business Suite and Query Optimization.....	19-1
Gathering Statistics for the CBO.....	19-2
Gather Table Statistics.....	19-3
Backup Table Statistics.....	19-4
Restore Table Statistics.....	19-4
Gather Schema Statistics.....	19-5
Gather Column Statistics.....	19-7
Gather All Column Statistics.....	19-8
Purge FND_STATS History Records.....	19-8
FND_STATS Package.....	19-8

20 Using and Administering Oracle E-Business Suite Secure Enterprise Search

Overview of Oracle E-Business Suite Secure Enterprise Search.....	20-1
Using Oracle E-Business Suite Secure Enterprise Search.....	20-18
Performing Administrative and Setup Tasks.....	20-25
Creating a Search Administrator.....	20-25
Setting Up Oracle E-Business Suite Secure Enterprise Search.....	20-26
Setting Up Oracle E-Business Suite Secure Enterprise Search for Oracle SES Integration	20-31
Securing Searchable Objects.....	20-38
Role-Based Access Control (RBAC) Security.....	20-39
Search Security Plug-ins.....	20-43
User Authorization Cache.....	20-63
Administering Searchable Objects.....	20-67
Securing Searchable Objects Using Security Grants.....	20-67
Deploying Searchable Objects to Oracle SES.....	20-68
Administering Crawls in Oracle SES.....	20-72
Testing Oracle E-Business Suite Secure Enterprise Search Setups.....	20-77
Additional Administrative Tasks.....	20-78
Managing Crawling Schedules.....	20-79
Optimizing Indexes.....	20-81
Error Messages.....	20-82

21 Administering Process Navigation

Overview of Process Navigation.....	21-1
Modifying Your Menu.....	21-1
Creating Process Navigator Processes.....	21-2

Index

Send Us Your Comments

Oracle E-Business Suite Maintenance Guide, Release 12.2

Part No. E22954-13

Oracle welcomes customers' comments and suggestions on the quality and usefulness of this document. Your feedback is important, and helps us to best meet your needs as a user of our products. For example:

- Are the implementation steps correct and complete?
- Did you understand the context of the procedures?
- Did you find any errors in the information?
- Does the structure of the information help you with your tasks?
- Do you need different information or graphics? If so, where, and in what format?
- Are the examples correct? Do you need more examples?

If you find any errors or have any other suggestions for improvement, then please tell us your name, the name of the company who has licensed our products, the title and part number of the documentation and the chapter, section, and page number (if available).

Note: Before sending us your comments, you might like to check that you have the latest version of the document and if any concerns are already addressed. To do this, access the new Oracle E-Business Suite Release Online Documentation CD available on My Oracle Support and www.oracle.com. It contains the most current Documentation Library plus all documents revised or released recently.

Send your comments to us using the electronic mail address: appsdoc_us@oracle.com

Please give your name, address, electronic mail address, and telephone number (optional).

If you need assistance with Oracle software, then please contact your support representative or Oracle Support Services.

If you require training or instruction in using Oracle software, then please contact your Oracle local office and inquire about our Oracle University offerings. A list of Oracle offices is available on our Web site at www.oracle.com.

Preface

Intended Audience

Welcome to Release 12.2 of the *Oracle E-Business Suite Maintenance Guide*.

This guide assumes you have a working knowledge of the following:

- The principles and customary practices of your business area.
- Computer desktop application usage and terminology.

If you have never used Oracle E-Business Suite, we suggest you attend one or more of the Oracle E-Business Suite training classes available through Oracle University.

See Related Information Sources on page xii for more Oracle E-Business Suite product information.

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.

Structure

- 1 Patching Overview
- 2 Patching Utilities
- 3 Patching Procedures

- 4 Patch Tracking Utilities
- 5 Patch Reporting Utilities
- 6 Basic DBA Tasks
- 7 Applications DBA System Maintenance Tasks and Tools
- 8 Applications DBA Reporting and Tracking Tasks
- 9 Troubleshooting Applications DBA Operations
- 10 Technology Inventory Utility
- 11 Managing Oracle Fusion Middleware Log Files
- 12 Logging Features in Oracle E-Business Suite
- 13 Running Diagnostics
- 14 Monitoring Oracle E-Business Suite with Oracle Applications Manager
- 15 Oracle Workflow Manager
- 16 License Manager
- 17 Functional Administrator and Functional Developer Tasks
- 18 Using Alerts
- 19 Managing Query Optimization Statistics
- 20 Using and Administering Oracle E-Business Suite Secure Enterprise Search
- 21 Administering Process Navigation

Related Information Sources

This book is included in the Oracle E-Business Suite Documentation Library, which is supplied in the Release 12.2 Media Pack. If this guide refers you to other Oracle E-Business Suite documentation, use only the latest Release 12.2 versions of those guides.

Online Documentation

All Oracle E-Business Suite documentation is available online (HTML or PDF).

- **Online Help** - Online help patches (HTML) are available on My Oracle Support.
- **PDF Documentation** - See the Oracle E-Business Suite Documentation Library for current PDF documentation for your product with each release.
- **Release Notes** - For information about changes in this release, including new features, known issues, and other details, see the release notes for the relevant product, available on My Oracle Support.
- **Oracle Electronic Technical Reference Manual** - The Oracle Electronic Technical Reference Manual (eTRM) contains database diagrams and a detailed description of database tables, forms, reports, and programs for each Oracle E-Business Suite product. This information helps you convert data from your existing applications and integrate Oracle E-Business Suite data with non-Oracle applications, and write custom reports for Oracle E-Business Suite products. The Oracle eTRM is available on My Oracle Support.

Related Guides

You should have the following related books on hand. Depending on the requirements

of your particular installation, you may also need additional manuals or guides.

Oracle E-Business Suite Concepts

This book is intended for all those planning to deploy Oracle E-Business Suite Release 12.2, or contemplating significant changes to a configuration. After describing the Oracle E-Business Suite architecture and technology stack, it focuses on strategic topics, giving a broad outline of the actions needed to achieve a particular goal, plus the installation and configuration choices that may be available.

Oracle Alert User's Guide

This guide explains how to define periodic and event alerts to monitor the status of your Oracle E-Business Suite data.

Oracle Application Framework Personalization Guide

This guide covers the design-time and run-time aspects of personalizing applications built with Oracle Application Framework.

Oracle Diagnostics Framework User's Guide

This manual contains information on implementing and administering diagnostics tests for Oracle E-Business Suite using the Oracle Diagnostics Framework.

Oracle E-Business Suite CRM System Administrator's Guide

This manual describes how to implement the CRM Technology Foundation (JTT) and use its System Administrator Console.

Oracle E-Business Suite Developer's Guide

This guide contains the coding standards followed by the Oracle E-Business Suite development staff. It describes the Oracle Application Object Library components needed to implement the Oracle E-Business Suite user interface described in the *Oracle E-Business Suite User Interface Standards for Forms-Based Products*. It provides information to help you build your custom Oracle Forms Developer forms so that they integrate with Oracle E-Business Suite. In addition, this guide has information for customizations in features such as concurrent programs, flexfields, messages, and logging.

Oracle E-Business Suite Flexfields Guide

This guide provides flexfields planning, setup, and reference information for the Oracle E-Business Suite implementation team, as well as for users responsible for the ongoing maintenance of Oracle E-Business Suite product data. This guide also provides information on creating custom reports on flexfields data.

Oracle E-Business Suite Installation Guide: Using Rapid Install

This book is intended for use by anyone who is responsible for installing or upgrading Oracle E-Business Suite. It provides instructions for running Rapid Install either to carry out a fresh installation of Oracle E-Business Suite Release 12, or as part of an upgrade from Release 11*i* to Release 12. The book also describes the steps needed to install the technology stack components only, for the special situations where this is applicable.

Oracle E-Business Suite Integrated SOA Gateway User's Guide

This guide describes the high level service enablement process, explaining how users can browse and view the integration interface definitions and services residing in Oracle Integration Repository.

Oracle E-Business Suite Integrated SOA Gateway Implementation Guide

This guide explains how integration repository administrators can manage and administer the service enablement process (based on the service-oriented architecture) for both native packaged integration interfaces and composite services (BPEL type). It also describes how to invoke Web services from Oracle E-Business Suite by employing the Oracle Workflow Business Event System, how to manage Web service security, and how to monitor SOAP messages.

Oracle E-Business Suite Integrated SOA Gateway Developer's Guide

This guide describes how system integration developers can perform end-to-end service integration activities. These include orchestrating discrete Web services into meaningful end-to-end business processes using business process execution language (BPEL), and deploying BPEL processes at run time.

This guide also explains how to invoke Web services using the Service Invocation Framework. This includes defining Web service invocation metadata, invoking Web services, and testing the Web service invocation.

Oracle E-Business Suite Security Guide

This guide contains information on a comprehensive range of security-related topics, including access control, user management, function security, data security, and auditing. It also describes how Oracle E-Business Suite can be integrated into a single sign-on environment.

Oracle E-Business Suite Setup Guide

This guide contains information on system configuration tasks that are carried out either after installation or whenever there is a significant change to the system. The activities described include defining concurrent programs and managers, enabling Oracle Applications Manager features, and setting up printers and online help.

Oracle E-Business Suite User's Guide

This guide explains how to navigate, enter and query data, and run concurrent requests using the user interface (UI) of Oracle E-Business Suite. This guide also includes information on setting user profiles and customizing the UI.

Oracle E-Business Suite User Interface Standards for Forms-Based Products

This guide contains the user interface (UI) standards followed by the Oracle E-Business Suite development staff. It describes the UI for the Oracle E-Business Suite products and how to apply this UI to the design of an application built by using Oracle Forms.

Oracle Workflow Administrator's Guide

This guide explains how to complete the setup steps necessary for any product that includes workflow-enabled processes. It also describes how to manage workflow processes and business events using Oracle Applications Manager, how to monitor the

progress of runtime workflow processes, and how to administer notifications sent to workflow users.

Oracle Workflow Developer's Guide

This guide explains how to define new workflow business processes and customize existing Oracle E-Business Suite-embedded workflow processes. It also describes how to define and customize business events and event subscriptions.

Oracle Workflow User's Guide

This guide describes how users can view and respond to workflow notifications and monitor the progress of their workflow processes.

Oracle Workflow API Reference

This guide describes the APIs provided for developers and administrators to access Oracle Workflow.

Oracle Workflow Client Installation Guide

This guide describes how to install the Oracle Workflow Builder and Oracle XML Gateway Message Designer client components for Oracle E-Business Suite.

Oracle XML Gateway User's Guide

This guide describes Oracle XML Gateway functionality and each component of the Oracle XML Gateway architecture, including Message Designer, Oracle XML Gateway Setup, Execution Engine, Message Queues, and Oracle Transport Agent. It also explains how to use Collaboration History that records all business transactions and messages exchanged with trading partners.

The integrations with Oracle Workflow Business Event System, and the Business-to-Business transactions are also addressed in this guide.

Integration Repository

The Oracle Integration Repository is a compilation of information about the service endpoints exposed by the Oracle E-Business Suite of applications. It provides a complete catalog of Oracle E-Business Suite's business service interfaces. The tool lets users easily discover and deploy the appropriate business service interface for integration with any system, application, or business partner.

The Oracle Integration Repository is shipped as part of the E-Business Suite. As your instance is patched, the repository is automatically updated with content appropriate for the precise revisions of interfaces in your environment.

You can navigate to the Oracle Integration Repository through Oracle E-Business Suite Integrated SOA Gateway.

Do Not Use Database Tools to Modify Oracle E-Business Suite Data

Oracle **STRONGLY RECOMMENDS** that you never use SQL*Plus, Oracle Data

Browser, database triggers, or any other tool to modify Oracle E-Business Suite 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 E-Business Suite data, you risk destroying the integrity of your data and you lose the ability to audit changes to your data.

Because Oracle E-Business Suite tables are interrelated, any change you make using an Oracle E-Business Suite form can update many tables at once. But when you modify Oracle E-Business Suite data using anything other than Oracle E-Business Suite, 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 E-Business Suite.

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

Part 1

Patching

Patching Overview

Release Maintenance: Patching Concepts

Patches are applied throughout the life cycle of an Oracle E-Business Suite system. This maintenance may be necessary for a number of reasons, including:

- Fixing an existing issue
- Determining the cause of a new issue
- Adding a new feature or functionality
- Updating to a higher maintenance level
- Applying the latest product enhancements
- Applying online help
- Providing interoperability with new or modified technology stack components or versions

Depending on its type, a patch may update the Oracle E-Business Suite file system, the database, or both.

Note: Oracle E-Business Suite patches are available from My Oracle Support [<http://support.oracle.com>].

Patch File Structure

Patches generally consist of a top-level directory that may contain several files, and one or more subdirectories. The top-level directory is named `<patchnum>`, where `<patchnum>` is the number of the patch. The most important files in the top-level directory are: README.txt, README.html and the unified driver file (named `u<patchnum>.drv`).

Readme File

The README.txt or README.html file describes what the patch does. If the patch contains manual steps, the readme file provides information on using Oracle Patch Application Assistant (PAA) to generate customized installation instructions. If the patch does not contain manual steps, the readme file provides instructions for applying the patch using the adop utility.

Unified Driver File

The unified driver, named `u<patchnum>.drv`, contains the commands necessary to change files and database objects, and to generate new objects. It contains a sequential list of copy, database, and generate instructions, which are arranged in sections. You typically run the unified driver on all APPL_TOPs. The adop utility runs only the actions that are required for the current APPL_TOP. However, there may be scenarios where you run only the applicable portion of the driver. In these cases, the readme file directs you to run PAA to generate the specific instructions.

Patch Formats

Patch format describes the way the patch is packaged and applied. If a patch format is described as *cumulative*, that patch contains a consolidation of updates for a given codeline from the inception of a release, up to, and including, the latest release level. Oracle E-Business Suite patches are released in the following formats:

Oracle E-Business Suite Patch Formats

Patches	Description
Individual bug fix	A patch that fixes an existing issue.
Product family release update pack (product family RUP)	An aggregation of patches on a given codeline created for all products in specific product family for a specific point release. For example, R12.FIN_PF.A.1.
Release update pack (RUP)	A cumulative aggregation of product family release update packs on a given codeline created across Oracle E-Business Suite after the initial release. For example, 12.1.1.
Pre-upgrade patch	All <i>upgrade-related</i> , high-priority patches consolidated from all the products within a product family. Pre-upgrade patches are released as needed.

Patches	Description
Consolidated upgrade patch	All upgrade-related patches consolidated from all the products in a product family. These patches are released as needed and are only available for upgrading a Release 12 system from one point release to another.

Patch formats can additionally be identified as high-priority. This means that the patch has an impact that is broad enough to merit application by all customers who have installed the affected product


Note: You can find the latest available patches by logging on to My Oracle Support. Click the Patches and Updates tab and choose the Quick Links to the Latest Patchsets, Mini Packs, and Maintenance Packs link.

Codelines

In Release 12, Oracle E-Business Suite patches are grouped into codelines. A *codeline* begins with a point release (for example, Release 12.0) consisting of a unique set of product features, and progresses to include all the patches created to maintain that point release. The initial Release 12.0 point release introduced codeline A. Additional point releases introduce new codelines, each identified by a unique letter. For example, Release 12.1 introduced codeline B, and Release 12.2 introduces codeline C.

Important: This discussion of releases, codelines, and codelevels (including the diagrams used as examples), is intended solely to illustrate the concepts of codelines and codelevels. It does not represent any release commitment on the part of Oracle.

Codelines




Codelines and their associated codelevels ease the tracking of patch prerequisites, dependencies, and compatibilities.

Codelevels

Patches associated with codelines not only implement a set of product features for that point release, but also provide fixes to that set of features. This unique set of product features for a point release is referred to as a *codelevel*, and assigned a unique number.

The following diagram illustrates the relationship between codelines and codelevels in the context of Oracle E-Business Suite Release 12.

Codelevels

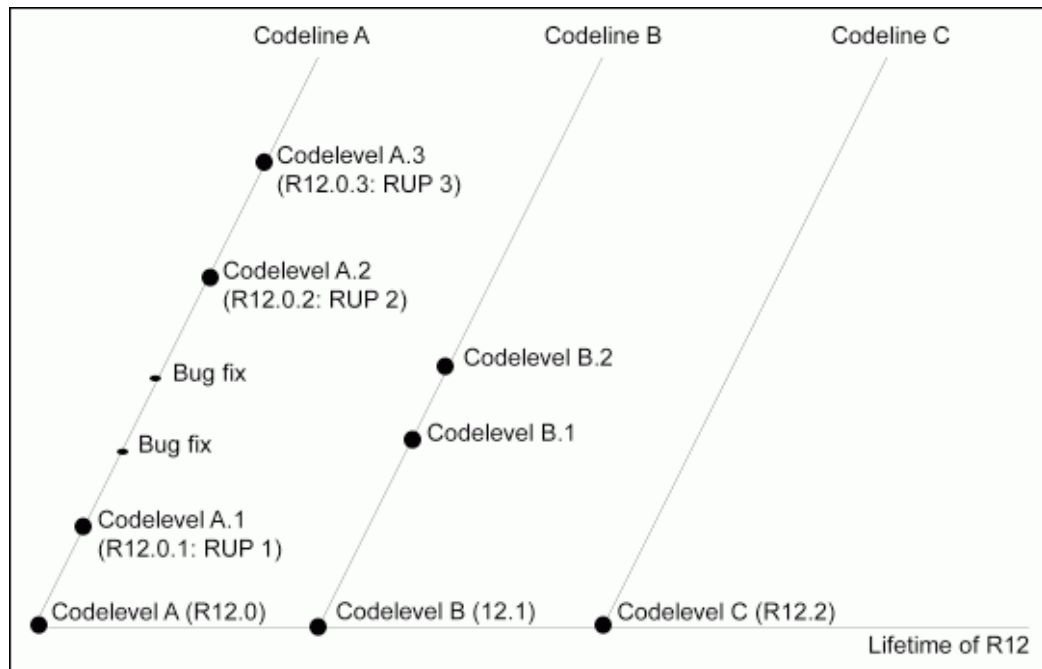


Further, codelevels identify patches for individual products. For example, if Oracle General Ledger (GL) is associated with your system, codelevel R12.GL.A.1 is the first set of fixes to codelevel R12.GL.A, R12.GL.A.2 is the second, and so on. Codelevels are cumulative - each one contains the initial set of features plus all the fixes created to date for that product or product family.

If individual bug fixes are created along the codeline, then subsequent codelevels will contain the bug fixes. For example, in the preceding diagram, the two bug fixes along codeline A will be included in codelevel A.2.

All codelevels created after the initial point release for a product or product family are aggregated into release update packs (RUPs). RUP1 is equivalent to R12.0.1, RUP2 is equivalent to R12.0.2, and so on. RUPs provide bug fixes. In addition, they may also provide feature enhancements, which provide new functionality that has a limited impact on your system.

Codelevels in RUPs



A new point release contains new features that will substantially impact your system and may change its operation. It starts a new codeline (for example, codeline B). At that point, you can choose to upgrade to the new codeline and adopt the new features, or stay on your existing codeline, where bug fixes and enhancements will continue to be provided for your existing features.

Warning: When applying individual bug fixes, make sure that you apply only the bug fixes for your existing codeline.

Note: For more information on determining the codeline and codelevel for each product installed in your system, see the Codelevels Summary page discussed in Codelevels Introduced by the Patch, page 4-17.

Patching Scope and Strategy

As described later in this book, patches to an Oracle E-Business Suite Release 12.2 system are applied online, while users are accessing the system. However, online patching does not support patching external systems: that is to say, software or metadata patches to external systems that are integrated with Oracle E-Business Suite cannot be applied as online patches.

This is because Oracle E-Business Suite Release 12.2 online patching applies exclusively

to the Oracle E-Business Suite Database, and usage of Edition-Based Redefinition (EBR) in other databases is transparent to Oracle E-Business Suite. Such usage includes schemas created using the Repository Creation Utility (RCU) in external databases for Oracle Fusion Middleware products that integrate with Oracle E-Business Suite. Among the products this applies to are Oracle WebCenter Portal, Oracle Access Manager, Oracle Internet Directory, and Oracle SOA Suite.

The same restriction exists for patches applied to external products that are integrated with Oracle E-Business Suite. Such products include Oracle Discoverer, SOA Suite/BPEL, OBIEE, ODI, Oracle EBS AccessGate, and any other Oracle Fusion Middleware products that are not embedded in the Oracle E-Business Suite technology stack. Finally, the restriction also applies to Oracle E-Business Suite patches that ship metadata to patch external systems.

When planning and executing a patching and maintenance strategy for an Oracle E-Business Suite system, you should perform the relevant tasks in the order shown below.

1. Apply the latest Oracle E-Business Suite Release Update Pack
2. Apply the latest Oracle E-Business Suite family packs, and all patches on the Recommended Patch List (ATG Release Update Packs, AutoConfig updates, and so on).
3. Upgrade all technology stack components to the latest certified levels.
4. Apply the latest Security Patch Updates (formerly known as Critical Patch Updates).
5. Apply the latest Oracle Database Patchset Updates (PSUs) and any concomitant Oracle E-Business Suite interoperability patches (which will differ for every site).
6. Apply specific interim patches (formerly known as one-off patches) if (and only if) it is not feasible to wait for the relevant fixes to be included in the release vehicles listed in points 1-4. This applies to both Oracle E-Business Suite and technology stack component patches.

Introduction to Patching Utilities

Patches are applied and tracked as needed by using one of the utilities designed specifically for that purpose. Some of these utilities are run from the command line, and others are Web-based. This section describes these utilities briefly. You can find a complete description of the utilities in later chapters of this book.

Command Line Patching Utilities

The utilities in this section are all run from the command line.

adop

adop is the utility you use to apply patches to the Oracle E-Business Suite file system or database.

Oracle E-Business Suite Release 12.2 introduces a new utility, *adop* (AD Online Patching).

AD Merge Patch (admrgpch)

When you apply patches individually, tasks common to all patches must be performed multiple times. Traditionally, the AD Merge Patch tool was used to merge multiple patches into a single patch, so that the common tasks only needed to be performed once.

In Oracle E-Business Suite Release 12.2, all the functionality of AD Merge Patch has been included in the *adop* patching tool. By default, *adop* will automatically merge all patches specified with the 'patches' parameter.

AD Merge Patch is still available for use if desired. If you want to do so, you should disable *adop*'s merging of patches, by adding 'merge=no' to the *adop* input file.

Patch Application Assistant (admsi.pl)

Oracle Patch Application Assistant (PAA) helps you track and perform manual steps during patching, and provides consistency in the format of manual steps. For patches that have manual steps, the patch readme file contains generic instructions for *all* systems. The readme file instructs you to use PAA to generate instructions specific to your system. For merged patches, PAA automatically combines the contents of the individual patch readme files into a single readme file.

PAA generates a custom set of instructions, specific to your installation, that consolidates and displays the relevant manual steps for all the patches you want to apply. After successfully performing each manual step, you can record that step as 'completed' in the PAA interface. Then, when applying patches in the future, you can refer to this record to see which steps you have already completed. Unless specified otherwise, you do not have to repeat the manual steps you have previously completed.

Web-Based Patching Utilities

The Web-based utilities described in this section are all accessed via Oracle Applications Manager. (OAM)

Applied Patches

Applied Patches enables you to query the patch history database for a list of patches that have been applied to your system. From the Applied Patches interface, you can view patch information such as patch number and type, driver file name, platform and version, location of applied patch, patch content and language, files changed or copied,

bug fixes in each driver file, whether patch application was successful, and timing information.

File History

With File History, you can view a history of the files that have been updated by a patch such as: APPL_TOP where the file resides, directory in which the file resides, product family that owns the file, name of the file, version of the file, date on which the file was changed, patch details report, and action summary report for the updates to the file.

Patch Wizard

An important part of the patching process is to keep abreast of new patches that are recommended, and analyze their effects before you actually apply them. With Patch Wizard, you can determine patches that have not been applied to your system, but are recommended to keep the system current. Patch Wizard also gives you a preview about the effects on your system of applying an individual patch.

Timing Reports

Timing Reports help you monitor a job that is running or provide you with a view of the statistics from completed adop and AD Administration maintenance sessions. You can view information such as task name, time taken to complete the task, start time and end time.

Register Flagged Files

Register Flagged Files provides a central register for your customizations. Use it to import, export, add, delete and view records of customized files. This utility replaces the applcust.txt file used in previous releases of Oracle E-Business Suite.

Software Updates

Software Updates is a portal from which you can view all the patching-related activities of your system.

Patching Utilities

Oracle Patch Application Assistant

For patches that have manual steps, the patch readme file instructs you to use Oracle Patch Application Assistant (PAA) by running the admsi.pl script. For merged patches, PAA automatically merges the contents of the individual patch readme files.

The Oracle Patch Application Assistant Interface

The Patch Application Assistant is started from the command line, and collects your input in a graphical user interface.

Running Oracle Patch Application Assistant

The following is a summary of the steps you use to run Patch Application Assistant. For a complete description of all the steps, see *Creating Customized Instructions for Patching Using PAA*, page 3-4.

Step 1: Set the environment

You must set the environment to apply the configuration parameters that define your system. This task is common to many AD utilities.

Step 2: Unzip the patch

Create a patch top directory, if it does not already exist. Download the patch into the patch top directory and unzip it.

Step 3: Review the information in the readme file

In the directory where you unzipped the patch, you will find a README.txt file and a README.html file. Review either of these files for information about the patch and for instructions on using Oracle Patch Application Assistant to generate customized

instructions for your system.

Step 4: Run Oracle Patch Application Assistant

Run PAA (admsi.pl) to generate customized instructions for your system. Follow the steps in the customized instructions to complete the patching process.

The adop Utility

You use the *adop* (AD Online Patching) utility to apply patches to the Oracle E-Business Suite file system or database. You can either allow *adop* to prompt for the information required to apply a patch, or enter the information without being prompted. Whichever method you choose, *adop* will then perform the tasks required to apply the patch:

- Reads patch metadata to determine patch dependencies and requirements
- Uploads patch information from a prior patch session to the database (if applicable)
- Reads and validate the patch driver file and reads the product driver files
- Compares version numbers of object modules from the product libraries and version numbers of the existing files against the patch files
- Backs up all existing files that will be changed by the patch
- Copies files
- Archive files in libraries
- Relinks executables
- Generates forms, reports, messages, graphics, and Java archive (JAR) files
- Compiles JSP files and invalid database objects
- Updates database objects
- Runs AutoConfig to update configuration files if any template files are introduced or updated by the patch
- Saves patch information to the database

Be aware of the following important points about *adop*:

- The *adop* utility always runs from the *run edition* file system. It automatically sets its environment correctly, regardless of the edition it is run from. Editions are described in more detail later.

- If a patch contains no new updates to files or database objects in your system, adop takes no action.
- If adop detects a previously failed patching session, it will attempt to recover that session.

adop Parameters

Run from the command line, adop accepts many parameters. Some are required, while others are optional. Some parameters override other parameters, and some have a higher order of precedence over others. All the parameters must be entered in `name=value` pairs.

adop Parameters

Parameter	Purpose	Values	Comments
phase	Used to tell adop which phases it is to run.	<ul style="list-style-type: none"> • prepare • apply • cutover • cleanup • finalize • actualize_all • fs_clone • abort 	<p>You can use a comma-separated list to specify multiple phases. For example, 'phase=prepare,apply'</p> <p>Note: Neither the abort nor fs_clone phases can be specified with any other phase.</p> <p>If you supply a phase other than those listed, the usage statement shown later will be printed and adop will exit.</p>

Parameter	Purpose	Values	Comments
loglevel	Used to specify the amount of information logged and displayed as adop performs its operations.	<ul style="list-style-type: none"> • STATEMENT • PROCEDURE • EVENT • WARNING • ERROR • UNEXPECTED 	<ul style="list-style-type: none"> • STATEMENT is only used for debugging. • PROCEDURE is only used for debugging high level procedures. • EVENT is used to display informational messages in normal processing. This is the default value. • WARNING is used to indicate an internal error that is handled by the system and does not affect processing. • ERROR indicates an action failed and will need to be reviewed by the user, but the system was able to continue processing. • UNEXPECTED indicates an unrecoverable error that halts processing and requires user intervention before processing can continue.

Parameter	Purpose	Values	Comments
cleanup_mode	Provides cleanup processing control.	<ul style="list-style-type: none"> full quick 	<p>cleanup_mode=full performs maximum cleanup, which includes dropping of covered objects and unused columns.</p> <p>cleanup_mode=quick performs minimum cleanup, and so requires least processing time. This is the default.</p>
finalize_mode	Used to specify whether the finalize or cutover phases should be performed in full mode or quick mode.	<ul style="list-style-type: none"> full quick 	<p>finalize_mode=full gathers statistics to help improve performance. Finalize will take about one hour longer if this mode is specified.</p> <p>finalize_mode=quick does not gather statistics, and therefore completes more quickly. This is the default.</p>
input_file	Used to specify the name of the input_file supplied to adop.	User-specified.	Must be an absolute file path.
runcontextfile	Used to specify the run context file supplied to adop	User-specified.	Must be an absolute file path.
patchcontextfile	Used to specify the patch context file supplied to adop	User-specified.	Must be an absolute file path.

Parameter	Purpose	Values	Comments
patches	Used to specify the patches adop is to apply.	<p>User-specified.</p> <p>There are two categories:</p> <ul style="list-style-type: none"> <p><i>Numbered only:</i> Most patches fall into this category.</p> <p>For example, you would specify patch number 111 as 'patches=111'. The directory and driver file will be inferred from the number.</p> <p><i>Containing a colon</i> : Should be used for HRMS Legislative Data patch (hrglobal.drv), online help patches, language-specific patches, and any other patches that are not all numbers.</p> <p>For example, you would specify the Korean language version of patch 222 as 'patches=222_KO:u222.drv'. Here, the directory is 222_KO, and the driver file to apply is u222.drv.</p> 	<p>You can use a comma-separated list to specify that multiple patches are to be applied in the same patching operation. The numbered-only and containing-a-colon categories of patch can be mixed.</p> <p>For example, you would specify patch number 111 and the Korean language version of patch 222 as 'patches=111,222_KO:u222.drv'.</p>

Parameter	Purpose	Values	Comments
workers	Used to specify the number of parallel workers to be employed.	User-specified integer.	<p>If you omit the 'workers' argument, a suitable number of workers will be chosen automatically.</p> <p>If you specify more workers than the machine can handle, adop will exit with an error.</p>
defaultsfile	Used to specify the path to the custom adop defaults file.	User-specified (but has a default value - see next column).	Must be an absolute file path. Defaults to \$TWO_TASK if not specified by user.
patchtop	Used to specify the location where the patches are unloaded.	User-specified (but has a default value - see next column).	Must be an absolute file path. Defaults to \$APPL_TOP_NE/EBS apps/patch if not specified by user.
merge	Used to merge multiple patches. You can merge the unified driver files into a single driver file that is passed to adop.	yes/no	Default value is 'no'.

Parameter	Purpose	Values	Comments
abandon	<p>Used to specify whether to restart the previous run of adop. May be useful if the previous action had an error.</p> <p>Note: If there was an error in the previous run, and 'abandon' is not set to 'yes', the same parameters will be re-used that were used in the failed run.</p>	yes/no	If you give a value for the 'restart' parameter, it cannot be the same as the value given for this parameter.
restart	<p>Used to specify whether to restart the previous run of adop. May be useful if the previous action had an error.</p>	yes/no	If you give a value for the 'abandon' parameter, it cannot be the same as the value given for this parameter.
action	<p>Used to specify whether to perform database actions.</p>	db nodb	<p>action=nodb - used to turn off database actions. This is useful if you are in a multi-node environment and adop has already updated the database, but you still need to update the file systems on the other nodes.</p> <p>action=db - must specified when allnodes=yes in a multi-node environment.</p>

Parameter	Purpose	Values	Comments
autoskip	Used to control whether the user is prompted about skipping actions in non-interactive patching. Allows operations to proceed with no user interaction. Particularly useful in multi-node systems where remote adop invocation means that user prompting is not possible.	yes/no	Default value is 'yes'. A report, ADZDPATCHSTAT.s ql, in the finalize phase gives the status of the patches that were applied in a particular session: SUCCESS, SKIPPED and SUCCESS, or SKIPPED and FAILED.
skipsyncerror	Enables the user to specify that any synchronization errors in the prepare phase are expected to be fixed automatically in the synchronization that takes place with subsequent patches.	yes/no	If the value of the parameter is passed as 'yes', the first patch to be synchronized will be done with the 'autoskip' flag set. It is the responsibility of the user to check the log files and correct any errors in the subsequent apply phase, or to confirm that synchronization with subsequent patches resolved the issue.
allnodes	Used to specify whether to run adop on all nodes or just a single node.	yes/no	Relevant to multi-node environments.
mtrestart	Used to specify whether to restart application tier services after cutover.	yes/no	Default value is 'yes'. If 'no' is specified, the services can later be restarted with the the adstrtal utility.

Parameter	Purpose	Values	Comments
allowcoredump	Used to specify that a core dump should be generated if adop crashes.	yes/no	Default value is 'no'. A value of 'yes' should be specified only if diagnostic information needs to be gathered.
analytics	Used to generate reports that can be helpful in debugging certain types of issue. Available with apply, finalize, cutover, and cleanup adop phases.	yes/no	Default value is 'no'. A value of 'yes' should be specified only if reports specifically need to be generated, because of the extra processing needed.

Online Help

To obtain help about the basics of adop operation, enter the command:

```
adop -help
```

This will display the following text:

Applications DBA Online Patching Tool (adop)

Usage: `adop [phase=<phase,phase,...>] [patches=<patch#,patch#,...>]`
`[<parameter>=<value> ...] [input_file=<filename>]`

Enter `adop -examples` for a detailed list of parameters and their usage.

See Oracle E-Business Suite Maintenance Guide for a full description of `adop` features, operation, and usage.

The `phase` parameter specifies the parts (phases) of the online patching cycle to be executed. The five standard phases are executed in the order shown below.

Standard phases:

<code>prepare</code>	- Prepare the instance for patch application.
<code>apply</code>	- Apply patches (to the patch edition).
<code>finalize</code>	- Ready the instance for cutover.
<code>cutover</code>	- Make the patch edition the new run edition.
<code>cleanup</code>	- Drop obsolete objects and data from old editions.

There are also three special phases, for use when needed.

Special phases:

<code>abort</code>	- Abort the current patching cycle.
<code>actualize_all</code>	- Create new copies of all code objects in the patch edition.
<code>fs_clone</code>	- Copy the run file system to the patch file system.

Phase-specific parameters control operation of a particular phase:

Apply parameters:

`patches=<patch#>[,<patch#>...]`

A single patch or comma-separated list of patches to apply.

This parameter is required when executing the `apply` phase.

For language patches, you must also specify the driver file:

`patches=10124646_AR:10124646.drv,10124646_KO:10124645.drv`

`restart=(yes|no) [default: no]`

Resume a failed apply action where processing left off.

`abandon=(yes|no) [default: no]`

Re-apply a failed patch from the beginning.

Finalize parameters:

`finalize_mode=(full|quick) [default: quick]`

Quick mode will provide the shortest execution time, by skipping non-essential additional actions.

Full mode performs additional actions such as gathering statistics, which may improve performance after cutover.

Cutover parameters:

`mtrestart=(yes|no) [default: yes].`

Specifies whether to restart application tier servers after cutover. Leave at default unless performing manual steps during downtime.

Cleanup parameters:

`cleanup_mode=(full|quick) [default: quick]`

Quick mode provides the shortest execution time, by skipping non-essential additional actions.

Full mode performs additional processing to remove all unused code, data, and old editions. May take a long time.

General parameters apply to all phases:

```
workers=<number> [default: computed]
    Number of parallel workers used to execute tasks.
    Default value is computed principally according to number of
    available CPU cores.
```

```
input_file=<file_name>
```

As well as being entered directly on the command line, adop parameters can be specified in a text file, with one <parameter>=<value> on each line of the file.

Examples:

```
phase=prepare,apply,finalize,cutover,cleanup
patches=123456
workers=4
```

Command line parameters override input file parameters.

```
loglevel=(statement|procedure|event|warning|error|unexpected)
    [default: event]
    Controls the level of diagnostic log detail displayed.
```

```
allnodes=(yes|no) [default: yes]
    Specifies whether actions should be executed on all
    application tier nodes of a multi-node system.
```

```
action=(db|nodb) [default: db]
    Specifies whether to execute database actions (as well as
    file system actions).
```

```
-status [<session_id>]
    Display status of the latest adop session, or a specified
    session.
```

```
-help
    This help screen.
```

```
-examples
    Additional help information with common usage examples.
```

Three examples (use adop -examples for more examples):

Complete patching cycle, running each phase separately:

```
adop phase=prepare
adop phase=apply patches=12345,67890 workers=4
adop phase=finalize workers=4
adop phase=cutover workers=4
adop phase=cleanup
```

Complete patching cycle, running all phases in a single adop command:

```
adop phase=prepare,apply,finalize,cutover,cleanup patches=12345
```

Complete patching cycle, specifying all parameters in an input_file:

```
adop input_file=adop2013_05_13.txt
```

adop exiting with status = 0

This help usage statement will also appear if you supply an invalid parameter on the adop command line.

Optionally, you can also display examples of the various adop parameters by entering

the command:

```
adop -examples
```

This will display the following text:

Applications DBA Online Patching Tool (adop)

Enter `adop -help` for usage syntax and an overview of phases and parameters.

See Oracle E-Business Suite Maintenance Guide for a full description of `adop` parameters and options.

Parameters relevant to all phases:

`workers` : Specifies number of parallel workers used to execute tasks.
Default value is computed, principally on the basis of available CPU cores.
If required, the number of workers can be specified explicitly.
Example:
`adop phase=prepare workers=8`

`input_file` : As well as being entered directly on the command line, parameters can be specified in a text file, with one `<parameter>=<value>` on each line of the file.
Example:
`phase=prepare,apply,finalize,cutover,cleanup`
`patches=123456`
`workers=4`
Command line parameters override input file parameters.

`loglevel` : Controls the level of detail displayed from the diagnostic log file on the console.
Takes values `(statement|procedure|event|warning|error|unexpected)` [default: `event`]
Examples:
`adop phase=prepare loglevel=statement`
Displays full details from the log file.
`adop phase=prepare loglevel=error`
Displays only errors from the log file.

`defaultsfile` : Path of the custom defaults file for use by `adop`.
Example:
`adop phase=apply`
`defaultsfile=adcustomdefaults.txt`
The `defaultsfile` should be in the `$APPL_TOP/admin/$TWO_TASK` directory on the patch file system, unless `adop` is being run in hotpatch mode. In this case, it should be in the same location on the run file system.

`allowcoredump` : Specifies whether `adop` should create a core dump if it crashes.
Takes values `'yes'` or `'no'` [default: `no`].
Example:
`adop phase=cutover allowcoredump=yes`
Only needs to be specified when debugging is required.

analytics : Takes values 'yes' or 'no' [default: no].
 Example:
 adop phase=finalize analytics=no [default]
 adop phase=finalize analytics=yes
 If analytics=yes, reports will be generated for the
 specified phase.

autoskip : Specifies whether adop should skip prompts and proceed
 further.
 Takes values 'yes' or 'no'. [default: no]

maxworkers : Specifies the maximum number of workers to be used to
 apply patches. The greater of the value specified and the
 computed value will be used.
 It This parameter can only be specified in an input file.
 cannot be specified on the command line.

Parameters relevant to prepare phase:

cleanup_mode : Specifies the mode for the cleanup that is implicitly
 executed as part of prepare.
 Quick cleanup performs minimum required processing for
 the patching cycle.
 Example:
 adop phase=prepare
 (same as adop phase=prepare cleanup_mode=quick)
 Full mode performs additional processing to remove all
 unused code, data, and old editions, but may take a long time.
 Example:
 adop phase=prepare cleanup_mode=full

skipsyncerror : Specifies whether synchronization of the first patch
 should be called with autoskip flag.
 Takes values 'yes' or 'no' [default: no].
 Example:
 adop phase=prepare skipsyncerror=yes

Parameters relevant to apply phase:

If the list of patches to be applied are provided to the apply phase of
 adop,
 adpatch will be invoked in non-interactive mode. Adpatch can be invoked
 interactively with the command
 adop phase=apply

To invoke adpatch non-interactively
 adop phase=apply patches=1234,5678 <other
 options>

patchtop : Path to the user-specified directory (on the

non-editioned file system) where patches are unzipped.
 Example:
 adop phase=apply patchtop=\$APPL_TOP_NE/./patch

hotpatch : Apply patches in hotpatch mode. Takes values 'yes' or 'no'
 [default: no].
 Example:
 adop phase=apply
 input_file=adopsession20130513.txt
 hotpatch=yes

merge : Merge patches before applying. Takes values 'yes' or 'no'
 [default: no].
 This replaces AD Merge Patch. Patches can be merged and applied to save time if the patches have some common files and actions.
 Examples:
 adop phase=apply patches=1234,5678 merge=yes
 adop phase=apply patches=1234,5678 merge=no
 [default]

abandon : Abandon a patching session. Takes values 'yes' or 'no' [default:none].

restart : Restart a patching session. Takes values 'yes' or 'no'. The abandon and restart parameters should be specified together, with one specified as 'no' and the other specified as 'yes'.
 Example of resuming a previous adop session, using the same parameter values:
 adop phase=apply abandon=no restart=yes
 Example of ignoring the previous failed adop session and starting a new one:
 adop phase=apply abandon=yes restart=no <other options>

options : Options that can be specified as a comma-separated list are as follows.
 Note that these options can be prefixed with "no", e.g. "nocheckfile".

checkfile - Skip running exec, SQL, and exectier commands if they are recorded as already [default: checkfile].

compiledb - Compile invalid objects in the database after running actions in the database [default: compiledb].

driver compilejsp - Sompile out-of-date JSP files, if the patch

has copy actions for at least one JSP file
[default: compilejsp].

copyportion - Run commands found in a copy driver
[default: copyportion].

databaseportion - Run commands found in a database driver
[default: databaseportion].

generateportion - Run commands found in a generate driver
[default: generateportion].

integrity - Perform patch integrity checking
[default: nointegrity].

autoconfig - Run AutoConfig [default: autoconfig].

actiondetails - Turn off display of action details
[default: actiondetails].

parallel - Run actions that update the database or
actions that generate files in parallel
[default: parallel].

prereq - Perform prerequisite patch checking prior
to
running patch driver files
[default: noprereq].

E-Business validate - Connect to all registered Oracle
application Suite schemas at the start of patch
[default: novalidate].

phtofile - Save patch history to file
[default: nophtofile].

forceapply - Reapply a patch that has already been
applied.
Useful in combination with "nocheckfile"
option
to rerun files that have already been
executed.

Examples:

adop phase=apply options=checkfile

This command shows an option that is the default, and therefore does not have to be specified.

adop phase=apply options=nocheckfile

This command can be used to execute sql, exec and exectier actions even if they have already been executed.

adop phase=apply options=forceapply,nodatabaseportion

This command shows the use of multiple options.

flags : Flags that can be specified as a comma-separated list
are as

follows.

Note that these flags can be prefixed with "no", e.g. "nologging".

file.	hidepw	- To omit the "HIDEPW:" comments in the log
		Default - hidepw.
file.	trace	- To log all database operations to a trace
		Default - notrace.
mode.	logging	- To create indexes in LOGGING or NOLOGGING
		Default - nologging.
some	autoskip	- To proceed with adpatch execution even if
		driver actions failed.
autoskip]		In non-interactive mode [default -
noautoskip]		In interactive mode [default -

The parameter can only be specified from the input file.

Example in input file:

```
flags=noautoskip,hidepw
```

This can be used to skip any failing driver actions and omit the hidepw comments in the log file.

Parameters relevant to finalize phase:

`finalize_mode` : Specifies mode for finalize phase. Takes values 'quick' and

'full'.

Quick mode provides the shortest execution time, by

skipping

non-essential additional actions.

Full mode performs additional actions such as gathering statistics, which may improve performance after

cutover.

Examples:

```
adop phase=finalize (assumes
finalize_mode=quick)
adop phase=finalize finalize_mode=quick
[default]
adop phase=finalize finalize_mode=full
```

Parameters relevant to cutover phase:

`finalize_mode` : Specifies the mode for the finalize that is implicitly performed during cutover.

Example of performing the minimum necessary actions:

```
adop phase=cutover
(same as adop phase=cutover
```

```
finalize_mode=quick)
```

Example of using full mode to perform additional

actions such

as gathering statistics:

```
adop phase=cutover finalize_mode=full
```

Full mode may improve performance after cutover.

`mtrestart` : Takes values 'yes' or 'no' [Default: no]
Whether to restart middle tier application servers
after
cutover. Leave at default unless performing manual
steps
during downtime.
Example:
If the user desires to perform some manual steps
during downtime the sequence of events would be
`adop phase=cutover mtrstart=no`
`<Perform downtime steps>`
`<Bring up services manually>`

Parameters relevant to cleanup phase:

`cleanup_mode` : Specifies the mode for the cleanup.
Quick cleanup performs minimum required processing for
the
patching cycle.
Example:
`adop phase=cleanup`
(same as `adop phase=cleanup cleanup_mode=quick`)
Full mode performs additional processing to remove all
unused
code, data, and old editions, but may take a long time.
Example:
`adop phase=cleanup cleanup_mode=full`

Parameters relevant to multinode instances:

`allnodes` : Takes values 'yes' or 'no' [default: no].
Can be used to perform actions in specific nodes in a
mutinode
instance
Example showing prepare actions on the current node
only:
`adop phase=prepare allnodes=no`
`action` : Takes values 'db' or 'nodb' [default: db].
Example showing only file system actions being run in
prepare
phase:
`adop phase=prepare action=nodb`

Example:

In a non-shared file system multi-node environment where tasks are to be
performed on different nodes separately:

On the first node to be patched, run the command:

```
adop phase=prepare allnodes=no
```

On all other nodes, run the command:

```
adop phase=prepare allnodes=no action=nodb
```

The same principle applies to the apply and cutover phases.

The other major phases (including `cleanup`, `actualize_all`, and `finalize`)
`all`
perform database actions.
Therefore, these phases only need to be performed on one node.

`adop` exiting with `status = 0`

The Input File

`adop` also accepts parameters in an *input file*. From the command line, you specify an input file by using the parameter `input_file=<myinput.txt>`, where `myinput.txt` is the name of your input file.

Input File Parameters

Note: You should always provide the full path to the input file.

Significant input file parameters include:

`patches`
`phase`
`patchtop`
`merge`
`defaultsfile`
`abandon`
`restart`
`workers`

To use these parameters, you must specify them in `name=value` pairs. For example, `phase=prepare, apply patches=12345,67890 workers=4`.

Note: If you supply a parameter to the input file twice (for example, `workers` is defined on both lines 2 and lines 5 of your input file), the last definition (in this example, on line 5) will be used.

The other arguments in the above list can be defined by using `options=<other_arguments>`. The `other_arguments` should be a comma separated list, and can be any combination of the following:

- `nodatabaseportion`
- `nogenerateportion`
- `hotpatch`

For example, `options=nodatabaseportion,nogenerateportion`.

The Defaults File

Parameters can also be passed to `adop` into `adop` through a *defaults file*. From the command line, you can specify a defaults file by using the parameter `defaultsfile=<mydefaults.txt>`, where `mydefaults.txt` is the name of your file.

Your own defaults file will be checked the validity of its contents, and if issues are found an error will be raised. If you do not specify a custom defaults file, adop will use the one that is automatically generated by the system (using AutoConfig).

If adop is being run in hotpatch mode, your own defaults file should be located on the *run* file system, under \$APPL_TOP/admin/\$TWO_TASK. Otherwise, the defaults file should be in the same location, but on the *patch* file system.

Note: Instead of using your own defaults file, it is generally preferable to supply your own parameters via the command line or in an input file. Parameters supplied in either of these ways take precedence over parameters in the the defaults file.

Parameters In Defaults File

Only one parameter, *patchtop*, can currently be defined in the defaultsfile. This parameter is used to specify the location where patches are unloaded. The default patchtop directory is on the non-editioned file system, at \$APPL_TOP_NE/EBSapps/patch.

If you wish to use the patchtop supplied in the defaults file, you must specify the defaults file as a parameter either on the command line or in the input file. If you do not specify the defaults file in one of these two locations, the file will not be read and the defaults file patchtop will not be used.

Order of Parameters

As described above, most parameters can be defined in at least two locations, with patchtop able to be defined in three different locations. If multiple different definitions are specified, the following order is used.

1. **Command Line:** An adop parameter specified on the command line will take precedence over all others.
2. **Input File:** An adop parameter given here will only be lower in precedence to a parameter specified on the command line.
3. **Defaults File:** Parameters defined here have the lowest level of precedence. A parameter defined here must not also be given either in the input file or on the command line if it is to be used.

Important: Because you can supply higher-priority parameters on the command line and in the input file, you should never need to edit the defaults file.

Preparing your System for Patching

Before you begin a patching cycle, there are some important tasks you need to

complete.

In Release 12.2, it is more appropriate to think in terms of a *patching cycle* than a single patching operation. The online patching cycle consists of a number of phases:

1. Prepare
2. Apply
3. Finalize
4. Cutover
5. Cleanup

First, you must set the environment by executing (sourcing) the run file system environment file:

```
$ source <run APPL_TOP path>/APPS<CONTEXT_NAME>.env
```

For more information, see *Setting the Environment in Running AD Utilities*, page 7-52.

You are now ready to invoke the adop utility, specifying the phase or phases you wish to run. The actions taken in these phases are described in *Oracle E-Business Suite Concepts*. This and the next chapter of this book provide details of the available options.

Summary of Fundamental Patching Operations

Patching is performed in several *phases*, which are specified on the adop command line as follows:

```
adop phase=<phase_name>
```

Prepare phase - The following example shows interactive mode, which prompts for inputs:

```
$ adop phase=prepare
```

Apply phase - The following example shows non-interactive mode, which uses an `input_file`:

```
$ adop phase=apply input_file=<yourfilename>.txt
```

The `input_file` is used to pass arguments such as the following:

```
patches=<patch number>  
patchtop=<location where the patches were downloaded and unzipped>  
workers=<number of workers>
```

Finalize phase - The following example shows the finalize command:

```
$ adop phase=finalize
```

Cutover phase - The following example shows the cutover command:

```
$ adop phase=cutover
```

Cleanup phase - The following example shows the cleanup command:

```
$ adop phase=cleanup
```

Abort phase - The following example shows the abort command:

```
$ adop phase=abort
```

Important: The abort phase can be run after either the prepare or apply phases have been run, but not after the cutover phase.

The adop phases are described in more detail in The Online Patching Cycle, page 3-6 section of the Patching Procedures chapter.

Patch Log Files

It is advisable to review the relevant log files after any patching operation. The adop log files are located on the non-edited file system (`fs_ne`), under:

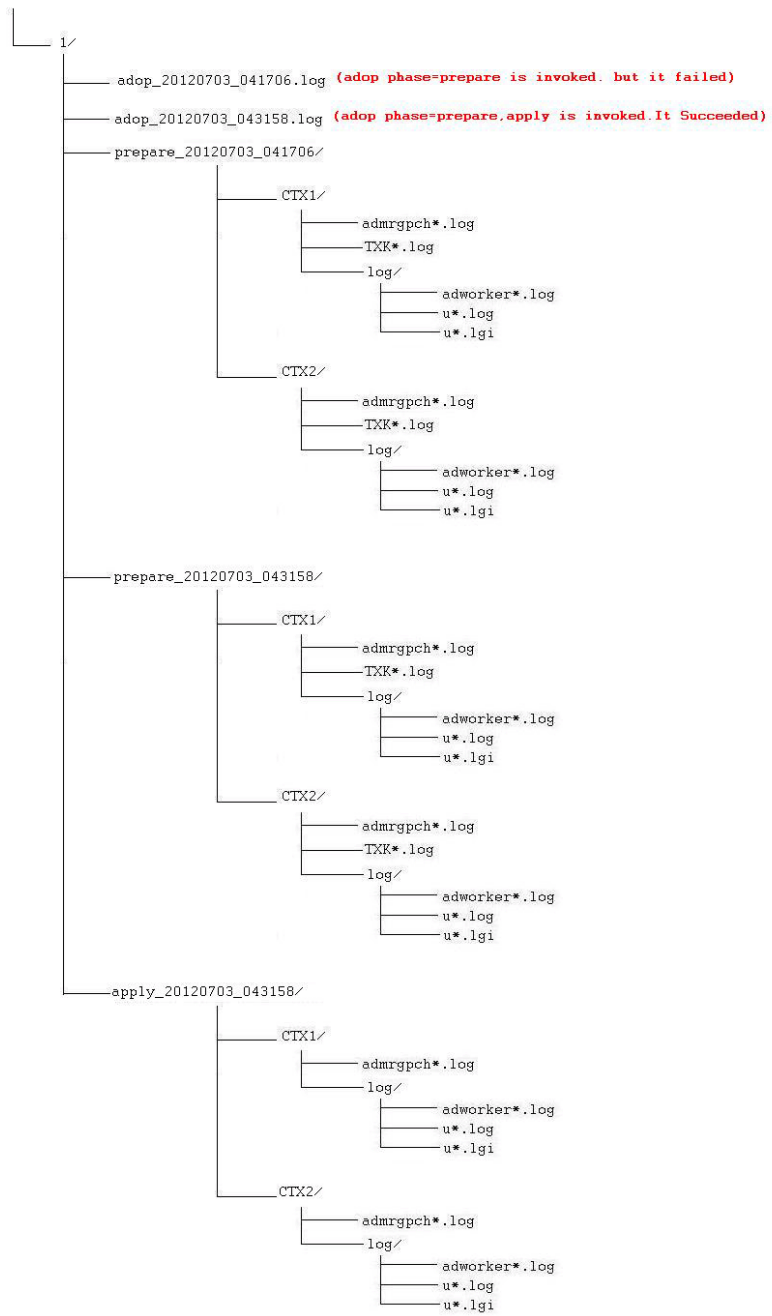
```
s_ne_base/EBSapps/log/adop/<adop_session_id>/<phase>_<date>_<time>/<context_name>/log
```

For example, if `s_ne_base` was `/u01/R122_EBS/fs_ne`, the session ID was 15, and the `<CONTEXT_NAME>` was `patch01_testsys`, the path to the adop log files from 9th June 2013 would resemble this:

```
/u01/R122_EBS/fs_ne/EBSapps/log/adop/15/apply_20130609_112226/patch01_testsys/log
```

Patch Log File Directory Structure

/s_ne_base/EBSapps/log/adop



Note: There is no environment variable for `s_ne_base`. The closest is `$APPL_TOP_NE`, which will have a value such as `/u01/R122_EBS/fs_ne/EBSapps/appl`.

Other log files are created for specific purposes, for example, to record all the actions associated with parallel workers. These worker log files are written to the non-editioned file system, under `EBSapps/log/adop/<adop_session_id>/<phase_timestamp>`. For example, `/u01/R122_EBS/fs_ne/EBSapps/log/adop`. Review these files when the patching session is complete.

Other AD log files include those shown in the following table:

Non-adop AD Log Files

Log File	Used For
<code>adrelink.log</code>	Relinking
<code>adlibin.log</code>	Moving C object files into the C library of a product
<code>adlibout.log</code>	Moving C object files out of the C library of a product
<code>adworkxxx.log</code>	Database operations run in parallel
<code><language>_<filename>_ldt.log</code>	Seed data loader files

If `adop` does not perform an action, it does not generate the log file associated with that type of action.

Note: You can also review log files using the View Log Files feature of OAM Timing Reports. See: View Log Files, page 5-28.

JAR File List

In the apply phase of an online patching cycle, `adop` creates a file called `jarlist.txt`. This file is provided in case you wish to perform your own JAR file signing using a very secure certificate. In such cases, you will need to specify the `adop` command line parameter `option=nojarsigning` in order to bypass the standard JAR file signing activity that will otherwise be performed by AD.

The `jarlist.txt` file is placed in the same directory as the patch log file. The

following example is for patch 13358502, which was applied during a patching session that had ID=14:

```
$APPL_TOP_NE/./log/adop/14/apply_20130515_125116/testenv_sys322  
0410/13358502/log/jarlist.txt.
```

Prompts

In addition to the standard prompts common to most AD utilities, adop also asks for information specific to the patching process. You must respond to all the prompts for each driver you run.

Important: You can only run one session of adop at a time.

Passwords Required

When run interactively, adop prompts for the Oracle Application Object Library (APPS), SYSTEM, and WLSADMIN passwords.

Note: You can change this behavior by adding `options=validate` to the `input_file`. See *Command Line Arguments*, page 2-33.

Patch Directory (Patchtop)

In interactive patching, adop asks you to specify the directory where the patch files have been unzipped. This is sometimes referred to as the *patchtop*. The default is the directory from which you started adop.

If specifying your own choice of directory, you must supply the full path, and the operating system user that is running adop must have write permissions to that directory.

Note: Oracle recommends using `fs_ne/EBSapps/patch` (i.e. on the non-edited file system) as the patchtop directory. For example, `/u01/R122_EBS/fs_ne/EBSapps/patch`.

If you have a multi-node environment, you must download and unzip the patches (under `$APPL_TOP_NE/EBSapps/patch`) on the respective nodes.

If you want to merge patches before applying them, you need to download and unzip all the individual patches in the same location as that of the merged driver file. For example, if you merge patches 111, 222, and 333 (using AD Merge Patch), and the merged patch driver file location is

`$APPL_TOP_NE/EBSapps/patch/mergetest/mergetest.drv`, you should then download and unzip the individual patches as

```
$APPL_TOP_NE/EBSapps/patch/mergetest/111,  
$APPL_TOP_NE/EBSapps/patch/mergetest/222, and
```

\$APPL_TOP_NE/EBSapps/patch/mergetest/333.

Patch Driver File

adop prompts for the name of the patch driver file. By default, it does not check the integrity of the patch - that is, whether the version of each file referenced in a driver file copy action matches the version in the patch - because Oracle E-Business Suite patches are always tested before release to ensure they contain the correct files.

The unified driver, named `u<patchnum>.drv`, contains the commands necessary to change files and database objects, and to generate new objects. It contains copy, database, and generate portions and performs the copy, database, and generate actions in the stated order. You typically run the unified driver on all APPL_TOPs and adop runs only the actions that are required for the current APPL_TOP. However, there may be cases where you run only the applicable portion of the driver.

Copy Portion of a Unified Driver

When the copy portion of a unified driver runs, adop performs the following actions:

- Extracts the appropriate files from the C library of each product.
- Compares the extracted object modules with their corresponding files in the patch directory. It also makes this type of comparison with files such as forms, reports, and SQL scripts.
- Backs up any product file with a more recent version in the patch directory to a subdirectory in the patch directory. For example, if `<patch_dir>` is the patch directory, `<system_name>` is the applications system name, `<appl_top_name>` is the APPL_TOP name, and `<prod>` is the name of the product being patched, it backs up:

```
<PROD>_TOP/<subdir(s)>/<old_file_name>
```

to

```
<patch_dir>/backup/<system_name>/<appl_top_name>/<prod>/<subdir(s)>/  
<old_file_name>
```

Note: The Applications system name and the APPL_TOP name are determined during the Rapid Install process.

- Replaces the outdated files of each product with newer files from the patch directory.
- Loads the new object modules into the C libraries.
- Relinks the Oracle E-Business Suite products with the operating system, Oracle server, and other Oracle products libraries.
- Applies changed Java class files and regenerates JAR files as needed.

- Copies any specified HTML or media files to their respective destinations.
- Compiles out-of-date Java Server Page (JSP) files (if any JSP files are included in the patch).

Database Portion of a Unified Driver

When the database portion of a driver runs, adop performs these actions:

- Gets a list of current invalid objects in the APPS schema.
- Determines whether the action was performed in a previous patch.
- Runs SQL scripts and EXEC commands, which change Oracle E-Business Suite database objects. By default, adop runs scripts and commands in parallel.
- Compiles invalid objects in the database.
- Assembles a list of current invalid objects in the APPS schema.

Note: As of Release 12, a separate MRC schema is not required, so Invoker's Rights processing (included in previous releases) has been removed.

Generate Portion of a Unified Driver

Apply the generate portion of a unified driver on all APPL_TOP directories containing one or more files being generated by the patch. When the generate portion of a driver runs, adop performs these actions:

- Generates Oracle Forms PL/SQL library files
- Generates Oracle Forms menu files
- Generates Oracle Forms executable files
- Generates Oracle Reports PL/SQL library files
- Generates Oracle Reports files
- Generates message files
- Generates Oracle Workflow resource files

Number of Parallel Workers

By default, adop runs database updates and file generation commands in parallel and prompts you for the number of workers. Tasks are assigned to workers, the workers run

the tasks to completion, and adop assigns new tasks.

The default value for the number of workers is twice the number of CPUs on the node from which you run adop. After you specify the number of workers, adop displays messages similar to the following as it begins to update the Oracle E-Business Suite products:

```
Performing version checking for driver files...
Copying driver files into installation area...
Determining valid on-site files...
Screening out files not valid for this installation...
Extracting object modules from product libraries...
Performing version checking...
Determining what executables to link...
Determining what Oracle Forms files to generate...
Determining what Oracle Reports libraries to generate...
Determining what Oracle Reports files to generate...
```

adop runs (adop) all database actions based on *phase order*, a grouping of actions in the database portion of the patch that minimizes dependencies. This order is not necessarily the order in which the commands are listed in the database portion of the patch driver.

Note: For more information, see Using Parallel Processing, page 7-55 in the Maintenance section of this book.

Customized Files

adop reviews the AD_FILES table to determine if any customized files (Register Flagged Files) will be replaced by the patch. If so, it displays a message listing the customized files it will replace.

Note: For more information, see Customization Standards, *Oracle E-Business Suite Developer's Guide*, and Register Flagged Files, page 1-9.

NLS

If the patch you are applying has an NLS-related version, and if you are an NLS customer, adop prompts you about the NLS-related version of the patch before allowing you to continue.

Preparing for Non-Interactive Patching

Non-interactive patching is a way to save time by avoiding some of the prompts and automating the patching process. It can be used with all the major phases of adop, including the apply phase. The only new parameter that can be specified in Release 12.2 is 'patchtop'.

You tell adop to run non-interactively by providing the location of the input_file as an argument. The input_file specifies the defaults file (which is generated automatically), as well as other options that would have been placed on the adpatch command line in

previous releases.

After the patching actions are complete, you perform any post-patching steps listed in the patch readme file. See *Performing Non-Interactive Patching*, page 3-38.

Messages

adop generates several types of messages. Each message is recorded in a log file. See *Log Files*, page 2-23 for a list and descriptions.

Informational Messages

Informational messages are written to the informational message file. This log file uses the same base file name as the main adop log file, but substitutes a .lgi extension for the .log extension. For example, if the adop log file is named u1234567.log, the adop informational log file is named u1234567.lgi.

For example, adop writes information pertaining to the files not updated because they are up-to-date in the informational log file.

```
File will not be copied to destination.
```

```
Version check:
```

```
/slot03/appmgr/prodappl/ad/12.2/xml/oam/patch/history/SearchFiles.uix  
version is equal to or lower than  
/slot03/appmgr/prodcomn/html/oam/patch/history/SearchFiles.uix.  
File will not be copied to destination.
```

```
Version check:
```

```
/slot03/appmgr/prodappl/ad/12.2/xml/oam/patch/history/SearchFilesCriteri  
aAdvanced.uix  
version is equal to or lower than
```

```
/slot03/appmgr/prodcomn/html/oam/patch/history/SearchFilesCriteriaAdvanc  
ed.uix
```

Error Messages

When adop is using parallel processing and an error occurs, the job fails. Review the main adop log file and the adworkxxx.log file to determine the source of the error, resolve the issues and continue. Restart adop using the adctrl command.

Note: See *Monitoring and Controlling Parallel Processing*, for details on using the adctrl command.

If you cannot resolve the issue, you must:

- Verify that all steps in the readme file were completed.
- Check My Oracle Support for additional information regarding the patch you are applying.

If the message indicates that a worker has failed its job, you can fix the problem and restart the worker while the manager is running. Some failed jobs are deferred (not immediately reassigned) by the manager. These jobs do not cause the manager or other workers to stop.

See: Managing Worker Processes, page 9-1 in this book.

Successful Completion Message

adop displays messages such as the following when processing is complete. If you do not see a completion message, you should investigate and identify the reason.

```
A job timing report has been generated for the current session.
You should check the file
    /slot03/appmgr/prodappl/admin/PROD/out/adt323790.lst
for details.
```

```
Purging timing information for prior sessions.
```

```
sqlplus -s APPS/*****
@/slot03/appmgr/prodappl/ad/12.0.0/sql/adtprurge.sql 10 1000
```

```
Done purging timing information for prior sessions.
```

```
adop is complete.
```

```
adop may have written informational messages to the file
/slot03/appmgr/prodappl/admin/PROD/log/adpatch.lgi
```

```
Errors and warnings are listed in the log file
/slot03/appmgr/prodappl/admin/PROD/log/adpatch.log
```

```
and in other log files in the same directory.
```

Backup Directory

When adop runs, a backup directory is created in the directory where you unzip the patch. The old version of each file updated by the patch is copied into the backup directory. When applying large patches (such as release update packs, product family RUPs, and pre-upgrade patches), ensure there is enough disk space on the system where you unzip the patch, or the patching process may fail. We recommend having at least twice the amount of disk space as the unzipped patch file uses.

Tip: When there is no patching cycle in progress, you can if desired delete the files in the backup directory to free the space.

adop Modes

adop can apply patches in two specialized modes: pre-install and test. The patch readme file instructs you when it is necessary to use either of these modes.

Pre-Install Mode

Pre-install mode is generally used during the upgrade process to update AD utilities, apply pre-upgrade patches, or work around other patching issues. adop asks all startup questions except those relating to the database.

Warning: Run adop in pre-install mode *only* if the patch readme instructs you to do so.

To run adop in pre-install mode, include *preinstall=y* on the adop command line. It performs the following actions:

- Compares version numbers
- Copies files
- Relinks FND and AD executables
- Saves patch information to the file system

Because adop does not read driver files in pre-install mode, it copies all product files in the patch to the APPL_TOP directory. Additionally, even if a file in the patch should be both in the APPL_TOP and in another directory (such as in \$OA_HTML), adop copies the file only to the APPL_TOP.

In preinstall mode, adop validates codelevels against the files Preinstall_Codelevel_AD.txt and Preinstall_Codelevel_MP.txt. These files are located in the \$APPL_TOP/admin directory, and contain codelevel information about AD and other products registered in the database tables.

Since no database connection is available in pre-install mode, adop tries to validate whether the current patch should be applied based on the codelevel information in these two files, as follows:

- If Preinstall_Codelevel_AD.txt is missing from the APPL_TOP, adop will apply the patch in pre-install mode without validating the patch for codelevel compatibility.
- If Preinstall_Codelevel_MP.txt is missing from the APPL_TOP, adop will proceed with patch application without validating the patch for codelevel compatibility of the entities.
- If both files are missing, adop will not validate codelevels in pre-install mode.

Note the following restrictions when applying a patch in pre-install mode:

- NLS patches cannot be applied on the instance.
- Baseline or codelevel-introducing patches cannot be applied on the instance.

- adop will not check to see if the patch is already applied on the system.

Test Mode

In test mode, adop does not apply the patch. Instead, it lists each file it would have copied, relinked, executed, or generated and shows exactly what actions it would have performed had it applied the patch. It also runs AutoConfig in test mode to determine any impending changes to the configuration files. This allows you to see the effects of the patch on your production system before you apply it.

To run adop in test mode, include *apply=no* in the `input_file`. This runs as if adop is applying the patch, except it does not. It only performs the following actions:

- Copies any files from the patch directory to the Oracle E-Business Suite file system
- Archives any object modules into the product libraries
- Relinks any executables
- Generates any forms, reports, PL/SQL libraries, or menu files
- Runs any SQL or EXEC commands (commands that change the database)
- Instantiates new configuration files
- Updates the patch information files
- Updates patch information and release version in the database

See: Testing a Patch Before Applying It, page 3-37.

Arguments

At present, you can adjust the way adop operates by adding modifiers to the `input_file`, which is supplied as a command line argument to adop in non-interactive patching:

```
$ adop phase=apply input_file=<input_file.txt>
```

Important: If patching interactively (with the command `adop phase=apply`), the traditional command line arguments such as `flags` are currently not supported. You must pass them via the `input_file`.

The arguments or options added to the `input_file` are in the "token=value" format, where *token* is the name of the modifier. You are advised to enter both the argument and the value in lower case: the "token" portion is converted to lowercase, but not the "value".

In the following example:

LOGLEVEL=WARNING

The token ("LOGLEVEL") will be converted to lowercase, but the value (WARNING) will not be recognized by the utility. The correct way to enter this command is:

```
loglevel=warning
```

You can enter more than one token=value argument on a single command line by separating them with a single space, as in the following example.

```
printdebug=y flags=hidepw
```

In some cases, you can include more than one value for a token. When doing so, you separate the values with commas and no spaces. For example:

```
flags=nohidepw,trace
```

is valid, but

```
flags=nohidepw, trace
```

is not valid.

The following *arguments* are specific to *adop*, and can be used to modify and refine its behavior. The default value is used if you do not specify a value.

Important: These arguments must be passed to *adop* as entries in the `input_file` that is specified on the command line. This is in contrast to earlier releases, where they were passed to *adpatch* directly on the command line.

adop Arguments

Argument	Description
apply	<p>Purpose: Tells <i>adop</i> whether to run in test mode.</p> <p>Values: y, meaning that <i>adop</i> does not run in test mode; n, meaning that <i>adop</i> does run in test mode.</p> <p>Default: y</p> <p>Example: apply=n</p>

Argument	Description
driver	<p>Purpose: Tells adop the name of the patch driver file. This is usually used during non-interactive processing. It is only valid when the patchtop option is also used.</p> <p>Values: A driver file name, or comma-separated list of patch driver file names.</p> <p>Default: None, meaning that adop prompts for the patch driver file name.</p> <p>Example: patchtop=/d01/prodappl/patches/1234567 driver=u1234567.drv</p>
patchtop	<p>Purpose: Tells adop the location where the patch was unloaded.</p> <p>Values: A fully qualified directory name on the non-edited file system.</p> <p>Default: \$APPL_TOP_NE/EBS apps/patch</p>
preinstall	<p>Purpose: Tells adop whether to run in pre-install mode. Pre-install mode is used to update AD utilities before an upgrade and to apply pre-upgrade patches.</p> <p>Values: y, meaning that adop does run in pre-install mode; n, meaning that adop does not run in pre-install mode.</p> <p>Default: n</p> <p>Example: preinstall=y</p>

Argument	Description
uploadph	<p>Purpose: Tells adop to upload patch history information from the patch information files (in \$APPL_TOP/admin/\$TWO_TASK) to the database. adop exits after uploading the patch history information.</p> <p>Values: y, meaning that adop uploads patch history information; n, meaning that adop does not upload patch history information.</p> <p>Default: none</p> <p>Example: ploadph=y</p>

adop Options

The *options=* argument is used to pass generic options to adop. It takes the form of a comma-separated list. Enter one option or a comma-separated list of options. For example, *adop options=nocopyportion,nogenerateportion*.

Note: As with adop arguments, there must be no space after the comma.

adop Options

Option	Description
actiondetails	<p>Purpose: Tells adop to print details of actions taken in a patching operation.</p> <p>Default: actiondetails</p> <p>Specify <i>options=noactiondetails</i> if you do not want the details to be printed.</p> <p>Comments: It is generally preferable to accept the default, as the action details can be useful in understanding adop operations and diagnosing issues.</p>

Option	Description
autoconfig	<p data-bbox="959 310 1373 369">Purpose: Tells adop to run AutoConfig automatically.</p> <p data-bbox="959 396 1162 424">Default: autoconfig</p> <p data-bbox="959 451 1463 573">Use <i>options=noautoconfig</i> if you are applying a number of patches in sequence and want to run AutoConfig once, after applying the last patch of the sequence.</p> <p data-bbox="959 600 1463 722">Comments: The dual file system in Release 12.2 means that there is no need to shut down application tier services before running AutoConfig.</p>
checkfile	<p data-bbox="959 768 1446 890">Purpose: Tells adop to either skip running EXEC, SQL, and EXECTIER commands if they are recorded as already run, or to record them as having run after running them.</p> <p data-bbox="959 917 1146 945">Default: checkfile</p> <p data-bbox="959 972 1446 1031">Use <i>options=nocheckfile</i> to turn off the checkfile feature.</p> <p data-bbox="959 1058 1349 1117">Comments: Using checkfile provides significant performance benefits.</p> <p data-bbox="959 1144 1446 1430">If you are reapplying a patch with <i>options=nocheckfile</i> and the patch contains a call for a seed table upgrade (AD_ZD_SEED.UPGRADE), you may receive the error message "ORA-20001: Cannot upgrade existing table from Patch Edition". You can safely ignore this message if the table already contains the column 'ZD_EDITION_NAME'.</p>

Option	Description
compiledb	<p>Purpose: Tells adop to automatically compile invalid objects in the database after running actions normally found in the database portion of the driver.</p> <p>Default: compiledb for standard patches. nocompiledb for standard patch translations, documentation patches, and documentation patch translations.</p> <p>Use <i>options=nocompiledb</i> to save time when multiple non-merged patches are applied in the same patching operation.</p> <p>Comments: Merging multiple patches and applying a single merged patch is usually a better strategy.</p>
compilejsp	<p>Purpose: Tells adop whether to automatically compile out-of-date JSP files. JSP files are only compiled if the patch contains copy actions for at least one JSP file.</p> <p>Default: compilejsp for standard patches. nocompilejsp for standard patch translations, documentation patches, and documentation patch translations.</p> <p>Use <i>options=nocompilejsp</i> to save time when multiple non-merged patches are applied in a maintenance window.</p> <p>Comments: Merging multiple patches and applying a single merged patch is usually a better strategy.</p>
copyportion	<p>Purpose: Tells adop whether to run commands normally found in the copy portion of the driver.</p> <p>Default: copyportion</p> <p>Comments: Use <i>options=nocopyportion</i> to tell adop not to perform copy actions of the driver.</p>

Option	Description
databaseportion	<p>Purpose: Tells adop whether to run commands normally found in the database portion of the driver.</p> <p>Default: databaseportion</p> <p>Comments: Use <i>options=nodatabaseportion</i> to tell adop not to perform database-related driver actions.</p>
forceapply	<p>Purpose: Tells adop to reapply a patch that has already been applied.</p> <p>Default: noforceapply</p> <p>Comments: Use the nocheckfile option in conjunction with forceapply to rerun files which may already have been executed.</p> <p>If you try to apply a patch that has already been applied and do not specify the forceapply parameter, adop will display an error like this:</p> <pre data-bbox="959 982 1458 1199">[WARNING] Skipping the application of patch 14125999_AR since it has been already applied [WARNING] Hint: Patches can be applied again by specifying options=forceapply when invoking adop</pre>
generateportion	<p>Purpose: Tells adop whether to run commands normally found in the generate portion of the driver.</p> <p>Default: generateportion</p> <p>Use <i>options=nogenerateportion</i> to tell adop not to perform generate actions of the driver.</p>

Option	Description
hotpatch	<p data-bbox="862 306 1341 333">Purpose: Tells adop to run in <i>hotpatch</i> mode.</p> <p data-bbox="862 363 1076 390">Default: nohotpatch</p> <p data-bbox="862 417 1313 478">Comments: Can only be specified with the apply phase. For example:</p> <pre data-bbox="862 489 1279 516">adop phase=apply hotpatch=yes</pre> <p data-bbox="862 543 1341 604">The finalize, cutover, and cleanup phases are run automatically.</p> <p data-bbox="862 632 1341 756">Warning: The hotpatch option can only be used to apply patches that provide explicit support for it. The readme of a patch will state if this is the case.</p> <p data-bbox="862 783 1341 909">For more information, see Hotpatch Mode Restrictions in the Patching and Management Tools chapter of <i>Oracle E-Business Suite Concepts</i>.</p>
integrity	<p data-bbox="862 1003 1354 1094">Purpose: Tells adop whether to verify that the version of each file referenced in a copy action matches the version present in the patch.</p> <p data-bbox="862 1121 1073 1148">Default: nointegrity</p> <p data-bbox="862 1176 1325 1234">Comments: Using <i>options=nointegrity</i> is safe and avoids some adop overhead.</p>
parallel	<p data-bbox="862 1283 1369 1407">Purpose: Tells adop whether to run actions that update the database in parallel (such as SQL) and actions that generate files in parallel (such as genform).</p> <p data-bbox="862 1434 1036 1461">Default: parallel</p> <p data-bbox="862 1488 1357 1577">Comments: Oracle does not recommend changing the default, as Oracle E-Business Suite patches are fully tested using this option.</p>

Option	Description
phtofile	<p>Purpose: Tells adop where to place patch history information after applying the patch.</p> <p>Default: nophtofile</p> <p>Use <i>options=phtofile</i> to tell adop to write patch history information to the patch information files in the file system (\$APPL_TOP/admin/\$TWO_TASK) instead of uploading it to the database.</p> <p>Comments: Using phtofile allows you to defer the uploading of patch history information to the database until after the system downtime.</p>
validate	<p>Purpose: Tells adop whether to connect to all registered Oracle E-Business Suite schemas at the start of the patch.</p> <p>Default: novalidate</p> <p>Use <i>options=validate</i> to validate password information for all Oracle E-Business Suite schemas.</p> <p>Comments: Useful for finding problems with incorrectly registered Oracle E-Business Suite schemas or schemas with invalid passwords.</p>

The adop Interface

adop is run from the command line. It prompts for any information required.

Running adop

The following is a summary of the steps you use to run adop. For a complete procedural description of all the steps, see *Creating Customized Instructions for Patching Using PAA*, page 3-4.

Step 1: Set the environment

You must set the environment to apply the configuration parameters that define your system. This task is common to many AD utilities.

Step 2: Unzip the patch

Create a patch top directory, if it does not already exist. Download the patch into the patch top directory and unzip it.

Step 3: Review the information in the readme file

In the directory where you unzipped the patch, you will find a README.txt file and a README.html file. Review either readme file for information about the patch and for instructions on using Oracle Patch Application Assistant (PAA) to generate customized instructions for your system.

Step 4: Run Oracle Patch Application Assistant

Run PAA (admsi.pl) to generate customized instructions for your system. Follow the steps in the customized instructions to start the patching process.

Step 5: Run adop

The customized instructions generated by PAA describe how to run adop using the adop command.

Note: You can add arguments on the command line to refine the way adop runs. See adop Modes, page 2-31 and Command Line Arguments, page 2-33.

Monitoring Status

You can obtain a brief report for the current patching session by running the command:

```
$adop -status
```

This will display information that includes phases completed and the time taken. In the example below, the current patching session ID is 5:

```
Current Patching Session ID: 5
Node Name      Node Type      Phase          Status          Started
-----
                Finished          Elapsed
-----
patchtest1     master
11:31:38 -07:00 07-MAY-13 12:27:51 -07:00 0:56:13
                APPLY          COMPLETED    07-MAY-13
04:19:17 -07:00 07-MAY-13 04:43:12 -07:00 0:23:55
                CUTOVER        COMPLETED    07-MAY-13
05:54:03 -07:00 07-MAY-13 05:57:57 -07:00 0:03:54
                CLEANUP        COMPLETED    07-MAY-13
09:14:33 -07:00 07-MAY-13 09:22:46 -07:00 0:08:13
```

The output is also recorded in a log file, called

```
/u01/R122_EBS/fs_ne/EBSapps/log/status_<YYYYMMDD>_<HHMMSS>/adzds  
howstatus.out. For example,  
/u01/R122_EBS/fs_ne/EBSapps/log/status_20130507_111647/adzdshows  
tatus.out.
```

Two additional options with this command are as follows.

- If you want information about a particular session, specify the relevant session ID:
`$adop -status <session ID>`
- If you want additional details of operations performed:
`$adop -status -detail`

This option will give a summary of last ten adop session IDs, and full details of the file system and database changes introduced by a patch. It also shows the log file location of the current patching cycle.

Stopping adop

You can abandon a patching cycle by using the command:

```
$ adop phase=abort
```

Important: Only the prepare or apply phases of the online patching cycle can be aborted. That is, you cannot run the abort phase after the cutover phase has been run.

Restarting adop

If you have shut down the workers, or if adop quits while performing processing actions, it saves all the actions completed up to that point in restart files. Investigate and resolve the problem that caused the failure, then restart adop. After you restart adop, it will ask if you want to continue with the previous session (at the point where the processing stopped), or start a new session.

Note: A difference from adpatch is that adop restart behavior is controlled by the abort=Y/N and restart=Y/N options in the `input_file` that can be passed to the adop command in the apply phase.

See: Restarting a Utility, page 7-52 in this book.

AD Merge Patch

Important: The functionality of AD Merge Patch is now included in the

adop tool. If you want adop to merge patches, you must explicitly specify `merge=yes` when invoking adop. AD Merge Patch is still supported, however, and its usage is described in this section.

When patches are applied individually, tasks such as responding to prompts and linking executables must be performed separately for every patch. This can be time-consuming and prone to error.

An alternative is to use *AD Merge Patch*. This utility merges multiple patches into a single patch, allowing you to reduce patch application time by eliminating the tasks you would otherwise have to have performed for each individual patch.


When merging compatible patches, AD Merge Patch bases its actions on metadata. It removes duplicate driver lines from the database portions of the driver. When merging two or more patches that have manual steps, the steps and readme files of both patches are also merged.

Source and Destination Directories

You extract the patches to be merged from the *source* directory. The *destination* directory is where the merged patch is created. AD Merge Patch reads the patch driver files for each patch in the source directory and merges them to create patch driver files in the destination directory. If a file exists in more than one source patch, only the highest revision of the file is copied to the destination directory.

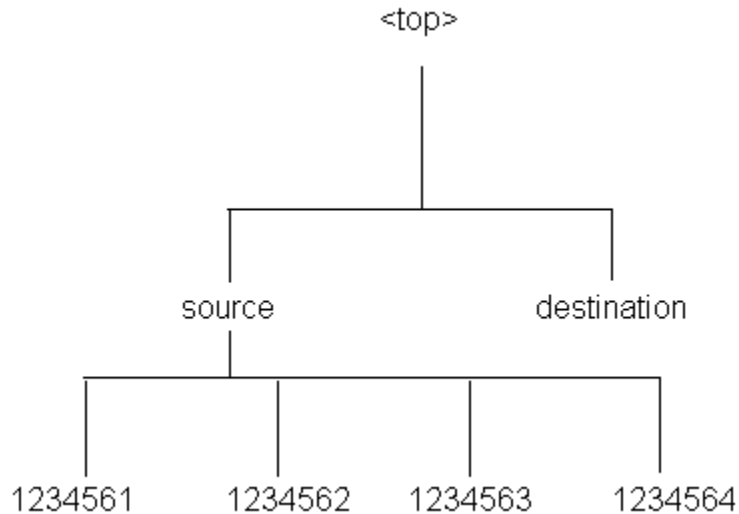
The source and destination directories should be created under the same parent directory. For example, if the parent directory is named `<top>`, both the source and destination directories should be subdirectories of `<top>`. The source and the destination directories cannot be child or parent directories of each other.

Directory Structure for Source and Destination Directories - Basic Example



The source directory must have all patches to be merged as immediate child directories. The patch directories cannot be in a lower directory. For example, a directory structure for merging four patches would look like this:

Directory Structure for Source and Destination Directories - Merging Four Patches



Naming the Merged Patch

You should indicate the name of the merged patch on the command line, using the `-merge_name` option to provide a meaningful name. If you do not use this option, the patch will be given the default name of `merged`.

Merging Zipped Patches

The manifest file is a text file in which you document the location and names of the patch zip files. The contents of a manifest file resemble the following:

```
/d01/prodappl/patches/p3903945_12_GENERIC.zip
/d01/prodappl/patches/p3892799_12_GENERIC.zip
/d01/prodappl/patches/p3874740_12_LINUX.zip
```

You can use the `-manifest` option to create a manifest file. AD Merge Patch references this file, and unzips the patches listed. It copies the unzipped files into the source directory and includes them, along with any other files in the source directory, in the merged patch.

The AD Merge Patch Interface

You run AD Merge Patch and supply the information it needs from the command line. There are no menus or input screens.

Running AD Merge Patch

AD Merge Patch is located in the `AD_TOP/bin` directory. However, you run it from the parent directory (`<top>`) of the source directory. The following is a summary of the steps

you use to run AD Merge Patch.

Step 1: Set the environment

You must set the environment to indicate the location of the configuration parameters that define your system. This task is common to many AD utilities.

Step 2: Run AD Merge Patch

From the parent directory (<top>), run AD Merge patch using the admrgpch command.

Patching Procedures

Preparing for Patching

For patches that have manual steps, the patch readme file instructs you to use Oracle Patch Application Assistant (PAA) to create customized instructions for your system. PAA consolidates and displays only the relevant manual steps for all the patches you want to apply, including steps that you have completed. It also automatically merges the contents of individual patch readme files for a merged patch.

Create Checklist of Product Functionality Disabled in Online Patching Cycle:

During an online patching cycle, the following product restrictions will apply. Before you commence patching, you should therefore ensure there will be no requirement for any these actions or features until the cycle is complete.

- **Payroll**
 - Users will not be able to define Fast Formulas or use the Fast Formula Assistant.
 - Users will not be able to perform dynamic trigger maintenance.
 - Users will not be able to create, update, or delete US Cities.
 - Data Pump meta-mapper generator will be disabled.
 - The Japanese Balance Dimensions concurrent program will be deferred to after the cutover phase is complete.
 - Pension Calculation Setup cannot be used.
 - US localization earnings and deduction setup cannot be used.

- Tax Withholding Rules Setup cannot be used.
- Wage Attachment Earnings Rules Setup cannot be used.
- Garnishment Rules Setup cannot be used.
- Quick Paint Reports cannot be used.
- Quantum Program Update Installer execution is unavailable.

- Order Management:
 - Creation of a new Defaulting Condition in the Attribute Defaulting Rules form is disabled, unless the same seeded condition already exists for a given attribute.

- Warehouse Management:
 - WMS Rule creation is restricted.

- Inventory:
 - Concurrent program "Generate Stock Locator Flexfield Definition for Mobile Transactions" will be disabled.

- Public Sector Financials International:
 - Users will not be able to run the following concurrent programs:
 - Subledger Security: Apply Security
 - Subledger Security: Import/Export Data Fix

- Subledger Accounting:
 - Users will not be able to Validate the Application Accounting definitions.

- Accounts Receivable:
 - Users will not be able to create new Transaction Sources.

- Incentive Compensation:
 - Transaction collection process for new mappings will not be available and any changed mapping will continue to use previous mapping rules.
 - Users will not be able to run the "Synchronize Classification Rulesets" program.

- Users will not be able to use the "Formula Generation" feature.
- Users will not be able to specify new formulas or changes to compensation rules.
- Oracle Demand Planning:
 - Demand plans will not be available for users.

Set Up Secure Shell on Application Tier Nodes:

If you are using the typical configuration of a multi-node application tier, and want to take advantage of *adop remote invocation*, you must set up *Secure Shell* (`ssh`) on all your application tier nodes. This will allow *adop* to be run automatically on all nodes.

Note: Rapid Install and Rapid Clone set up the `ssh` key infrastructure.

Principles

The `ssh-keygen` command is used to generate a *private/public key pair*. The *private key* is for the node from where all the remote nodes will subsequently be accessible by an `ssh` login that requires no password. The *public key* must be copied to each remote node's `<User Home Dir>/ .ssh` directory.

In essence, the sequence is as follows:

1. The following command initiates creation of the key pair:

```
$ ssh-keygen -t rsa
```

Note: The `<Enter>` key should be pressed instead of a passphrase being entered.

2. The private key is saved in `<User Home Dir>/ .ssh/id_rsa`.

Important: As this read-only file is used to decrypt all correspondence encrypted with the public key, its contents must not be shared with anyone.

3. The public key is saved in `<User Home Dir>/ .ssh/id_rsa.pub`.
4. The contents of the public key are then copied to the `<User Home Dir>/ .ssh /authorized_keys` file on the systems you subsequently wish to `ssh` to without being prompted for a password.

The following example demonstrates the steps:

1.

```
$ ssh-keygen -t rsa
Generating public/private rsa key pair.
Enter file in which to save the key (/u01/user2/.ssh/id_rsa):<Enter>

Enter passphrase:<Enter>
Enter same passphrase again:<Enter>
Your identification has been saved in /u01/user2/.ssh/id_rsa.
Your public key has been saved in /u01/user2/.ssh/id_rsa.pub.
The key fingerprint is:
16:d0:e2:dd:37:2f:8e:d5:59:3e:12:9d:2f:12:1e:5a
```
2.

```
$ scp -pr /u01/user2/.ssh/id_rsa.pub \
user2@system1:/u01/user2/.ssh/authorized_keys
user2@system1's password:<password>
id_rsa.pub 100% 398 0.4KB/s 00:00
```
3.

```
$ ssh user2@system1
```

Note: If you receive this message, it can safely be ignored:
Warning: untrusted X11 forwarding setup failed:
xauth key data not generated Warning: No xauth
data; using fake authentication data for X11
forwarding.

Once this has been done for the relevant operating system account on all nodes - that is, ssh can log in from the primary node to each secondary node without entering a password - so you are ready to run adop on multiple application tier nodes. It must be run on at least the master (admin) node: from there, it will attempt to contact all the other application tier nodes that are part of the same Oracle E-Business Suite instance, and will run the required steps remotely on those nodes. If it cannot contact any of those nodes, or if the administrator prohibits it (by specifying the allnodes=no option), the administrator must manually connect to those nodes and run adop with the options allnodes=no and action=nodb.

Important: If you do not wish to set up ssh, you will need to run the relevant adop operation on each of your application tier nodes individually.

Important: If you change the password for the relevant operating system account on one or more nodes, you must regenerate the ssh credentials either using the \$AD_TOP>/patch/115/bin/txkRunSSHSetup.pl script, or your own native solution if you prefer. The txkRunSSHSetup.pl script has a -help option that shows relevant usage options such as enablessh.

Creating Customized Instructions for Patching Using PAA:

Requirement: How do I know which manual steps associated with a patch apply to my

system?

Sorting through the manual steps in a patch readme file to determine which ones apply to your system can be time-consuming. The Patch Application Assistant allows you to create a customized set of steps to that apply to your unique instance. Using the information on this list reduces the possibility of performing steps that are not necessary or that have been completed previously during the application of another patch.

When you download and unzip a patch, it delivers a static README.html file that advises you if the patch requires manual steps. If manual steps are required, you can generate a list of the steps by running a Perl script (admsi.pl) to initiate PAA. Once you have generated the list, use the PAA interface to see a full list of steps, or only those steps that apply to your system.

After successfully performing each manual step, you can record that it was *completed*. When applying patches in the future, this information is displayed in the PAA interface so that you can see which manual steps you have already performed. Unless specified otherwise, you can previously completed steps.

To run PAA

1. Download the patch that you want to apply and set (source) your environment. On UNIX systems, you must also set the environment variable DISPLAY to an active and authorized display.

For instructions on setting your environment, see: Running AD Utilities, page 7-52 in this book.

2. Run the admsi.pl script to generate customized installation instructions.

```
$ admsi.pl
```

The Oracle Patch Application Assistant welcome page appears:

You can select:

- View instance-specific instructions for a new patch.
 - View generic instructions as shipped by Oracle for a new patch - to view all the generic manual steps for a particular patch, including the completed steps.
 - Look at all incomplete tasks from previous patches - to view all the manual steps that have not been completed from previous patches.
3. Select View instance-specific instructions for a new patch. Enter the APPS password, and select the location where the patch is staged. Click Next.

The Summary of Installation Instructions page appears:

This page summarizes all the manual steps for the patch, grouped into the following categories: Preparation Tasks, Pre-Install Tasks, Apply the Patch, Post-Install Tasks, Finishing Tasks, and Additional Information. This page displays

only those categories in which there are manual steps.

4. Click the plus-sign icon in each category for more detailed information. For example, if you click the plus-sign icon next to Best Practices, the Preparation Tasks screen appears with the tasks suggested for preparing your system for patching.
5. After you have completed all the manual steps in a category, check the Completed box to record the completion status in the database, then click Next. If a patch that you apply in the future contains any of the same manual steps, it will be marked as *completed* to inform you that you do not have to perform that task again.

After you have completed all manual steps in all categories, the system returns you to the Summary of Installation Instructions page.

Note the column of Completed boxes that corresponds to each task in a category. Check marks appear in the boxes for which you have completed manual steps.

6. Click Save to record tasks completed in the database. Click Cancel to exit PAA.

The Online Patching Cycle

This section describes the online patching cycle from beginning to end, illustrating the actions taken in the different phases and putting into context the more detailed description of online patching in the following sections.

Important: This section is designed to be read in conjunction with the important background material provided in the Patching and Management Tools chapter of *Oracle E-Business Suite Concepts*.

Applying Oracle E-Business Suite patches without a significant system downtime is referred to as *online patching*, and a new utility, *adop*, is used to apply patches.

Online patching is supported by the capability of storing multiple application editions in the database, and the provision of a dual application tier file system. At any given point in time, one of these file systems is designated as 'run' (part of the running system) and the other is the 'patch' (either being patched or awaiting the start of the next patching cycle). Whichever is the current run file system appears to the user in exactly the same way as the single application tier file system did in Oracle E-Business Suite releases prior to 12.2.

Important: The existence of the dual file system has implications for where general (non-patching) maintenance activities are carried out. For important information on choosing the appropriate file system to run AD tools from, refer to: Choosing the Correct File System For Maintenance Tasks, page 7-1 in Chapter 7 of this book.

A new environment variable, `$FILE_EDITION`, shows the current designation of a given dual file system member. Three other new environment variables designate the root directories of the run (`$RUN_BASE`), patch (`$PATCH_BASE`), and non-editioned (`$APPL_TOP_NE`) file systems.

For example:

- `$FILE_EDITION = patch`
- `$RUN_BASE = /u01/R122_EBS/fs1`
- `$PATCH_BASE = /u01/R122_EBS/fs2`
- `$APPL_TOP_NE = /u01/R122_EBS/fs_ne/EBSapps/appl`

When a patch is being applied, the Oracle E-Business Suite system is running in normal production mode (full functionality) in the run edition of the file system and database. Full application functionality is retained as patch execution proceeds, until the cutover phase is reached (as described later in this section).

Important: Do not attempt to clone a system while an online patching cycle is in progress.

The online patching cycle consists of a number of high level phases:

1. prepare
2. apply
3. finalize (called automatically)
4. cutover
5. cleanup (called automatically)

A high level overview of an online patching cycle would, programmatically, look like this:

```

# Source patch edition environment
source /u01/R122_EBS/fs2/EBSapps/appl/APPS<CONTEXT_NAME>.env

# Prepare patch edition
adop phase=prepare

# Apply patches
adop phase=apply patches=<patch number>

# Apply customizations to patch edition
sqlplus apps/apps @my_custom_script_01
sqlplus apps/apps @my_custom_script_02
...

# Run finalize
adop phase=finalize

# Run cutover
adop phase=cutover

# Perform user acceptance testing via application UI

# Run cleanup
adop phase=cleanup

```

Additional points:

- Customizations are applied to the patch edition during the apply phase, normally after any Oracle E-Business Suite patches have been applied.
- Maintenance Mode is not needed for online patching, and so is not available in Oracle E-Business Suite Release 12.2.

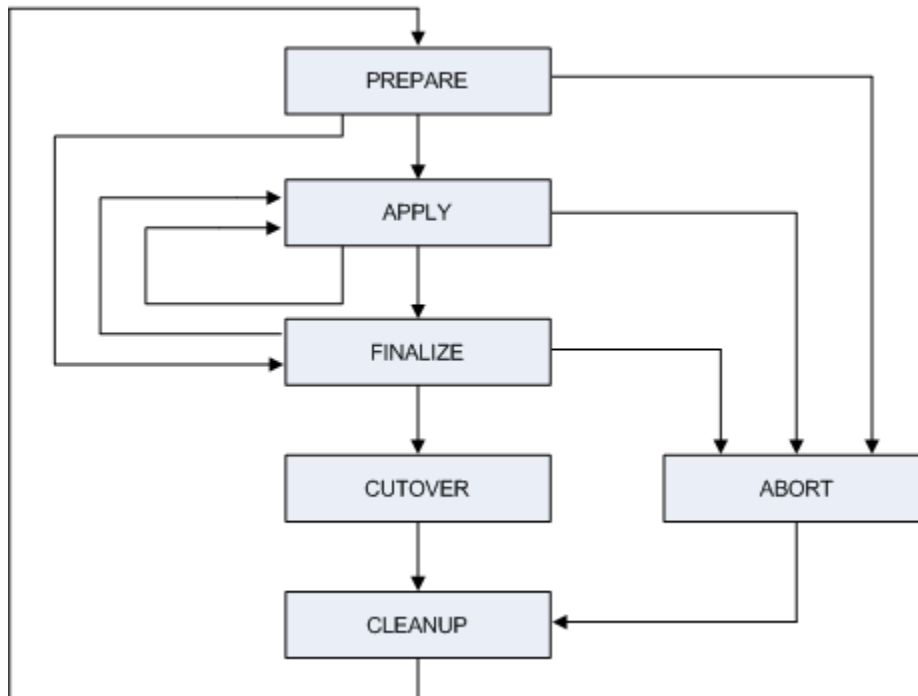
Special Phases

Two additional phases are provided for specialized use. Neither can be run in conjunction with any other phase. Further details of these phases are described in later sections.

- The *abort* phase is used to abandon a patching cycle before it is complete, and roll back any changes that have been made.
- The *fs_clone* phase is a stand-alone command used for file system cloning.

Note: Standard cloning (using `adcfgclone.pl`) cannot be used to synchronize the run and patch file systems. It can only be used for traditional cloning scenarios, for example cloning a development system to a test system. Conversely, *fs_clone* is only for use with `adop`.

The Online Patching Cycle



Running the adop tool is similar to other AD command-line utilities such as adctrl. In the simplest case, you can just enter the following command on the run file system, and be prompted for further input:

```
$ adop
```

Note: adop will automatically set its environment to the run file system. It cannot be executed from the patch file system.

The adop tool can execute patching cycles interactively or non-interactively. In either case, it can run all the applicable phases from a single command, or one phase at a time. In addition, it can be used to generate patching reports.

The adop tool is typically run as follows:

```
adop phase=<phase_name>
```

The phase parameter accepts the following phase names. These names can be specified individually, or (except where otherwise noted) with other phase names in a comma-separated list:

- `prepare` - Prepares the environment for patching.
- `apply` - Applies the patch the environment.
- `finalize` - Performs the final steps in patch application. Can be run at any point

after prepare and before cutover. Facilitates verification of customizations or other objects that may be affected by a patch.

- `cutover` - Performs the actions needed to make the former patch edition the new run edition, and the former run edition the new patch edition. Cutover also prepares the environment for users, and is the only phase that involves a (brief) period of downtime.
- `cleanup` - Drops now-redundant columns and other objects.
- `abort` - Abandons the patch application process and restores the previous state of the environment. Cannot be specified with any other phase.
- `fs_clone` - Used for file system cloning. This special phase must be run from the run file system, and cannot be specified with any other phase.

A special parameter, `input_file=<file>`, must be used in conjunction with `phase=apply` when non-interactive mode is required. The parameters that can be specified in this input file include some of the `<name>=<value>` pairs that could be supplied on the `adpatch` command line in pre-12.2 releases.

Mandatory and optional `input_file` parameters include:

- **defaultsfile**: Full path to defaults file. Mandatory.
- **workers**: Parallel worker count. Mandatory.
- **patches**: Patches to be applied. Mandatory. There are two scenarios:
 - *Containing numbers only*: Most patches are specified using only numbers. If you wanted to apply patch 111, you would specify it as `patches=111`. The associated low-level directory and driver file will be extrapolated from this number.
 - *Containing a colon*: Patches can also be specified using a colon. For example, you could (as noted earlier) supply the Korean language version of patch 111 as `patches=111_KO:u111.drv`. The low level directory is `111_KO`, and the unified driver file to apply is `u111.drv`. This specification should be used for `hrglobal.drv`, online help patches, language-specific patches, and any other patches whose names are not all numbers.

You can use a comma-separated list to specify that multiple patches are to be applied in a single patching operation. In addition, you can mix and match number-only patches and patches containing a colon: for example, `patches=111,222_KO:u222.drv`.

- **merge**: Optional. Takes the values 'yes' or 'no' without quotes. Merges the patches specified in the parameter before they are applied. All such patches must be individual ones: i.e. none should already be merged.

- **abandon:** Optional. Takes the values 'yes' or 'no' (without the quotes).
- **restart:** Optional. Takes the values 'yes' or 'no' (without the quotes).

Note: The 'abandon' and 'restart' parameters can be used if a patching operation had an error. They should always be specified together, with one as 'no' and the other as 'yes':

- `abandon=no restart=yes` will retry a patching operation
- `abandon=yes restart=no` will *not* retry a patching operation

There are two phase-independent adop options, as follows.

- **logfile**

To create a user-specified log file of patching activity, enter a command such as:

```
$ adop phase=prepare logfile=<path to log file name>
```

- **analytics**

To run a script that will generate a report of activities performed in a given phase, enter a command such as:

```
$ adop phase=apply analytics=yes
```

The report can be useful in helping with diagnosis of certain issues.

Note: The 'analytics' option should only be used when you specifically want to check on the activities performed in a particular phase.

For a full description of adop options and usage, see adop Parameters, page 2-3 in Chapter 2.

Important: Many Oracle E-Business Suite systems employ multiple application tier nodes. On such systems, you run the adop prepare, apply, and cutover phases on the different nodes depending on whether you have set up adop remote invocation via ssh (see Set Up Secure Shell on Application Tier Nodes, page 3-3). The finalize, actualize_all, and cleanup phases all perform database actions, so only need to be run on one node (the master node).

Specific adop commands for multi-node systems are described under the individual phases.

Online Patching Cycle Steps - Prepare Phase:

This section describes the principles of adop operation in the prepare phase, followed by the steps you take to run this phase.

Note: The exact actions taken during the prepare phase are context-dependent: for example, the first time it is ever run on a system; when it is run after an apply phase has been aborted; and when it has been run after cutover.

Principles

In the prepare phase, adop:

1. Checks that the environment is set to the run APPL_TOP.
2. Checks whether to perform a cleanup, which will be needed if the user failed to invoke cleanup after the cutover phase of a previous online patching cycle.
3. Checks to see if the database is prepared for online patching:
 - Checks that the FILE_EDITION environment variable value is set to 'run', and if not displays the message: `Error: Environment variable FILE_EDITION is expected to have value as 'run'. Hint: Ensure that you are executing the tool from the run APPLTOP.`
 - Checks to see if enough space is available in the database.
 - Checks if the database user is edition-enabled (at least one user, the PL/SQL API, will return 'true'). If no edition-enabled users exist, adop displays the message: `Error: Users are not edition-enabled. Hint: Please prepare the ENV for Online Patching.`
 - Checks to see if the patch service has been created. If the check fails, adop displays the message: `Error: Patch service check failed Hint: Try to create the patch service.`
 - Checks to see if logon trigger exists and is enabled. If the logon trigger is missing or the patch service has not been created, adop will automatically try to fix the issue so that it can proceed. If it cannot do so, it will exit with an error message.
4. Checks the file system, using the TXK script `$AD_TOP/patch/115/bin/txkADOPPreparePhaseSanityCheck.pl`. This script checks for the file system space, database connections, and so on.
5. Produces a report showing information about the most important tablespaces is

generated. This report is created in `$APPL_TOP/admin/$TWO_TASK/out`.

6. Checks for the existence of the "Online Patching In Progress" (ADZDPATCH) concurrent program. This program prevents certain predefined concurrent programs from being started, and as such needs to be active while a patching cycle is in progress (that is, while a database patch edition exists). The flow of control is as follows.
 1. If the ADZDPATCH program exists, a message to this effect is displayed to the user. If it does not exist, it is started.
 2. The status of ADZDPATCH is determined. If it is pending, it may be waiting for an incompatible program to finish. At that point, its status will change to running, and it will allow the prepare phase to proceed. A message to this effect is displayed to the user.
 3. The next stage depends on whether the concurrent managers are running:
 1. If the concurrent managers are all down, the prepare phase continues, with ADZDPATCH entering a status of pending (with the highest priority) until the managers are started.
 2. If the concurrent managers are partially up, but there is no manager defined that can run ADZDPATCH, then the prepare phase will exit with an error.
 3. If the concurrent managers are up, and there is one defined that can run ADZDPATCH, processing will loop until ADZDPATCH changes status from pending to running (that is to say, as noted in Step 2, no incompatible programs are found). The prepare phase then continues.

Note: ADZDPATCH is cancelled when the cutover phase is complete.

7. Invokes the TXK script
`$AD_TOP/patch/115/bin/txkADOPPreparePhaseSynchronize.pl` to synchronize the patches which have been applied to the run `appltop`, but not the patch `APPL_TOP`. The script depends on the adop repository for patches that have been applied on the run `APPL_TOP` but not the patch `APPL_TOP`.
8. Checks the database for the existence of a patch edition, and creates one if it does not find one.
9. Calls the `$AD_TOP/patch/115/bin/txkADOPPreparePhaseSanityCheck.pl` script again to confirm that the database connection to the patch edition is working.

If any of these checks fail, adop will exit with an error message.

Steps

You perform the following steps in the prepare phase.

Important: Before you run the prepare phase, you should ensure that the following space requirements are met:

- SYSTEM tablespace has a minimum of 25 GB of free space
- APPS_TS_SEED tablespace has a minimum of 5 GB of free space

You can check the available space by running the \$AD_TOP/sql/ADZDSHOWTS.sql script. For information about increasing the size of a tablespace, refer to the Altering and Maintaining Tablespaces section in the Managing Tablespaces chapter of *Oracle Database Administrator's Guide*.

1. Set the environment by executing (sourcing) the run file system environment file:

```
$ source <run APPL_TOP path>/APPS<CONTEXT_NAME>.env
```

For more information, see *Setting the Environment in Running AD Utilities*, page 7-52

You can confirm that the environment is properly set by examining the relevant environment variables:

```
$ echo $FILE_EDITION
run
$ echo $TWO_TASK
dbSID
```

If you had sourced the *incorrect* environment file (i.e. from the patch file system), the environment variables would show as:

```
$ echo $FILE_EDITION
patch
$ echo $TWO_TASK
dbSID_patch
```

2. Download the patches to be applied, and place them in a central location on the non-edited file system:

```
$APPL_TOP_NE/EBSapps/patch
```

This is the 'patchtop' location, to which you will download all the patches you want to apply.

3. Unzip the patch:

```
$ unzip p99999999.zip
```

4. Prepare the instance for patching by running the following command to create the

new database edition:

```
$ adop phase=prepare
```

In this step, the patch file system APPL_TOP is synchronized with the run file system APPL_TOP. This can be done by either of two methods, the first of which is the default:

- **Method 1** - Identify the patches that were applied to the run APPL_TOP and apply them to the patch APPL_TOP. The following steps are performed *automatically*:
 1. The list of patches that need to be applied to the patch APPL_TOP are identified from the database.
 2. The merged patches are applied using the adop utility.

This lightweight process is the preferred approach. The adop tool identifies the delta patches to be applied, and applies them silently to the current patch APPL_TOP. As this procedure only requires the application of delta patches, it is relatively fast compared to Method 2 (below).

- **Method 2** - Create a new patch file system by cloning the run file system.

This method is useful if the APPL_TOPs have become very unsynchronized (meaning that there would be a large number of delta patches to apply). It is a heavyweight process, taking a backup of the entire current patch APPL_TOP and then cloning the run APPL_TOP to create a new patch APPL_TOP. A total of at least 75 GB of free disk space is required. Also, you will need at least 25 GB of free space in your temporary directory (typically /tmp). As well as being resource-intensive, this method is very time-consuming and requires several manual actions by the user. It should therefore be used only when necessary.

The steps are as follows:

1. The existing APPL_TOP, COMMON_TOP, 10.1.2 Oracle Home and FMW_Homes are renamed.
2. The APPL_TOP, COMMON_TOP, and 10.1.2 Oracle Home are copied from the run edition.
3. The fs_clone phase is run using the following command:

```
adop phase=fs_clone
```

This command must be invoked from the *run* file system, before the next prepare phase is run.

Important: Windows users should refer to their platform-specific release notes for restrictions that currently

apply when running `fs_clone`.

Any machine that is abandoned must be re-imaged (via a clone from another application tier machine) before it can rejoin the Oracle E-Business Suite instance. For more information, see the Cloning Oracle E-Business Suite with Rapid Clone chapter in *Oracle E-Business Setup Guide*

Using the `skipsyncerror` parameter

This parameter enables you to specify that you expect any synchronization errors in the prepare phase to be fixed automatically in the synchronization that takes place with subsequent patches.

If the value of the parameter is passed as 'yes', the first patch to be synchronized will be done with the 'autoskip' flag set.

Important: It is your responsibility to check the log files and correct any errors in the subsequent apply phase, or to confirm that synchronization with subsequent patches resolved the issue.

An example of using this parameter would be as follows.

1. You run `adop phase=prepare`.
2. The phase fails with an error when trying to synchronize the run and patch file systems. That is, of the patches gives an error on an attempt to apply it with the 'noautoskip' flag set.
3. You examine the log files and conclude that the synchronization errors will be fixed automatically in the synchronization that takes place with subsequent patches.
4. You run the command `adop phase=prepare skipsyncerror=yes` to restart the prepare phase. This time, application of the patch that failed in the previous prepare will be retried with the 'autoskip' flag set.

Alternatively, if you are not confident that the error will be fixed (for example, you cannot identify the cause from examination of the log files), you should:

1. Run the command `adop phase=abort`
2. Run the command `adop phase=fs_clone`
3. Rerun the command `adop phase=prepare`

Synchronizing Customizations

You need to ensure that any customizations you may have are preserved when the

run and patch file systems are being synchronized. Examples include:

- Compiling user-defined JSPs
- Copying some third-party libraries
- Executing user-defined SQL files
- Compiling user-defined forms

To add such actions, you need to edit the file

<s_ne_base>/EBSapps/appl/ad/custom/adop_sync.drv. You will add your customizations to this section of the file:

```
#Begin Customization
...
#End Customization
```

All the actions defined in this file will be performed by adop automatically during the prepare phase. Be aware that there are two categories of custom command in adop_sync.drv: those that are run one time only, and those that are run at each file system synchronization (during the adop prepare phase).

Important: The adop_sync.drv file is not currently reset to its template file at any point. Consequently, after cutover (and before the next prepare phase), you should review the contents of adop_sync.drv and ensure the requirements for your custom commands continue to be met.

This is only an outline of the steps you need to take to preserve customizations. For full details, refer to *Oracle E-Business Suite Developer's Guide*

Prepare Phase in Multi-Node Environments

In a multi-node environment, one application tier node will be designated as the *primary node*. This is the node where the Admin Server is located, and will usually also be the node that runs Oracle HTTP Server. All other application tier nodes are designated as *secondary nodes*.

The recommended strategy in a multi-node environment is to enable the ssh utility, and thereby enable adop to be run remotely on other nodes via *adop remote invocation*. The ssh setup steps are described in Set Up Secure Shell on Application Tier Nodes, page 3-3. When ssh is used, the primary (master) node is the node from which adop is always run, communicating with all the secondary nodes via ssh.

An alternative strategy is to perform the steps on each node individually, and not take advantage of adop remote invocation.

Different parameters are used with adop, according to the precise configuration. The following scenarios for the prepare phase will illustrate this:

- **Shared Application Tier Filesystem:** On the primary node, run the command:

```
adop phase=prepare input_file=<input_file name>
```
- **Shared Application Tier Filesystem Without ssh:**
 1. On the primary node, run the command:

```
adop phase=prepare allnodes=no action=db  
input_file=<input_file name>
```
 2. On the other (secondary) nodes, run the command:

```
adop phase=prepare allnodes=no action=nodb  
input_file=<input_file name>
```
- **Non-Shared Application Tier File System:** On the primary node, run the command:

```
adop phase=prepare input_file=<input_file name>
```
- **Non-Shared Shared Application Tier Filesystem Without ssh:**
 1. On the primary node, run the command:

```
adop phase=prepare allnodes=no action=db  
input_file=<input_file name>
```
 2. On the other (secondary) nodes, run the command:

```
adop phase=prepare allnodes=no action=nodb  
input_file=<input_file name>
```

Note: The equivalent commands are used with the apply phase.

If a node unexpectedly becomes inaccessible via ssh, it will be *abandoned* by adop, and the appropriate further actions taken. Consider a scenario where the adop phase=prepare command is run in a system with ten application tier nodes. The command is successful on nine nodes, but fails on the tenth. In such a case, adop will identify the services enabled on nodes 1-9. If they are sufficient for Oracle E-Business Suite to continue to run normally, adop will mark node 10 as abandoned and then proceed with its patching actions. If they are *not* sufficient, adop will proceed no further.

Online Patching Cycle Steps - Apply Phase:

This section describes the principles of adop operation in the apply phase, followed by the steps you take to run this phase.

Principles

In the apply phase, adop actually applies the specified patches. Either interactive or non-interactive mode can be used.

Steps

In this phase, you will apply the patches that you designated for inclusion in this patching cycle. You can apply as many patches as you want per patching cycle. The adop utility supports both interactive and non-interactive mode.

- When using interactive mode, you can allow adop to merge the patches and then apply the merged patch.
- When using non-interactive mode, you can use the merge option in the input_file for the patches to be automatically merged and applied.

The following two steps will illustrate the options. Follow either Step 1 for interactive mode or Step 2 for non-interactive mode.

1. To apply your patch in *interactive mode*, run the command:

```
$ adop phase=apply
```

2. To apply your patch in *non-interactive mode*, run the command:

```
$ adop phase=apply input_file=<inputfile.txt>
```

This uses the input_file that was mentioned earlier in this section.

An example input_file might look like this:

```
defaultsfile=<defaults file on patch APPL TOP>
workers=<number of workers>
patchtop=<central location for patches>
patches=<patch number 1>:<driver file 1>.drv, <patch number
2>:<driver file 2>.drv ...
```

Note: The 'patchtop' is the directory where you unzipped the patches. The default (and recommended) location is
\$APPL_TOP_NE/EBSapps/patch.

An example defaults file is located in:

```
$APPL_TOP/admin/<SID>/adalldefaults.txt
```

Reports under the \$APPL_TOP/admin/<SID>/out directory can help you identify and diagnose problems that may occur in the online patching cycle. These reports list the proposed changes to database objects, both new and modified.

The key files to examine are:

- \$APPL_TOP/admin/<SID>/out/adzdcmped.out
- \$APPL_TOP/admin/<SID>/log/u<patch_number>.log

Note: For merged patches, the log file name will be derived from the timestamp when merging was performed.

Using the analytics parameter in apply

If you want to use the `analytics` parameter (see `adop Parameters`, page 2-3) with the `apply` phase, enter the command:

```
adop phase=apply analytics=yes
```

Specifying this option will cause `adop` to run the following scripts and generate the associated output files (reports):

- `ADZDCMPED.sql` - This script is used to display the differences between the run and patch editions, including new and changed objects. The output file location is:
`/u01/R122_EBS/fs_ne/EBSapps/log/adop/<adop_sessionID>/<apply_directory>/<context_name>/adzdcmped.out.`
- `ADZDSHOWED.sql` - This script is used to display the editions in the system. The output file location is:
`/u01/R122_EBS/fs_ne/EBSapps/log/adop/<adop_sessionID>/<apply_directory>/<context_name>adzdshowed.out.`
- `ADZDSHOWOBSJS.sql` - This script is used to display the summary of editioned objects per edition. The output file location is:
`/u01/R122_EBS/fs_ne/EBSapps/log/adop/<adop_sessionID>/<apply_directory>/<context_name>adzdshowobjs.out`
- `ADZDSHOWSM.sql` - This script is used to display the status report for the seed data manager. The output file location is:
`/u01/R122_EBS/fs_ne/EBSapps/log/adop/<adop_sessionID>/<apply_directory>/<context_name>adzdshowsm.out`

Note: The `analytics` parameter should only be used when required, because of the extra processing needed.

Online Patching Cycle Steps - Cutover Phase:

This section describes the principles of `adop` operation in the cutover phase, followed by the manual steps you can optionally execute to run this phase.

Important: You should ensure that no users remain on the system during cutover, as there will be a short downtime period while the application tier services are automatically shut down and restarted. In addition, any third-party processes connected to the old run edition of the database should be shut down, otherwise they will be terminated automatically.

Principles

The key actions performed in the cutover phase are:

1. *Finalize*: Invoke Finalize API in the database.

Note: If required, you can run finalize as a separate phase instead of allowing it to be run automatically as part of cutover. Refer to adop Parameters, page 2-3 for details.

2. *Application Tier Services Switch, Stage 1*: Shut down services on current run APPL_TOP, and restart services on current patch APPL_TOP.
3. *Database Cutover*: Promote patch database edition to become the new run database edition, using `adzdpmgr.pl` script.
4. *File System Cutover*: Promote patch file system to become the new run file system, switching the `$FILE_EDITION` values in the patch and run environments. The current patch APPL_TOP becomes the new run APPL_TOP, and the current run APPL_TOP becomes the new patch APPL_TOP.
5. *Retire Old Editions*: Disable access to old database editions.
6. *Terminate Old Database Sessions*: Terminate any database connections to the old run edition of the database.
7. *Application Tier Services Switch, Stage 2*: Shut down services on old run APPL_TOP, and start up services on new run APPL_TOP.

Note: The adop utility invokes the TXK script `txkADOPCutOverPhaseCtrlScript.pl` to perform tasks 1, 2, 3, 5, and 6. Task 4 is performed by AutoConfig.

1. Ensure you are ready to commit to application of the selected patches.
2. In most cases (but see below for an important exception), you execute the cutover phase by running the command:

```
$ adop phase=cutover
```

This will promote the patch edition to be the new run edition, as well as switching the patch and run labels on the file systems (and thereby, as noted above, changing the patch file system to be the new run file system and the run file system to be the new patch file system).

Important: Cutover will take longer if it has to wait for long-running concurrent processes to complete. In such a case, you can expect to see an informational message of the form
[STATEMENT] [END 2013/05/28 23:47:16] Waiting for

Using the analytics parameter in cutover

If you want to use the `analytics` parameter (see `adop Parameters`, page 2-3) with the cutover phase, enter the command:

```
adop phase=cutover analytics=yes
```

Specifying this option will cause `adop` to run the following script and generate the associated output file (report):

- `ADZDCOBSJS.sql` - This script is used to display the summary of covered objects per edition. The output file location is:
`/u01/R122_EBS/fs_ne/EBSapps/log/adop/<adop_sessionID>/<cutover_folder>/<context_name>/adzdcobjs.out.`

Note: The `analytics` parameter should only be used when required, because of the extra processing needed.

Patching AutoConfig or TXK Code

An exception to normal cutover is if you are patching the database tier with AutoConfig or TXK code using `adop`. In such cases, you must perform certain manual steps immediately after completion of cutover, and without the application services having been started. To be able to do this, you must execute cutover by running the command:

```
$ adop phase=cutover mtrestart=no
```

You must then also run the steps in the following section, `Patching the Database Tier`:

Patching the Database Tier

These steps are performed post-cutover.

1. On the application tier, as the `applmgr` user:
 1. Change directory to the run file system `$APPL_TOP` and source your environment file.

2. Run the following command:

```
$ perl <AD_TOP>/bin/admkappsutil.pl
```

This will create the `appsutil.zip` file in `<INST_TOP>/admin/out`.

2. On the database tier, as the `oracle` user:

Copy or ftp the `appsutil.zip` file to the `RDBMS_ORACLE_HOME`, then run the following commands:


```
$ cd <RDBMS_ORACLE_HOME>
$ unzip -o appsutil.zip
```

3. Run AutoConfig on the database tier.
4. Run AutoConfig on the run file system of each application tier node.
5. Start the application tier services.

Manual Cutover Steps

Normally, you can allow adop to undertake all cutover tasks itself. Sometimes, however, you may need to undertake cutover manually. For example, if you have some nodes behind a firewall or in a DMZ, ssh may not be able to communicate with them. In such cases, you can run cutover manually.

Important: You must have completed all the prerequisite (prepare, apply, and finalize) phases on *all* the nodes before you issue the requisite cutover command on the admin and other nodes.

In a multi-node environment that has ssh enabled (as described in Set Up Secure Shell on Application Tier Nodes, page 3-3), cutover is performed on the admin node with the `adop phase=cutover` command noted above.

In a multi-node environment that does not have ssh enabled, additional steps are required:

1. On the admin node, run the command: `adop phase=cutover allnodes=no action=db`
2. On the other (secondary) nodes, run the command: `adop phase=cutover allnodes=no action=nodb`

Important: In a multi-node environment that does not have ssh enabled, you must run the above cutover commands at the same time (in parallel) on all the nodes, including the admin node and all secondary nodes.

JAR Files and Cutover

In an online patching cycle, the requisite JAR files are initially stored in the `$APPL_TOP/admin/<SID>/out` directory, and then uploaded into the database during the cutover phase. Therefore, the out directory must not be deleted at least until cutover is complete.

Online Patching Cycle Steps - Cleanup Phase:

This section describes the principles of adop operation in the cleanup phase, followed by the steps performed in this phase.

Important: Do not omit the cleanup phase.

Principles

This phase must be run from the run APPL_TOP.

The following actions are performed during cleanup:

1. Drop covered objects
2. Drop indexes
3. Drop forward cross-edition triggers
4. Clean up seed data

Steps

1. Set the environment by executing (sourcing) the run file system environment file:

```
$ source <run APPL_TOP path>/APPS<CONTEXT_NAME>.env
```

For more information, see *Setting the Environment in Running AD Utilities*, page 7-52

2. Cleanup is performed with the command:

```
$ adop phase=cleanup
```

The adop parameter `cleanup_mode` provides control of cleanup processing:

- `cleanup_mode=quick` - Performs minimum cleanup, and so requires least processing time. This is the default, so does not need to be specified.
- `cleanup_mode=full` - Performs maximum cleanup, which includes dropping of covered objects and unused columns.

Important: On a multi-node system, the cleanup command only needs to be issued on the primary node.

Using the analytics parameter in cleanup

If you want to use the `analytics` parameter (see adop Parameters, page 2-3) with the cleanup phase, enter the command:

```
adop phase=cleanup analytics=yes
```

Specifying this option will cause adop to run the following script and generate the associated output file (report):

- `ADZDCLEANUPRP.sql` - This script is used to display the display the cleanup status. The output file location is:

```
/u01/R122_EBS/fs_ne/EBSapps/log/adop/<adop_sessionID>/<cleanup_folder>/<context_name>/adzdcleanuprp.out.
```

Note: The analytics parameter should only be used when required, because of the extra processing needed.

Online Patching Cycle Steps - Abort Phase:

If for some reason either the prepare or apply phase failed or gave problems, you can abort the patching cycle at either of these points by running the following command:

```
$ adop phase=abort
```

You will be prompted for the APPS username and password.

In the abort phase, adop does the following:

1. Checks that the environment is set to the run APPL_TOP.
2. Checks if any previous session has an incomplete session, i.e. whether the abort call is valid.
3. Checks for the existence of a patch edition and drops one if it exists.
4. Checks concurrent program status and if necessary cancels a concurrent program submitted in earlier run.
5. Deletes the rows inserted for the pending session ID from the ad_adop_sessions and ad_adop_session_patches tables.

Be aware of the following important points:

- After running abort, you must always run a full cleanup (with the command `adop phase=cleanup cleanup_mode=full`). This will remove any columns that were added by the patch but are no longer needed because of the abort. If they are not removed, they may cause problems in a later patching cycle.
- If you run the prepare phase on a single node of a multi-node environment, and later have to abort that adop session, you must run the abort command on **all** the nodes (not just on the one where you ran prepare).
- In a multi-node environment that uses a non-shared file system, you must run the command `adop phase=abort autoskip=yes` rather than simply `adop phase=abort`

What's Next

This section covers a variety of tasks that may apply either to individual online

patching operations, or to your system setup as a whole. Diagnostic, troubleshooting, and reporting features are also described.

Manual Post-Installation Tasks

Traditionally, some patches have associated post-installation tasks, including recompilation of invalid packages, regenerating JAR files, and running AutoConfig. In an online patching environment such as Release 12.2 such tasks will normally be performed automatically in the apply phase.

If a post-installation patch step mentions any tasks that need to be performed explicitly, where they are run from depends on the type of patching:

- In a normal patching cycle, the steps should be executed from the *patch* file system after the apply phase.
- If the patch is being applied in hotpatch mode, the steps should be executed from the *run* file system after the apply phase.

Considerations When Applying Patches That Have Already Been Applied

As mentioned in the adop Options, page 2-36 section of Chapter 2, if you try to apply a patch that has already been applied and do not specify the forceapply parameter, adop will display an error like this:

```
[WARNING] Skipping the application of patch 14125999_AR since it has
been already applied
[WARNING] Hint: Patches can be applied again by specifying
options=forceapply when invoking adop
```

There are two more scenarios that may apply in this kind of situation.

- When a failed patch session is restarted with `abandon=no, restart=yes`, the patches applied in current adop session will *not* be applied even if `options=forceapply` is specified. For example, you run the command `adop phase=apply options=forceapply patches=1111,2222`, and application of patch 1111 is successful but patch 2222 fails. After correcting the problem, you try to rerun adop with the command `adop phase=apply options=forceapply patches=1111,2222 abandon=no, restart=yes`. In this example, patch 1111 would be skipped as it had successfully been applied, and application of patch 2222 would resume. If you wanted to apply patch 1111 again, you would need to specify `abandon=yes, restart=no`.
- If you apply multiple patches with `merge=yes`, and you do not specify `options=forceapply`, the patches will be applied *only* if at least one of the patches has not been successfully applied before.

Note: This check will be performed for AD and non-AD patches separately, as adop applies these two categories of patch in different sessions.

Example Multi-Node Patching Scenarios

The following scenarios illustrates various permutations in patching systems that employ multi-node application tiers. For each phase, they show whether adop should be run on a particular node, and if so, with what options.

Scenario 1: Shared file system, allowing adop to ssh to other node

	prepare	apply	finalize	cutover	cleanup
Run on admin node?	Y (allnodes=yes action=db)	Y (allnodes=yes action=db)	Y (allnodes=yes action=db)	Y (allnodes=yes action=db)	Y (allnodes=yes action=db)
Run on other node?	N	N	N	N	N

Scenario 2: Shared file system, not allowing adop to ssh to other node

	prepare	apply	finalize	cutover	cleanup
Run on admin node?	Y (allnodes=no action=db)	Y (allnodes=no action=db)	Y (allnodes=no action=db)	Y (allnodes=no action=db)	Y (allnodes=no action=db)
Run on other node?	Y (allnodes=no action=nodb)	Y (allnodes=no action=nodb)	N	Y (allnodes=no action=nodb)	N

Scenario 3: Non-shared file system, allowing adop to ssh to other node

	prepare	apply	finalize	cutover	cleanup
Run on admin node?	Y (allnodes=yes action=db)	Y (allnodes=yes action=db)	Y (allnodes=yes action=db)	Y (allnodes=yes action=db)	Y (allnodes=yes action=db)
Run on other node?	N	N	N	N	N

Scenario 4: Non-shared file system, not allowing adop to ssh to other node

	prepare	apply	finalize	cutover	cleanup
Run on admin node?	Y (allnodes=no action=db)	Y (allnodes=no action=db)	Y (allnodes=no action=db)	Y (allnodes=no action=db)	Y (allnodes=no action=db)
Run on other node?	Y (allnodes=no action=nodb)	Y (allnodes=no action=nodb)	N	Y (allnodes=no action=nodb)	N

Configuration Management and Online Patching

The following guidelines apply to making configuration changes to Oracle E-Business Suite in the context of online patching. They particularly apply to the technology stack and application components that reside in the file system.

The two basic scenarios are *online* and *offline* configuration changes. Each will be considered in turn.

Online configuration changes are performed within the context of an online patching cycle. This is the recommended strategy.

First, you prepare your system by running the `adop phase=prepare` command. You then make the desired configuration changes to the patch file system. They may

include:

- Oracle WebLogic Server configuration changes
- HTTP Server configuration changes
- File system changes performed by the AD utilities

After making the configuration changes, you must run the command `adop phase=cutover` to promote them.

You must also run the command `adop phase=fs_clone` to propagate the configuration changes to the secondary file system.

Offline configuration changes are applied directly to the run file system, outside an online patching cycle. You can use the `adop -status` command to verify that no patching cycle is currently active. After making the desired configuration changes, you must explicitly run the `adop phase=fs_clone` command to propagate the changes to the patch file system.

Important: This offline scenario will require a period of downtime for users.

Customizing Patch File System Backup Count

A new AutoConfig context variable, `s_fs_backup_count`, is used to specify the 'Patch File System Backup Count'. This variable denotes the number of backups of the patch file system that are to be preserved by `adop`. The variable is used during the `fs_clone` phase, where the existing patch file system is backed up before it is recreated from the run file system.

Valid values for the `s_fs_backup_count` variable are 0-9. A value of 0 (the default) will not preserve any patch file system backups. A value of 1 will preserve the latest patch file system backup, a value of 2 will preserve the latest two backups, and so on. You can set this value as needed for your system.

Note: For more information about AutoConfig, see Technical Configuration chapter of *Oracle E-Business Suite Setup Guide*.

Integrating Your Custom Tasks Into the Online Patching Cycle

You may have business-specific tasks specific that need to be performed before, during or after a patching cycle. Support for such tasks is currently provided by callout points at the beginning and end of the cutover phase of the online patching cycle. This support will be extended in future releases of Oracle E-Business Suite.

Diagnostics and Troubleshooting

This list highlights known issues and other problems. It will be expanded as necessary in future editions of this book.

- **Defaults File Becomes Corrupt**

As mentioned earlier, an example defaults file is located here:

```
$APPL_TOP/admin/<SID>/adalldefaults.txt
```

If this file becomes corrupt, running AutoConfig will automatically instantiate a new copy.

- **Cutover Fails**

Consider a patching operation where cutover fails with `cutover_status=3`. If you attempt to resume the failed session it may fail with an 'Invalid Credentials' error: if so, this will be because the database patch edition has already been promoted to be the new run edition.

To resume and complete cutover successfully, you need to run the command:

```
$ adop phase=cutover action=nodb
```

If cutover fails in a multi-node environment during force startup of the application tier services, you need to restart it on each node using the command:

```
adop phase=cutover allnodes=no
```

- **Concurrent Processing Log Files Cannot Be Created**

If you have chosen to set the `$APPLLDM` environment variable to 'product' (so the concurrent processing log and out files will be written to product-specific directories under `$APPLCSF`), you may experience a known issue where concurrent processing log files cannot be created. This will be fixed in a future release. In Release 12.2.0, the workaround is to create the following directories under `$APPLCSF`: `sqlap/log`, `sqlap/out`, `sqlgl/log`, `sqlgl/out`, `ofa/log`, and `ofa/out`.

Note: For more information on options for setting `$APPLLDM`, refer to the Patching and Management Tools chapter of *Oracle E-Business Suite Concepts*.

- **File System Cloning Fails**

Consider the following scenario:

1. You attempt to apply patch 1111, but patch application fails.
2. A new patch, 2222, is supplied as a replacement for 1111.
3. You run the `adop prepare` phase for 2222, which performs an `fs_clone`. This fails.

The result is that the online patching cycle is stuck in the `fs_clone` phase. Using hotpatch mode is not a solution, as it cannot be initiated part way through a

patching cycle.

The preferred solution to this situation is as follows:

1. You abort the adop session with the failed fs_clone.
2. You apply patch 2222 in hotpatch mode to the run file system.
3. Oracle fixes the issue with fs_clone and provides patch 3333.
4. You apply patch 3333 in hotpatch mode to the run file system.
5. You run adop with the fs_clone phase. If it fails, you go to step 3 again.
6. You run an empty patch cycle to switch to the other file system and apply all hotpatch mode patches there as well.

- **Nodes Are Abandoned**

In a multi-node environment using remote adop invocation, there are different scenarios that can apply to abandoned nodes.

The most general case is where adop is running remote commands from the master node, and a command fails on a secondary node. You will see the message: 'Remote execution failed on Node: <nodename>'. In such a case, adop will not attempt to run further remote commands on that node. You will need to check the adop log files for that node, identify the problem, and then run the relevant adop commands manually.

A more complex scenario is as follows. Assume there are three nodes, A, B, and C, with Node A being the admin (master) node and also the only essential node (that is, it is running all the essential Oracle E-Business Suite services). If a command fails on Nodes B and C, but the patching cycle is complete on Node A, you have the following three options.

- You can run the relevant adop commands manually on Node B and Node C, and thereby complete the patching cycle.
- You can ignore Node B and Node C, and start a new patching cycle. In this case, when adop detects an abandoned node it displays two prompts.

The first prompt allows you to continue with the current patching cycle, but gives a warning:

```
This node has been abandoned in a previous patching cycle.  
You may choose to continue with current patching cycle by  
answering 'yes' below. This may lead to a state where this  
node will be out of sync with the others.
```

If you choose to continue with the current patching cycle, Node B and Node C will be marked as abandoned. You will subsequently either have to clone them

from the master node, or manually run the relevant adop commands on them (also see third option below).

The second prompt gives a further option:

If you want to abandon <nodename> node for this session as well please select 'yes' below.

This allows you to abandon that particular node in the current patching cycle as well. If you select 'yes', adop abandons the node for the current cycle. If you selects 'no', adop assumes that you fixed that particular node by cloning from another node before the patching cycle began, and brings the node to the level of admin node.

- You can clone Node B and Node C from Node A after the patching cycle is complete. In this case, adop will subsequently display the two prompts from option 2 above when it detects the abandoned nodes. Answer 'yes' and 'no' respectively to the prompts will cause adop to consider Node B and Node C as no longer abandoned, and resume normal processing.

Reporting

To help diagnose issues, or simply gain knowledge about the status of your system, you can run the *Online Patching Diagnostic Reports* utility, `$AD_TOP/bin/adopreports.sh`.

The `adopreports` utility is invoked by entering the command:

```
$ adopreports <APPS username> <APPS password>
```

This displays the `adopreports` Main Menu:

```
Online Patching Diagnostic Reports Main Menu
-----
1.  Run edition reports
2.  Patch edition reports
3.  Other generic reports
4.  Exit
```

Choosing option 1 from the Main Menu displays the Run Edition Reports Sub Menu:

```
Run Edition Reports Sub Menu
-----
1.  All
2.  Count of uncovered objects per edition
3.  List of uncovered objects per edition
4.  Cleanup status - summary of covered objects per edition,etc.
5.  Show covered objects count per edition.
6.  Show list of covered objects per edition.
7.  Back to main menu
```

Choosing option 2 from the Main Menu displays the Patch Edition Reports Sub Menu:

Patch Edition Reports Sub Menu

1. All
2. Patch status - new/changed objects
3. Objects patch in the current edition
4. Table manager status
5. Back to main menu

Choosing option 3 from the Main Menu displays the Other Generic Reports Sub Menu:

Other Generic Reports Sub Menu

1. Editions summary
2. Editioned objects summary
3. Free space in important tablespaces
4. Status of critical AD_ZD objects
5. Actual objects in current edition
6. Objects dependencies
7. Objects dependency tree
8. Editioning views column mappings
9. Index details for a table
10. Inherited objects in the current edition
11. All log messages
12. Materialized view details
13. Database sessions by edition
14. Table details (Synonyms, EV, etc.)
15. Count and status of DDL execution by phase
16. Back to main menu

Performing Interactive Patching

Patches and updates to the Oracle E-Business Suite file system or database are applied using the *adop* utility, which identifies the servers set up during your installation and performs the actions required by the patch on each APPL_TOP. In a shared application tier file system, changes made during patching sessions on one node are immediately available on all nodes.

You can apply a patch interactively or non-interactively. *Interactive* patching (the default patching method) means that you supply all the information that *adop* needs by responding to a series of prompts. You can also apply a patch *non-interactively* to avoid having to respond to some of the *adop* prompts and to accommodate special types of patches. See: *adop*, page 2-2 and Performing Non-Interactive Patching, page 3-38.

See: Preparing your System for Patching, page 2-21.

Applying a Patch Interactively:

Requirement: How do I apply a patch?

After you have determined that you need to patch your system, download the patch.

Patches may require prerequisite patches and manual steps. Use PAA to generate customized instructions. These instructions contain all the required manual steps that are specific to your system.

Use adop to apply the patch. Apply the unified driver to all APPL_TOPs. adop determines which actions are required for the current APPL_TOP.

Note: Some of the installation instructions generated by PAA may specify pre-install mode. If so, follow the instructions in Pre-Install Mode, page 2-32.

To apply a patch

This procedure describes the typical steps for applying a patch. Subsequent procedures describe command line options that change the default behavior of adop.

1. Log in as applmgr (Applications file system owner) and set the environment:

UNIX:

The environment file is typically APPS<CONTEXT_NAME>.env, and is located under the APPL_TOP. From a Bourne, Korn, or Bash shell, enter the following:

```
$ . APPS<CONTEXT_NAME>.env
```

Windows:

Run %APPL_TOP%\envshell.cmd using either Windows Explorer or the Run command from the Start menu. This creates a Command Prompt window that contains the required environment settings for Oracle E-Business Suite. Run all subsequent commands in this Command Prompt window.

If you are running on a Windows platform, ensure that all necessary tools are installed properly. In addition, all %JAVA_TOP% and %JAVA_TOP%\loadjava.zip files are included in the set CLASSPATH statement of %APPL_TOP%\admin\adovars.cmd.

2. Create a patch top directory, if it does not already exist. Download the patch into a staging directory and unzip the patch into the patch top directory. Do not use the patch subdirectory under the <PROD>_TOP directories as your patch top directory.
3. Change directory to the patch top directory where you unzipped the patch.
4. Review the readme file in the patch top directory.

Review the readme file (README.txt or README.html). It contains an abstract of the patch. If the patch contains manual steps, the readme file will contain instructions for running Oracle Patch Application Assistant (PAA) to generate customized manual steps for your system.

5. If indicated in the patch readme, run PAA to generate customized instructions for applying the patch. You will need to provide the location of your patch top directory and the applmgr password.

```
$ admsi.pl
```

Perform the manual steps contained in the customized instructions generated by PAA. Additional steps may also be detailed depending on the patch, the state of your system, and the products you have installed.

Perform the following steps, in addition to the steps detailed in the customized instructions, to apply the patch.

6. Start `adop` from the patch top directory (the directory where you downloaded the patch files).

Important: When invoked, `adop` sets the environment so that it always runs from the run edition of the file system. You can identify which edition you are in by checking the value of the `FILE_EDITION` environment variable.

For interactive mode, use the command

```
$ adop phase=apply
```

and follow the subsequent `adop` prompts.

See *Command Line Arguments*, page 2-33.

7. Respond to the `adop` prompts that appear in interactive mode. The following information is required to apply the patch:
 - Name of the `APPL_TOP` where you want to apply the patch
 - Log file name: Select a name for the log file, for example, `u<patchnum>.log`
 - Email where you want to receive notifications
 - Batch size (default is 1000)
 - Database name
 - Patch top directory where you unzipped the patch
 - Driver file name: Provide the name of the driver file located in the patch top directory, for example, `u<patchnum>.drv`

8. Apply the driver.

At the `adop` prompt for the driver name, specify the name of the driver.

9. Review customizations.

Customized files are registered on the OAM Register Flagged Files page. If `adop` displays a message indicating that previously registered, customized files will be replaced by the patch, review those files in the Register Flagged Files page to

determine if the customizations need to be reregistered or removed. See: Register Flagged Files, page 1-9.

10. After adop exits, review the log files.

Review the adop log file after the application of each driver file for warnings or errors. The log file (named u<patchnum>.log in step 9 is located in <APPL_TOP>/admin/<SID>/log. In addition, some patch tasks may create separate log files in the same directory. If the patching process used multiple workers, each worker creates its own log file. You can also use the View Log Files feature in Timing Reports to view the files. See: Log Files, page 2-23 and View Log Files, page 5-28.

11. Preallocate space for packages, functions, and sequences (optional).

If adop has modified Oracle E-Business Suite database objects, you may want to run ADXGNPIN.sql and ADXGNPNS.sql to allocate space ("pin") for new packages and sequences in the Oracle System Global Area. These scripts are located in AD_TOP/sql.

See: Pre-allocating Space for Packages and Functions in Maintaining the Database, page 7-14.

12. Verify that the patch was applied successfully.

Applying Unified Drivers:

Requirement: I received a patch that contains a unified driver. However, the instructions state that I should only run the database portion of the patch.

To apply only a portion of a unified driver

1. Follow the instructions in Steps 1 through to 5 in Applying a Patch Interactively, page 3-33.

2. Run adop in interactive mode with the command:

```
$ adop phase=apply
```

When prompted by adop, enter an applicable option such as:
options=databaseportion

See: Command Line Arguments, page 2-33.

3. At the prompt for the driver name, specify the unified (u) driver. adop runs the driver, applying only the database portion of the patch.
4. Respond to the adop prompts. See: adop, page 2-2.
5. Finish applying the patch as directed in Applying a Patch Interactively, page 3-33.

Testing a Patch Before Applying It:

Requirement: How do I test the effects of a patch on my system before I apply it?

One way to see how applying a patch will affect your system is to first apply it on a test system. If you do not, or cannot, use a test system, you can apply the patch on your production system using the adop test mode argument, *apply=no*, in the input_file. adop lists the actions it would take, but does not actually perform any of the actions.

In test mode, adop reads and validates the patch driver file, reads product file driver files, extracts object modules from product libraries (for version checking), performs version checking, and runs AutoConfig (in test mode). It does not, however, update the database or file system..

Note: For more information on AutoConfig, see Oracle E-Business Suite Technical Configuration in *Oracle E-Business Suite Concepts*.

To determine how a patch will affect the files in your system, use the Patch Impact Analysis Report in Patch Wizard. See: Determining Patch Impact on System Files, page 3-50.

To test a patch

1. Follow steps 1-5 in Applying a Patch Interactively, page 3-33.
2. When directed to run adop, do so with the test mode argument `apply=no` in the input_file.
3. Follow steps 7-10 in Applying a Patch Interactively, page 3-33.

Enabling Password Validation:

Requirement: How can I validate passwords before I apply a patch?

To reduce the time it takes to apply a patch, adop (by default) does not validate Oracle schema passwords. If you need to enable password validation, you can do so by supplying the validate option, *options=validate*, in the input_file when you run adop.

If you are applying multiple patches, you can still use AD Merge Patch, page 2-43 to combine the patches (where compatible) so that you apply them in a single online patching cycle. In this case, you need to validate passwords only once.

If you have several patches that cannot be merged, you can save time by turning on the validate option only for the application of the first patch, and then leaving it off for the subsequent patches.

Important: The preferred strategy in Oracle E-Business Suite Release 12.2 is to allow the adop utility to perform patch merging itself, and thereby avoid using AD Merge Patch.

To validate passwords

1. Follow the instructions in Applying a Patch Interactively, page 3-33.
2. When directed to run `adop`, specify the `validate` option by adding it to the `input_file` as `validate=yes`.
3. Complete the remaining steps in Applying a Patch Interactively, page 3-33.

Applying Emergency Patches:

Requirement: How can I apply an emergency patch in an online patching environment?

An emergency patch may be applied using `adop` in *hotpatch mode*, instead of in an online patching cycle.

Performing Non-Interactive Patching

You can apply patches interactively or non-interactively. *Interactive* patching means that you supply basic information that `adop` needs by responding to a series of prompts. See: Performing Interactive Patching, page 3-33.

Applying a patch *non-interactively* substantially reduces the need for user intervention when `adop` processes patching tasks. You create a defaults file that contains much of the information you would have supplied at the `adop` prompts. Then, when you run `adop`, you specify the name of the defaults file, the location of the patch top directory, the name of a driver file, and other parameters on the command line. `adop` performs the remaining actions based on the information in the defaults file.

Caution: You should always back up the file system and database before you apply large patches such as release update packs (RUPs), product family RUPs, or pre-upgrade patches.

Applying a Patch Non-Interactively:

Requirement: How do I apply a patch non-interactively?

Instead of responding to `adop` prompts each time you initiate a patching session, you can store the responses in a defaults file. Then you specify the name of the defaults file when you run patches non-interactively. As it runs, `adop` uses the responses to complete the information for the corresponding prompts, and completes patch processing with little or no user intervention.

To apply a patch non-interactively

1. Ensure you have a suitably-customized defaults file, `$APPL_TOP/admin/<SID>_patch/adalldefaults.txt`, on both the run

APPL_TOP and patch APPL_TOP.

Important: The defaults file is created and automatically maintained by adop. It does not have to be created manually as was the case in releases prior to 12.2.

The defaults file is not specified on the adop command line, but in the input_file that is included on the command line. The input_file contents must include the following required parameters:

```
patches=<patch number>
workers=<number of workers>
patchtop=<directory where patches are staged>
```

Note: You can specify patches just as numbers, or (for some specialized patches) a more complex format containing a colon. Both formats can be used in a single patching command. For example, patches=111,222_KO:u222.drv. See adop, page 2-2 in this book.

2. Run the command:

```
$ adop phase=apply input_file=<input_file.txt>
```

Applying a Single Patch Driver:

Requirement: How do I apply a single patch driver non-interactively?

If you have created a defaults file, specify adop to run non-interactively and specify the location and name of the defaults file and the driver.

To apply one or more patch drivers

1. Ensure the defaults file exists, as described in Applying a Patch Non-Interactively, page 3-38. Check its contents meet you needs.
2. Follow steps 1 through to 5 in Applying a Patch Interactively, page 3-33.
3. Run the adop command, adding the location of the input_file.

```
$ adop phase=apply input_file=<input_file.txt>
```

The input_file file might have contents like this:

```
defaultsfile=<defaults file on patch APPL TOP>
workers=<number of workers>
patchtop=<location of patches to be installed>
patches=<patch number 1>:<driver file 1>.drv, <patch number
2>:<driver file 2>.drv ...
```

4. Perform the remaining steps in Applying a Patch Interactively, page 3-33 (as necessary).

Applying a Non-Standard Patch:

Requirement: I need to apply a patch that was not created with the standard patch naming convention. I would like to apply it non-interactively.

A *non-standard* patch is one where the structure is standard, but the naming convention is not. That is, the last component of the patch directory is not a 6- to 8-digit number, or the patch driver files are not named **<patchnum>.drv*, or both. Most merged patches are non-standard because of the way they are named.

To apply a non-standard patch

1. Create the *input_file* as described in Applying a Patch Non-Interactively, page 3-38.
2. Follow steps 1 through to 5 in Applying a Patch Interactively, page 3-33.
3. Run the *adop* command as described in Applying a Single Patch Driver, page 3-39. For the *driver=<values>* argument, use a comma-separated list of the patch driver names.
4. Perform the remaining steps in Applying a Patch Interactively, page 3-33 (as necessary).

Restarting a Non-Interactive adop Session:

Requirement: *adop* failed with an error while I was applying patches non-interactively. I have resolved the issue that caused the error and want to restart the failed session.

When *adop* is running non-interactively and encounters an error, it exits to the operating system and reports a failure. The *restart* argument is intended specifically for this circumstance. When *adop* sees the *restart=yes* argument, it assumes that there is an old session, and expects to find one. If it does not, it will fail.

To restart a non-interactive adop session

1. Look through the log files, diagnose the error, and fix it.
2. Run the same *adop* command you used initially, but add *restart=yes* to the *input_file*. For example:

```
defaultsfile=<defaults file on patch APPL TOP>
workers=<number of workers>
patchtop=<location of patches to be installed>
patches=<patch number 1>:<driver file 1>.drv, <patch number
2>:<driver file 2>.drv
restart=yes
```

Abandoning a Non-Interactive adop Session:

Requirement: *adop* failed with an error while I was applying patches non-interactively. I do not want to restart the failed session, but would rather apply another patch non-interactively.

Running adop with the *interactive=no* and *restart=yes* input_file options will restart the previously incomplete session, which is not what you want.

To start a completely new adop session when there is an existing failed session, specify *abandon=yes* in the input_file. With this option, adop deletes the restart files and any leftover database information from the failed session.

Caution: If you use the *abandon=yes* argument, you cannot subsequently restart the failed session as the restart files are no longer available. Do not specify *abandon=yes* if you might want to restart the session later.

To abandon a non-interactive adop session

1. Verify that you do not want to restart the previous failed session.
2. Run the same adop command you used initially, but add *abandon=yes* to the input_file. For example:

```
defaultsfile=<defaults file on patch APPL TOP>
workers=<number of workers>
patchtop=<location of patches to be installed>
patches=<patch number 1>:<driver file 1>.drv, <patch number
2>:<driver file 2>.drv
abandon=yes
```

Patching HRMS Legislative Data

Special instructions apply to installation of the HRMS Legislative Data patch. These include adop options that will:

- Prevent attempts to synchronize the run and patch file systems when applying a patch file (in this case hrglobal.drv) that includes database operations only.
- Ensure that adop will proceed with the relevant patching request, instead of deeming the upgrade to have been installed already.

Important: Before applying the HRMS Legislative Data patch, refer to the full instructions provided in My Oracle Support Knowledge Document 1469456.1, *Datainstall and HR Global Application: 12.2 Specifics*.

Patching NLS Systems

These patching procedures apply regardless of whether you are running American English (US) and one additional language, or American English (US) and several additional languages. If your system uses multiple languages, you can use AD Merge Patch to create merged patches in whichever of the following ways suits you best:

- A single, merged patch that contains all languages (including US English).
- One merged patch for US English and a second merged patch for all other languages.
- A separate merged patch for each language.

Before the introduction of online patching, the choice of which strategy to follow largely depended on the downtime that was acceptable: for example, the first option was straightforward but required the greatest downtime, and the third was the most complex but allowed users of a particular language to resume their work as soon as the relevant patch was applied. The second option often provided the best compromise between easy application and minimum downtime. Now, however, use of online patching means that downtime is greatly reduced as a factor when determining the strategy that most closely suits the organization's needs

When merging multiple language patches, AD Merge Patch converts the character set according to the NLS_LANG variable in the Oracle E-Business Suite environment file. If you changed your character set since the initial installation, you might need to update the NLS_LANG variable. If this variable is not set properly, run AutoConfig from Oracle Applications Manager to update the Oracle E-Business Suite context with the correct character set information, then run the appropriate AutoConfig command to recreate the Oracle E-Business Suite environment file. Reset the environment using the new environment file before merging patches.

For more information, see: AD Merge Patch, page 2-43.

Applying a Single Patch to an NLS Installation:

Requirement: I need to apply a single patch to an Oracle E-Business Suite NLS installation.

If an Oracle E-Business Suite system contains languages other than American English (US), the recommended method is to apply the US patch first and then apply the translation patch for each installed language. If you have installed more than one additional language, you can merge all the translation patches and apply them as a single, merged NLS patch.

You can also merge US patches with the additional language patches. However, depending on your system topology, it may be necessary to keep the US and non-US patches separate.

To apply a single patch to an NLS installation

This procedure assumes that you will apply US and language patches separately.

1. Use adop to apply the patch driver of the US patch.
2. Use adop to apply the patch drivers of each NLS patch. If you have merged the individual NLS patches for a system that runs multiple languages, apply the driver

for the merged NLS patch. See: *Applying a Patch Interactively*, page 3-33.

Applying Multiple Patches to an NLS Installation:

Requirement: I need to apply several patches to an Oracle E-Business Suite NLS installation.

If an Oracle E-Business Suite system contains multiple languages other than American English (US) and you are applying multiple patches for each language, the recommended method is to merge all US patches into a single patch and all patches for every non-US language into a single patch. Then apply the merged US patch, followed by the merged language patch.

You can also merge US patches with the additional language patches or merge each language in separate language-specific patches. Depending on your system topology, it may be necessary to keep the US and non-US patches separate. This procedure assumes that you will apply US and language patches separately.

To apply multiple patches to an NLS installation

This example assumes the system has American English, French, and German installed.

1. Use AD Merge Patch to merge the US (American English) patches into a single patch.
2. Use AD Merge Patch to merge the French and German patches into a single NLS patch.
3. Use adop to apply all drivers of the merged US patch.
4. Use adop to apply all drivers of the merged NLS patch.

Keeping Your System Current

Each time you apply a patch, AutoPatch stores the associated information in the Oracle Applications Manager (OAM) patch history database. The OAM Patch Wizard and Applied Patches tools provide graphical user interfaces that you can use to query the database for a complete history of patches applied to your system, to search for the patches you have already applied, and to determine existing patches that should be applied to keep your system current. Patch Wizard determines which recommended patches you should apply to your system, and the impact of applying these patches.

Before running Patch Wizard, you must set up My Oracle Support credentials. You must also set up preferences and filters that govern the way you download patches. To see how to complete these one-time tasks, as well as learn about navigating the Patch Wizard pages and submitting requests, see: *Patch Wizard*, page 2-2.

Creating a List of Recommended Patches:

Requirement: How do I determine if there are patches that I have not yet applied?

Patch Wizard creates a list of patches by comparing the patches in the patch history database against a list of recommended patches in a Patch Information Bundle file downloaded from My Oracle Support. It then determines which of the recommended patches you should apply to your system and reports the contents of the patch and the files that it will update when applied.

It does not report on all available patches, but only patches at the current codeline, such as high-priority patches, and those that update your system to a new codeline (pre-upgrade patches).

To see a list of patches recommended for your system

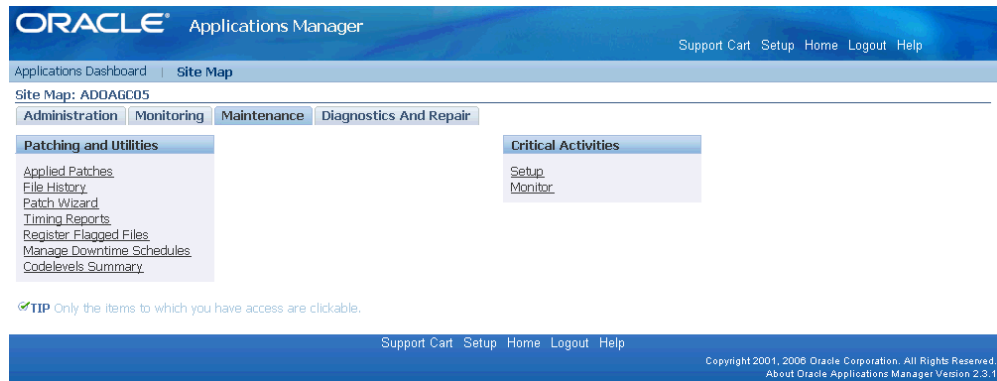
1. Access Oracle Applications Manager.

Follow the instructions in *Accessing Patch Wizard*, page 2-2 to access OAM. All procedures in this section begin with the Site Map.

2. Access the Patch Wizard home page.

From the Site Map (Maintenance tab), click Patch Wizard under the Patching and Utilities heading.

Site Map Page



The Patch Wizard home page appears.

3. Submit a request for recommended patches.

From the Recommend Patches page, select a patch filter. Use the magnifying glass icon to see a list of available patch filters.

Recommend Patches Page - Top

ORACLE Applications Manager

Support Cart Setup Home Logout Help

Applications Dashboard | Site Map

Applications System:121test > Patch Wizard >

Recommend Patches : 121test

Last Updated : 18-Mar-2009 14:39:28
Staging Directory /slot/ems1227/appmgr/stage
Oracle MetaLink User ID murali.kumanduri@oracle.com

Options

Patch Wizard automatically downloads patches or the Patch Information Bundle from MetaLink before using them for analysis or recommendations. If it is unable to download files from MetaLink, it will try to use existing files in the staging directory.

OK Cancel

Select

Create Recommendation

Using Patch Filter

Analyze Aggregate Patch Impact

The Patch Information Bundle created on 18-Mar-2009 08:03:49 was uploaded on 18-Mar-2009 13:19:59.

Analyze Specific Patches

Patch

(Enter Patch numbers, separated by commas ex: 1234567, 8765432_R12.AD.A, 8888888.AD.B, Maximum number of allowed characters is 175)

Analyze Aggregate Patch Impact

The Patch Information Bundle created on 18-Mar-2009 08:03:49 was uploaded on 18-Mar-2009 13:19:59.

After you have entered the requested information, click *OK*. The results of your request are shown in the Results section of the Patch Wizard main page. You can also schedule the request for a future date.

4. Track the status of your request.

From the main page, you can track the status of your recommended patch request. Click the Job Status icon for the Recommend/Analyze Patches task.

Job Status Page

ORACLE Applications Manager

Support Cart Setup Home Logout Help

Applications Dashboard | Site Map

Applications System:ADOACC04 > Patch Wizard >

Job Status:Recommend/Analyze Patches:ADOACC04

Last Updated : 11-Aug-2006 14:37:05

New Search Modify Search

Request ID	Program	Phase	Status	Requestor	Requested Start Date	Duration	Wait Time
<input type="checkbox"/> Hide S265804	PatchWizard - Recommend Patches (Request Set SubmitAdvisorCriteria)	Running	Paused	SYSADMIN	11-Aug-2006 14:35:44	00:00:01	00:00:01

Short Name: FDRSUB1243
Application Name: Applications DBA
Submission Date: 11-Aug-2006 14:35:44
Actual Start Date: 11-Aug-2006 14:35:45
Completion Date: 11-Aug-2006 14:35:45
Style: Portrait, Program Name: PatchWizard - Recommend Patches
Priority: 50
Responsibility: System Administration
Language: AMERICAN
Territory: AMERICA
Numeric Characters

Print to: noprint, Copies: 0, Language:

Notify Parameters: 50, 1408
Repeat Interval: PatchWizard - Recommend Patches
Completion Text:

Sub Requests Cancel View Diagnostics Manager Log Request Log Detailed Logs Environment

TIP Duration is the total time(HRS:MI:SS) the request has been running or the request ran.
 TIP Wait Time is the time(HRS:MI:SS) the request has waited.
 TIP To display the available actions on a request, please click on show details.

Add to Support Cart

Support Cart Setup Home Logout Help

The Job Status page displays summary information. If you click the Show/Hide icon corresponding to your request ID, the page displays more details. For more information about the fields and functions on this page, see: Patch Wizard, page 2-2.

Downloading Recommended Patches:

Requirement: How do I use Patch Wizard to download patches?

Patch Wizard can download patches based on either the list created by the "recommend patches" request or any list of patches entered in the Download Patches page.

The Download Patches page prompts you for information about the patches to download, then downloads them directly from My Oracle Support. The Merge Options section of this page defines how patches should be merged after they are downloaded.

To download patches using Patch Wizard

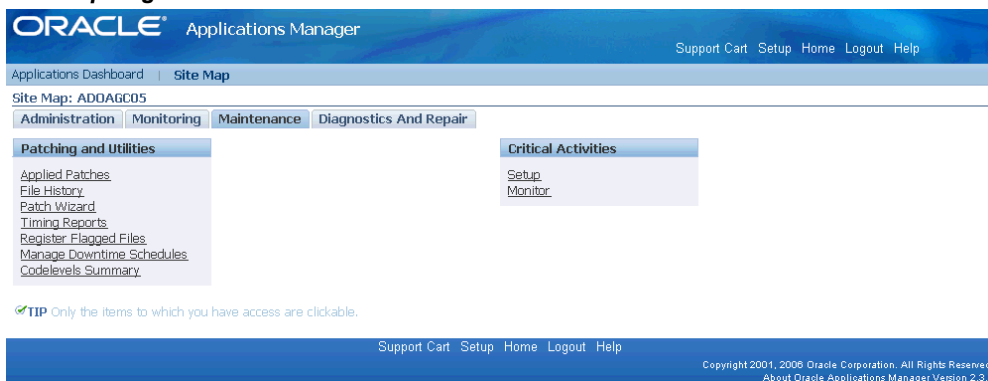
1. Access Oracle Applications Manager.

Follow the instructions in Accessing Patch Wizard, page 4-4 to access OAM. All procedures in this section begin with the Site Map.

2. Access the Patch Wizard home page.


From the Site Map (Maintenance tab), click Patch Wizard under the Patching and Utilities heading.

Site Map Page



The Patch Wizard home page appears.

Patch Wizard Home Page



- Submit a request to download patches.

Click the Tasks icon for Download Patches. The Download Patches page appears.

Download Patches Page - Top

[Support Cart](#) [Setup](#) [Home](#) [Logout](#) [Help](#)

Applications Dashboard | [Site Map](#)
 Applications System: r121test > [Patch Wizard](#) >

Download Patches : r121test
 Last Updated : 18-Mar-2009 14:48:20
 Staging Directory : /slot/ems1227/appmgr/stage
 Oracle MetaLink User ID : murali.kumanduri@oracle.com

Patch Selection

You must set up your Metalink credentials before downloading patches.

Patch List

(Enter Patch numbers, separated by commas ex: 1234567,8765432_P12.AD.A,888888.AD.B). Maximum number of allowed characters is 175)

Options:

- Download only
(Download the exact list of patches above)
- Download and Analyze
(Download only patches listed above that have not been applied and their prerequisite patches)
- Download, Analyze and Aggregate Patch Impact
(Download only patches listed above that have not been applied and their prerequisite patches, analyze them, and compute aggregate patch impact)

Merge Options
 Automatically merge downloaded patches
(Merge happens only if all patches are downloaded successfully)
 Merged Patch Name:
 Merging Strategy:

- One merged patch: US and non-US
- Two merged patches: US; non-US
- Multiple merged patches: US; language1;language2;...

On this page, list the patches you want to download in the Patch List field.

Another option is to click the Details icon for a recommended patch request in the Results section of the Patch Wizard home page.

Patch Wizard Home Page - Recommended Patches Results

Recommended Patches Results							
ⓘ Previously submitted Filter Names/ Patch Lists that do not appear in the Recommended Patches Results section have been purged according to the frequency setting in Purge Concurrent Request. Change the frequency setting in Applications Dashboard > Critical Activities if needed.							
Filter Name/ Patch List	Total (Applied & Unapplied)	Unapplied	Requested By	Completion Date	Run Status	Request Set	Details
Recommended Patches and New Codelevels	1	1	SYSADMIN	2006/Oct/11 02:51:41	Normal	280040	
Recommended Patches	0	0	SYSADMIN	2006/Oct/11 02:32:16	Normal	280014	
4502603	1	1	SYSADMIN	2006/Oct/09 00:09:09	Warning	279849	
4502400	1	1	SYSADMIN	2006/Oct/08 23:52:48	Warning	279809	
Recommended Patches and New Codelevels	1	1	SYSADMIN	2006/Oct/08	Normal	279792	

The Recommended Patches Results page for the recommended patch request appears.

Recommended Patches Results Page

ORACLE Applications Manager

Support Cart Setup Home Logout Help

Applications Dashboard | Site Map

Applications System: ADOAGC03 > Patch Wizard >

Recommended Patches Results : ADOAGC03

Last Updated : 10-Aug-2007 21:11:44

Patch Filter/Patch ID: Recommended Patches
 Requested By: SYSADMIN
 Completion Date: 10-Aug-2007 20:38:43
 View Aggregate Patch Impact:

Show Hidden Patches (with the check mark in the Hide Patch column)

I If the Show Hidden Patches checkbox is not selected, the number of patches displayed may be less than the number listed on the Patch Wizard page.
 Only patches selected on the current page can be downloaded.

Recommended Patches Results

Select Patch and ...

Select All | Select None

Select Patch	Product	Prerequisites	Codelevel Introduced	Status	PAA	Reason Recommended	Patch Description	Hide Patch	Included in Aggregate Patch Impact	Impact
<input type="checkbox"/>	6166150.A.2x	0	No	Unapplied	Yes	High Priority Patch	E-Business Tax: Consolidated upgrade script changes post R12	<input type="checkbox"/>	No	<input type="button" value="I"/>

Patches that Introduce New Codelevels

									Included in Aggregate
									<input type="button" value="I"/>

Select any number of recommended patches on this page and click the Download button. This populates the Patch List field in the Download Patches page with the selected patch numbers.

4. Set download options.

On the Download Patches page, set Merge options and indicate information about languages and platforms. If you choose to automatically merge patches while downloading, specify the merged patch name and the merging strategy. You can select the languages and platform of the patches to download. When you provide information in this section of the page, Patch Wizard only downloads patches that match the selected languages and platform. You can also schedule the download for a future date.

5. Submit request.

After you have entered the patch information, click OK.

6. Track the status of your request.

From the main page, you can track the status of your patch request. Click the Job Status icon for Download Patches.

Job Status Page

ORACLE Applications Manager

Support Cart Setup Home Logout Help

Applications Dashboard | Site Map

Applications System:ADOACC04 > Patch Wizard >

Job Status:Download Patches:ADOACC04

Last Updated : 11-Aug-2006 14:32:28

New Search Modify Search

Details	Request ID	Program	Phase	Status	Requestor	Requested Start Date	Duration	Wait Time
<input type="checkbox"/> Hide	S265600	DownloadPatches (Request Set) DownloadPatches	Completed	Error	SYSADMIN	11-Aug-2006 14:28:01	00:01:40	00:00:11

Short Name: FNDRSSUB1623
Application Name: Applications DBA
Submission Date: 11-Aug-2006 14:28:00
Actual Start Date: 11-Aug-2006 14:28:12
Completion Date: 11-Aug-2006 14:29:52
Priority: 50
Responsibility: System Administration
Language: AMERICAN
Territory: AMERICA
Numeric Characters:
Print to: Style: Portrait, Program Name: noprint, Copies: 0, Language:
Notify:
Parameters: 50, 2152
Repeat Interval: DownloadPatches
Completion Text: The set completed normally with outcome Error. The outcome was determined by the stage Submit Download Patches (10).

Restart Sub Requests View Diagnostics Manager Log Request Log Detailed Logs Output

TIP Duration is the total time(HRS:MI:SS) the request has been running or the request ran.
 TIP Wait Time is the time(HRS:MI:SS) the request has waited.
 TIP To display the available actions on a request, please click on show details.

Add to Support Cart

The Job Status page displays. If you click the Show/Hide icon corresponding to you request ID, the page displays more details. For more information about the fields and functions on this page, see: Patch Wizard, page 2-2.

Determining Patch Impact on System Files:

Requirement: Before I apply a patch, can I see which system files will be affected?

Patch Wizard provides a Patch Impact Summary page that shows the impact of a specific patch if applied to your system. It contains the following information: Patch Impact Analysis, Direct Impact Summary, and Indirect Impact Summary. By reviewing these results, you can see detailed information about files included in a patch, as well as the effect a specific patch will have on your existing system files. For example, you can see information about total files in the patch, the number and type of files that will be installed, and which existing files will be changed. See: Patch Wizard, page 2-2.

To view the information on the Patch Impact Summary page

1. Access Oracle Applications Manager.

Follow the instructions in Accessing Patch Wizard, page 2-2 to access OAM. All procedures in this section begin with the Site Map.

2. Access the Patch Wizard home page.

From the Site Map (Maintenance tab), click Patch Wizard under the Patching and Utilities heading.

Site Map Page

3. View recommended patches results.

From the home page, click the Details icon for an item in the Results section.

Recommended Patches Results Page

Task Name	Description	Tasks	Job Status
Patch Wizard Preferences	Set download, merge, and stage area preferences		
Define Patch Filters	Create custom patch filters		
Recommend/Analyze Patches	Submit requests for patch advice or analysis		
Download Patches	Submit requests to download patches		
Aggregate Patch Impact	Aggregate Patch Impact		

Filter Name	Operator	Value
Filter Name	contains	<input type="text"/>
Completion Date	is	<input type="text"/> (yyyy/MMM/dd)

Filter Name/ Patch List	Total (Applied & Unapplied)	Unapplied	Requested By	Completion Date	Run Status	Request Set	Details
Recommended Patches and New Codelevels	1	1	SYSADMIN	2006/Oct/11 02:51:41	Normal	280040	
Recommended Patches	0	0	SYSADMIN	2006/Oct/11 02:32:16	Normal	280014	
4502603	1	1	SYSADMIN	2006/Oct/09 00:09:09	Warning	279849	
4502400	1	1	SYSADMIN	2006/Oct/08 23:52:48	Warning	279809	

The Recommended Patches Results page for the recommended patch request appears.

Recommended Patches Results Page

ORACLE Applications Manager

Support Cart Setup Home Logout Help

Applications Dashboard | Site Map

Applications System: ADOAGC03 > Patch Wizard >

Recommended Patches Results : ADOAGC03

Last Updated : 10-Aug-2007 21:11:44

Patch Filter/Patch ID: Recommended Patches
 Requested By: SYSADMIN
 Completion Date: 10-Aug-2007 20:38:43

View Aggregate Patch Impact

Show Hidden Patches (with the check mark in the Hide Patch column)

! If the Show Hidden Patches checkbox is not selected, the number of patches displayed may be less than the number listed on the Patch Wizard page.
 Only patches selected on the current page can be downloaded.

Recommended Patches Results

Select Patch and ...

Select All | Select None

Select	Patch	Product	Prerequisites	Codelevel	Introduced	Status	PAA	Reason	Recommended	Patch Description	Hide Patch	Included in Aggregate Patch	Impact
<input type="checkbox"/>	6166150.A.zx	0		No		Unapplied	Yes	High Priority Patch		E-Business Tax: Consolidated upgrade script changes post R12	<input type="checkbox"/>	No	<input type="button" value="Impact"/>

Patches that Introduce New Codelevels

												Included in Aggregate

4. Access the Patch Impact Analysis page.

Clicking the Impact icon in the Recommended Patches Results page opens the Patch Impact Analysis page for the selected patch.

Patch Impact Analysis Page

The screenshot shows the Oracle Applications Manager interface. At the top, it says 'ORACLE Applications Manager' with navigation links for 'Support Cart', 'Setup', 'Home', 'Logout', and 'Help'. Below this is the 'Applications Dashboard' and 'Site Map'. The main content area is titled 'Patch Impact Analysis for Patch 4630372-R12: ADOAGC03'. It displays the patch description 'R12 izu: R12.IZU.A' and a 'Patch Readme' icon. Summary statistics show 'Total Files in Patch: 4' and 'Files to install: 4 (100.00%)'. Two summary tables are present: 'Direct Impact Summary' and 'Indirect Impact Summary'. The Direct Impact Summary table lists: Applications Patched (2), File Types Installed (2), New Files Introduced (4), Existing Files Changed (0), Flagged Files Changed (0), Existing Files Unchanged (0), and Non-US Language Patches Required (0). The Indirect Impact Summary table lists: Unchanged Files Affected (0 JSPs), Menu Navigation Trees Affected (0 Responsibilities, 0 Paths), and Diagnostics Tests to Re-Run (0 Test(s)). Below these tables are three tips (TIP) regarding analysis on unchanged files, prerequisite patches, and non-US language patches. An 'Add to Support Cart' button is located at the bottom left of the main content area. The footer contains 'Support Cart Setup Home Logout Help' and copyright information: 'Copyright 2001, 2006 Oracle Corporation. All Rights Reserved. About Oracle Applications Manager Version 2.3.1'.

Many of the line items on this page are links to detailed information about the impact of the patch on the system. For example, the File Types Installed value is a link to a page that lists the file types and the number of unchanged, changed, and new files in the file system as a result of applying the selected patch.

Creating Patch Recommendations Without an Internet Connection:

Requirement: How do I use the features of creating patch recommendations if I do not have access to an Internet connection?

You can run Patch Wizard without access to an Internet connection, if necessary, by downloading the Patch Information Bundle to a system which has Internet access. Once the download is complete, copy the Patch Information Bundle file to the Patch Wizard's staging directory. Then run Patch Wizard as you normally would based on the files you copied to the staging directory.

To create recommendations without using an Internet connection

1. Download the Patch Information Bundle to a system which has Internet access.
2. Set up a staging directory on a system that does not have Internet access. Patch Wizard must be able to read from and write to this staging directory.
3. Copy the Patch Information Bundle zip file to the staging directory. The zip file must be copied to a system that can access the Patch Wizard staging directory. If the staging directory is on a local disk, the zip file must be copied to the system where you run Patch Wizard. If the staging directory is on a shared (network) disk, it can be copied to any system with access to the shared disk.

4. Run Create Recommendations as you normally would from this point.

Analyzing Applied Patches

As you apply patches, AutoPatch records the actions in the Oracle E-Business Suite patch history database. You can query this database using the Oracle Applications Manager (OAM) Applied Patches feature, which provides easy access to reports based on your search criteria.

Note: Patch information is not stored in the database if the patch is applied in pre-install mode or test mode. Also, if patch application does not run successfully to completion, the associated information is neither uploaded to the patch history database nor available in the Applied Patches feature.

You enter search criteria on a search patches page, either Simple Search or Advanced Search. A summary report is displayed at the bottom of the search page.

Several detailed reports are also available, including Timing Details, Files Copied, Bug Fixes, and Action Summary. Most of these detailed reports have a standard layout. The top portion displays the criteria that were used for the search, and the bottom portion displays the results of the search.

See also Applied Patches in the OAM Help system.

Determining If a Patch Was Applied:

Requirement: Can I determine if a specific patch has been applied to my Oracle E-Business Suite system?

To determine which patches were applied, enter a patch ID in the Applied Patch Check area of the Software Updates page. You can perform a simple search by entering an ID or a series of IDs separated by commas.

To determine if a patch was applied

1. Access Oracle Applications Manager.

Follow the instructions in Accessing Patch Wizard, page 4-4 to access OAM.

2. Access the Software Updates page.

From the Applications Dashboard, click the Software Updates tab. The Software Updates page appears.

Software Updates Page

The screenshot displays the Oracle Applications Manager interface for the Software Updates page. At the top, there is a navigation bar with 'Support Cart', 'Setup', 'Home', 'Logout', and 'Help'. Below this, the 'Applications Dashboard' is visible, showing the current dashboard as 'ADDAGC01'. The main navigation tabs include 'Overview', 'Performance', 'Critical Activities', 'Diagnostics', 'Business Flows', 'Security', and 'Software Updates'. The 'Software Updates' section is active, showing 'Applications System Version: 12.0.0'.

The 'Applied Patch Check' section contains a help icon and a description: 'Use this function to determine if a patch has been or has not been applied.' Below this is a 'Patch' input field with a 'Go' button and a note: '(Enter Patch numbers, separated by commas)'. There are two columns: 'Applied' and 'Not Applied', both currently empty.

The 'Maintenance Activities' section features a table with the following data:

Task Name	Status	Last Updated	Details	Log Files
AutoPatch - u6655114.drv	✓	19-Dec-2007 18:22:56	[Details Icon]	[Log Files Icon]

3. Enter a patch ID.

In the Applied Patch Check area of the Software Updates page, enter a patch ID or a series of IDs separated by commas. Your queried ID appears in the corresponding column depending on whether it has been applied.

Searching for Patch Details:

Requirement: What information is available on the Patch Details report? How do I create the report?

From any Patch Summary report, you can click the Details icon for a selected row to open the Patch Details report, which displays summary information carried over from the Results portion of either the Simple Search or Advanced Search page.

This report also contains more specific information about the patch, including:

- Name of the driver file and the date and time it was applied
- Command line options used to run the file
- Platform of the driver file
- Location where the driver was run
- Report on whether a codelevel was introduced, and if so, which one

From the Patch Details page, you can also access additional information about a patch, including timing details, files copied, bug fixes, and a summary of actions performed.

See: Applied Patches, page 5-1. See also Applied Patches in the OAM Help.

To review patch details

1. Access Oracle Applications Manager.

Follow the instructions in Accessing Patch Wizard, page 4-4 to access OAM.

2. Create a Patch Summary report.

From the Site Map (Maintenance tab), click Applied Patches under the Patching and Utilities heading. From either the Simple Search or Advanced Search page, enter a patch number or a date range to create a Patch Summary report. Click *Go*.

3. Select the patch.

Click the Details icon in any selected row of the Patch Summary report. The Patch Details report appears.

Patch Details Report

The screenshot shows the Oracle Applications Manager interface. At the top, there's a navigation bar with 'ORACLE Applications Manager' and links for 'Support Cart', 'Setup', 'Home', 'Logout', and 'Help'. Below this is a breadcrumb trail: 'Applications Dashboard > Site Map > Applications System: ADOAGC05 > Applied Patches > Patch Details : 5336717.A : ADDAGC05'. The main content area displays patch details: 'Last Updated : 2006/Oct/11 15:44:46', 'APPL_TOP Name : ap6049rt', 'Patch : 5336717.A', 'Merged Patches : No', and 'Language : US'. On the right, it shows 'Drivers Applied : 1', 'Completion Date : 2006/Oct/06 13:45:39', and 'Patch Description : AD OAM DOC FOR RELEASE 12'. A note states: 'If the same patch is applied multiple times, this page displays all of them.' Below this is a table with tabs for 'Select Driver File and view...', 'Timing Details', 'Files Copied', 'Bug Fixes', and 'Action Summary'. The table has columns: 'Select', 'Driver File', 'Start Date', 'End Date', 'AutoPatch Options', 'Platform', 'Patch Top', and 'Codelevel Introduced'. One row is selected, showing 'u5336717.drv' with start date '2006/Oct/06 13:44:41' and end date '2006/Oct/06 13:45:39'. The 'AutoPatch Options' are 'novalidate, hotpatch', and the 'Platform' is 'GENERIC'. The 'Patch Top' is '/ SLOTS/ slot05/ appmgr/ patches/ 5336717'. There is an 'Add to Support Cart' button at the bottom right. The footer contains 'Support Cart Setup Home Logout Help' and copyright information: 'Copyright 2001, 2006 Oracle Corporation. All Rights Reserved. About Oracle Applications Manager Version 2.3.1'.

The report displays patch details such as driver files, start and end dates, and platform. It also provides access to other patch details related to the driver files, such as files copied and bug fixes. You can select a driver from the list, and click one of the additional detail buttons to see other reports.

4. View additional details.

As an example of the details that are available for a selected driver, click *Files Copied*.

Files Copied Report

The screenshot displays the Oracle Applications Manager interface. At the top, it says "ORACLE Applications Manager" with navigation links for "Support Cart", "Setup", "Home", "Logout", and "Help". Below this, there are tabs for "Applications Dashboard" and "Site Map". The main content area shows the "Files Copied" report for driver file "u5251362.drv : ADOAGC04". It includes a "Last Updated" timestamp of "08-Jun-2006 21:11:25". The report details include "Start Date: 26-May-2006 10:31:16", "End Date: 26-May-2006 10:34:19", "AutoPatch Options: hotpatch", "Platform: LINUX", "Patch Top: / SLOTS/ slot03/ appmgr/ patches/ 5251362", and "Driver File: u5251362.drv". A filter section shows "Product" set to "AD" and "Directory" set to "lib". Below this is a table with columns "Product", "Directory", "File", and "Version".

Product	Directory	File	Version
AD	lib	adphisto	120.14

At the bottom right of the report area, there is a button labeled "Add to Support Cart". The footer of the page contains the text "Copyright 2001, 2006 Oracle Corporation. All Rights Reserved. About Oracle Applications Manager Version 2.3.1".

For each file, the Files Copied report shows the product short name, the directory where the file was copied, the name of the file, and the version number. To view other information associated with the driver file, click the Patch Details link at the top of the page to return to the previous page.

As another example, click the Bug Fixes button.

Bug Fixes Report

The screenshot displays the Oracle Applications Manager interface. At the top, it says "ORACLE Applications Manager" with navigation links for "Support Cart", "Setup", "Home", "Logout", and "Help". Below this, there are tabs for "Applications Dashboard" and "Site Map". The main content area shows the "Bug Fixes" report for driver file "u5251362.drv : ADOAGC04". It includes a "Last Updated" timestamp of "08-Jun-2006 21:12:38". The report details include "Start Date: 26-May-2006 10:31:16", "End Date: 26-May-2006 10:34:19", "AutoPatch Options: hotpatch", "Platform: LINUX", "Patch Top: / SLOTS/ slot03/ appmgr/ patches/ 5251362", and "Driver File: u5251362.drv". A filter section shows "Bug Fix" set to "5251362" and "Product" set to "ad". Below this is a table with columns "Bug Fix", "Product", "Applied", and "Remarks".

Bug Fix	Product	Applied	Remarks
5251362	ad	Y	
4654046	ad	Y	

At the bottom right of the report area, there is a button labeled "Add to Support Cart". The footer of the page contains the text "Copyright 2001, 2006 Oracle Corporation. All Rights Reserved. About Oracle Applications Manager Version 2.3.1".

The Bug Fixes report lists all bug fixes included in the selected driver file. It contains the bug number, the associated product, and whether the bug fix was applied. If the fix was not applied, the Remarks column explains why.

5. View the Action Summary report.

You can create a report that summarizes the actions of a selected driver file. Click the Patch Details link at the top of the page to return to the Patch Details page. (You can also access the Action Summary report by clicking the bug fix number on the Bug Fixes report.)

From the Patch Details page, select a driver and click the Action Summary button.

Action Summary Report

The screenshot shows the Oracle Applications Manager interface. At the top, it says 'ORACLE Applications Manager' with navigation links like 'Support Cart', 'Setup', 'Home', 'Logout', and 'Help'. Below that, there's a breadcrumb trail: 'Applications Dashboard > Site Map > Applications System: ADOAGC04 > Applied Patches > Patch Details >'. The main title is 'Action Summary : u5251362.drv : ADOAGC04'. It shows the last updated time as '08-Jun-2006 21:13:58'. The start date is '26-May-2006 10:31:16' and the end date is '26-May-2006 10:34:19'. The platform is 'LINUX'. The patch top is '/ SLOTS/ slot03/ appmgr/ patches/ 5251362'. The driver file is 'u5251362.drv'. There's a filter section with 'Bug Fix' selected and a 'Go' button. A note says 'The details icon displays additional information about database updates.' Below that is a table with columns: Product, Directory, File, Action, Phase, Run, Bug Fix, and Details. The table shows 12 rows of patch actions.

Product	Directory	File	Action	Phase	Run	Bug Fix	Details
FND	include	afugai.h	copy		N	5251362	
AD	bin	adadmin	link		Y	5251362	
AD	bin	adident	link		Y	5251362	
AD	bin	adjkey	link		Y	5251362	
AD	bin	admrgpch	link		Y	5251362	
AD	bin	adncrv	link		Y	5251362	
AD	bin	adpatch	link		Y	5251362	
FND	include	wfmir.h	copy		N	5251362	
FND	include	wfntf.h	copy		N	5251362	
FND	include	afpcrm.h	copy		N	5251362	

The Action Summary report shows more information about the driver and its actions. For definitions of the column headings, see Action Summary, page 5-13.

If the driver selected contains a database portion, the Patch Summary report shows the driver actions, such as sql and exec. If the driver performed actions on the database, the Details icon is active. Click it to see the Action Details report.

Searching for Translation Patches:

Requirement: My Oracle E-Business Suite system operates in multiple languages. I want to make sure translation patches have been applied successfully.

If a patch has an associated translation patch, you apply the translation patch separately. AutoPatch stores information in the patch history database about all translation patches you apply.

To search for translation patches

1. Access the Oracle Applications Manager.

Follow the instructions in Accessing Patch Wizard, page 4-4 to access OAM.

2. Access the Simple Search page.

Site Map Page

ORACLE Applications Manager

Support Cart Setup Home Logout Help

Applications Dashboard | Site Map

Site Map: ADDAGC05

Administration Monitoring Maintenance Diagnostics And Repair

Patching and Utilities

- Applied Patches
- File History
- Patch Wizard
- Timing Reports
- Register Flagged Files
- Manage Downtime Schedules
- Codelevels Summary

Critical Activities

- Setup
- Monitor

TIP Only the items to which you have access are clickable.

Support Cart Setup Home Logout Help

Copyright 2001, 2006 Oracle Corporation. All Rights Reserved.
About Oracle Applications Manager Version 2.3.1

From the Site Map (Maintenance tab), click Applied Patches under the Patching and Utilities heading. The Simple Search page appears.

Enter the search criteria. For details about using the Simple Search page, see: Determining If a Patch Was Applied, page 3-54, or click the OAM Help button.

3. Specify the patch.

On the Simple Search page, enter the ID of the translation patch in the Patch field. Click Go.

4. Review the Patch Summary report

All applications of the patch are displayed. If multiple translations were applied, there will be multiple rows. The Language column shows the languages applied.

Applied Patches Page - Simple Search

ORACLE Applications Manager

Support Cart Setup Home Logout Help

Applications Dashboard | Site Map

Applications System: ADOAGC05 >

Applied Patches : ADOAGC05 Select Feature Patch Wizard Go

Last Updated : 2006/Oct/11 15:46:45

Simple Search

At least one field must be completed. Advanced Search

Querying by a specific patch ID will return all patches that are included in the specified patch ID.

Patch

(Enter either a 7-digit patch number or merge patch name)

Applied Within Last Days

Applied From Date To Date

(yyyy/MM/dd) (yyyy/MM/dd)

Language

Go Reset

A Bug Fix resolves a specific issue and a patch may contain one or more Bug Fixes.

Patch Name	Patch Description	Merged Patches	APPL_TOP Name	Language	Completion Date	Details
5336717.A	AD OAM DOC FOR RELEASE 12	No	ap6049rt	US	2006/Oct/06 13:45:39	
5553100.A	CONNECTION LEAK	No	ap6049rt	US	2006/Oct/04 14:11:36	
4461237.0		No	ap6049rt	KO	2006/Oct/03 07:14:18	
4461237.0	R12 ATG Family Pack A	No	ap6049rt	US	2006/Oct/02 16:36:18	

Viewing Applied Patches in a Report Format:

Requirement: Can I review applied patches information without the OAM screens?

There may be times when you want to view applied patch history without running the Oracle Applications Manager. For example, you may need to view large amounts of data, or you may just need a list of patches without the detail provided in the OAM Patch History reports. In these cases, you can run command line scripts that list all patches applied in each AutoPatch session, all files affected by a patch, or all patches applied within a certain date range. The scripts, and a description of the reports they produce, are listed in the following table.

Patch Report Scripts

Script Name	Report Content	Output Format
adphrept.sql	Lists patches applied in individual AutoPatch sessions, and includes details.	XML
adfhrept.sql	Displays information about files changed by patches.	XML

Script Name	Report Content	Output Format
adpchlst.sql	Lists patches applied in a given date range.	Text

The XML reports produced by adphrept.sql and adfhrept.sql can either be processed as XML or viewed as HTML.

To run a report that provides a listing of applied patches, follow the appropriate instructions in this section.

To see a list of all completed AutoPatch sessions with patch details

Run the adphrept.sql script (located in <AD_TOP>/patch/115/sql). This script produces an XML report showing individual AutoPatch sessions. If a patch was applied more than once, this report lists each application of the patch. If a merged patch was applied, it lists the merged patch by patch name. It does not list the individual patches within the merged patch.

To run adphrept.sql, use the following parameters:

```
<query_depth> <bug_number or ALL> <bug_product or ALL> \
<start_date_from (mm/dd/rr or ALL)> <end_date_to (mm/dd/yyyy or ALL)> \
<patchname/ALL> <patchtype/ALL> <level/ALL> <language/ALL> \
<apptop/ALL> <limit to forms server? (Y/N)> \
<limit to web server? (Y/N)> \
<limit to node server? (Y/N)> \
<limit to admin server? (Y/N)> \
<only patches that change DB? (Y/N)> <report_name>.xml
```

For <query_depth>, specify 1 (details of patches only), 2 (details of patches and their bug fixes only), or 3 (details of patches, bug fixes, and bug actions).

At the command prompt, enter the report command and enter values for the parameters and prompts. For example, to see complete patch details for AutoPatch sessions that were run during January 2009, enter the following, using the mm/dd/yyyy date format:

UNIX:

```
$ cd $AD_TOP/patch/115/sql
$ sqlplus <APPS username>/<APPS password> \
@adphrept.sql 3 ALL ALL 01/01/2009 01/31/2009 \
ALL ALL ALL ALL ALL N N N N N jan09.xml
```

Windows:

```
C:>\ cd %AD_TOP%\patch\115\sql
C:>\ sqlplus <APPS username>/<APPS password> @adphrept.sql 3 ALL ALL
01/01/2009 01/31/2009 ALL ALL ALL ALL ALL N N N N N jan09.xml
```

The <AD_TOP>/html directory contains the adpchrep.xml style sheet for displaying the XML output file in HTML format. To view the XML file as HTML, copy both the adpchrep.xml style sheet and XML output report to a directory accessible by a browser.

Open the directory in your browser and click the XML filename.

To display information about files changed by patches

Run the `adfhrept.sql` script (located in `<AD_TOP>/patch/115/sql`) to produce an XML report named `adfilerep.xml`. Use the following parameters:

```
<filename> <latest file version only? (Y/N) \
<start_date (mm/dd/rr or ALL)> <end_date (mm/dd/yyyy or ALL)> \
<patchtype/ALL> <language/ALL> \
<appltop/ALL> <limit to forms server? (Y/N)> \
<limit to web server? (Y/N)> \
<only patches that change DB? (Y/N)>
```

At the command prompt, enter the report command and enter values for the parameters and prompts. For example, to see the complete file version history for `admorgb.pls` considering only patches applied in January 2008, enter the following, using `mm/dd/yyyy` format:

UNIX:

```
$ cd $AD_TOP/patch/115/sql
$ sqlplus <APPS username>/<APPS password> \
@adfhrept.sql admorgb.pls N 01/01/2008 01/31/2008 \
ALL ALL ALL N N N N N
```

Windows:

```
C:>\ cd %AD_TOP%\patch\115\sql
C:>\ sqlplus <APPS username>/<APPS password> @adfhrept.sql admorgb.pls
N 01/01/2008 01/31/2008 ALL ALL ALL N N N N N
```

The `<AD_TOP>/html` directory contains the `adfilerep.xsl` style sheet for displaying the XML output file in HTML format. To view the XML file as HTML, copy both the `adfilerep.xsl` style sheet and XML output report to a directory accessible by a browser. Open the directory in your browser and click on the XML filename.

To see a list of all patches in a given date range

The `adpchlst.sql` report (located in `<AD_TOP>/patch/115/sql`) produces a list (`adpchlst.lst`) of all patches in a date range, without patch detail. It differs from `adphrept.sql` in two ways: it lists a patch only once regardless of how many times it was applied, and it lists individual patches included within a merged patch. For example, if you combine patches 123, 124, and 125 in a merged patch called `merged1`, the report lists patches 123, 124, and 125, but not `merged1`.

At the command prompt, enter the report command and enter the date parameters in `mm/dd/yyyy` format. For example, to see a list of patches applied in October 2008, enter the following:

UNIX:

```
$ cd $AD_TOP/patch/115/sql
$ sqlplus <APPS username>/<APPS password> \
@adpchlst.sql 10/01/2008 10/31/2008
```

Windows:


```
C:>\ cd %AD_TOP%\patch\115\sql
C:>\ sqlplus <APPS username>/<APPS password> @adpchlst.sql 10/01/2008
10/31/2008
```

Monitoring Patches in Progress:

Requirement: Can I monitor the progress of a patch while it is being applied?

Depending on the size and complexity of a patch, it may take from several minutes to several hours to completely apply it to your system. It is useful to know what a patch is currently doing and how long individual steps are taking.

In order to access the Timing Reports to track an in-progress patching session, the Web server must be started in restricted mode and OAM accessed through a restricted mode URL.

When using Timing Reports to track an in-progress patching session, the timing report displays the most recently performed tasks. Use the Refresh icon to get the latest running tasks.

To monitor patches in progress

1. Set up the ad_monitor user account. Use the ad_monitor user account to log in to OAM in restricted mode.
 - Log in to SQL*Plus as SYSTEM.
 - Unlock the ad_monitor user.

```
SQL> alter user ad_monitor account unlock;
```
 - Log in to SQL*Plus as the ad_monitor user and reset the password. The default password is 'lizard'.
2. Run AutoPatch (adop) to start the patch session.
3. Access OAM through the restricted mode URL:

```
<host>:<port>/servlets/weboamLocal/oam/oamLogin
```
4. Log in to OAM as the ad_monitor user.
5. Navigate to the Timing Reports (Navigation: Sitemap > Maintenance > Patching and Utilities > Timing Reports).

See: Timing Reports, page 5-15.
6. When the patching session is complete, shut down the restricted mode Web server.

UNIX:

```
$ adaprstctl.sh stop
```

Windows:

```
C:\> adaprstctl.cmd stop
```

You can also monitor the progress of the patching process by reviewing:

- AutoPatch messages

As AutoPatch runs, it displays messages on the screen about the status and progress of the patching process.

- Patch log files

AutoPatch creates log files in the current directory. These files contain information about patching actions that have been performed.

- Worker status

For jobs run in parallel, use AD Controller to view the status of the concurrent manager and workers assigned to process jobs. See AD Controller examples in *Managing Worker Processes*, page 9-1.

Analyzing Patches Without an Internet Connection:

Requirement: How do I analyze specific patches if I do not have access to an Internet connection?

You can run Patch Wizard to analyze specific patches without access to an Internet connection, if necessary, by downloading the patches to a system which has Internet access. Once the download is complete, copy the patches to the Patch Wizard's staging directory. Then run Patch Wizard as you normally would based on the files you copied to the staging directory.

To analyze specific patches without using an Internet connection

1. Download the patch zip file(s) to a system which has Internet access.
2. Set up a staging directory on a system that does not have Internet access. Patch Wizard must be able to read from and write to this staging directory.
3. Copy the patch zip file(s) to the <staging directory>/ad directory, if the downloaded patch is an AD product patch. Otherwise, copy the patch zip file(s) to <staging directory>/nonad directory. The zip file(s) must be copied to a system that can access the Patch Wizard staging directory. If the staging directory is on a local disk, the zip file(s) must be copied to the system where you run Patch Wizard. If the staging directory is on a shared (network) disk, it can be copied to any system with access to the shared disk.
4. Run Analyze Specific Patches as you normally would from this point.

Patch Tracking Utilities

Patch Wizard

With Patch Wizard, you can determine patches that have not been applied to your system. It does not report on all available patches. It compares the patches you have already applied against a list of all recommended Oracle E-Business Suite patches. Recommended patches can include high-priority patches or patches that update to a new codelevel, such as release update packs (RUPs), product family RUPs, and pre-upgrade patches.

How Patch Wizard Works

Patch Wizard supplies you with an interface from which you can:

- Set preferences, both site-specific and general, that include the staging directory and various defaults that will apply to the patches you download.
- Set up filters that report only those patches that may affect your system.
- Submit a request for a report of recommended patches, based on the filter(s) you set up.
- Analyze the impact of specific patches from a list you supply to Patch Wizard.
- Download patches, specify the language of the patches to download, and merge patches.

Before running any of the Patch Wizard tasks, set up your My Oracle Support (Metalink) credentials on the OAM Update Metalink Credentials page. To access this page, go to the Patch Wizard main page, click Setup in the title bar. Alternatively, click the Setup link at the top right of any page of the OAM interface.

The Dashboard Setup page appears. Click the Metalink Credentials link on the left side. The Update Metalink Credentials page appears.

OAM Update Metalink Credentials Page

Update your My Oracle Support credentials by providing your user ID, password, email address, proxy server host name, proxy server port, proxy bypass domains, proxy user name, and proxy password. The recommend, analyze, and download patches features typically require that your My Oracle Support user ID and password are set on the OAM Update Metalink Credentials page. However, the recommend and analyze features can also be used when Patch Wizard does not have access to a direct Internet connection. If this is the case, you can leave the Metalink user ID and password empty.

The Patch Information Bundle

The *Patch Information Bundle* file contains the zip files of recommended patches, the list of recommended patches (Recommended.xml), the latest codelevel patches (Codelevels.xml), and information on products and product families (ProductInfo.xml). Each patch zip file contains a readme file, a patch LDT file, and a patch metadata file (patch_metadata.xml).

The Patch Information Bundle file is updated daily. When you submit a patch analysis request, this file is automatically downloaded (if it is not specified otherwise in your Metalink credentials.)

Patch Wizard loads the Patch Information Bundle data, including LDT files and readme files, into the Oracle E-Business Suite database. It uses the metadata to provide patch recommendations.

Concurrent Programs

When you submit a request for patch analysis, Patch Wizard performs the following tasks using a set of concurrent programs:

- Uploading patch information from the Patch Information Bundle to Patch Wizard tables
Patch Wizard loads the Patch Information Bundle metadata, including LDT files and readme files, into the Oracle E-Business Suite database.
- Recommending patches based on the current environment and the Patch Information Bundle
Patch Wizard reports which patches update Oracle E-Business Suite at the current codelevel and which update to a new codelevel.
- Downloading patches (ad hoc or based on the list of recommended patches)
Patch Wizard can download patches from My Oracle Support, and then merge the patches in the Patch Wizard staging directory.
- Analyzing lists of patches after downloading them from My Oracle Support
Patch Wizard uploads the metadata for a specific patch or set of patches for you to view information reported from the metadata. For example, you can submit a request for patch analysis, and then view any recommended patches that have not yet been applied and the impact of applying this new patch.

Running Patch Wizard Without Access to an Internet Connection

You can run Patch Wizard without access to an Internet connection, if necessary, by downloading the Patch Information Bundle to a system which has Internet access. Once the download is complete, copy the Patch Information Bundle file to the Patch Wizard's staging directory. Then run Patch Wizard as you normally would, to recommend and analyze patches, based on the files you copied to the staging directory.

The Patch Wizard Interface

Patch Wizard is a Web-based utility in Oracle Applications Manager (OAM). The OAM interface gives Patch Wizard pages a uniform look and feel.

Main Page

From this page, you have access to task icons used to set up the Patch Wizard staging directory, manage patch filters, submit concurrent requests, and view recommended patches. In addition, the Recommended Results section of this page displays a list of patches based on submitted requests.

Task Icons

From the main page, access the other Patch Wizard pages by clicking on Task icons. The icons provide links to the following pages: Patch Wizard Preferences, Define Patch Filters, Recommended/Analyze Patches, Download Patches, and Aggregate Patch

Impact.

Details Icons

On many Patch Wizard pages, you can drill down to see more detail. For example, from the Recommended Patches Results section of the main page, click the Details icon for a specific recommended patch request to view the recommended patch results

Accessing Patch Wizard

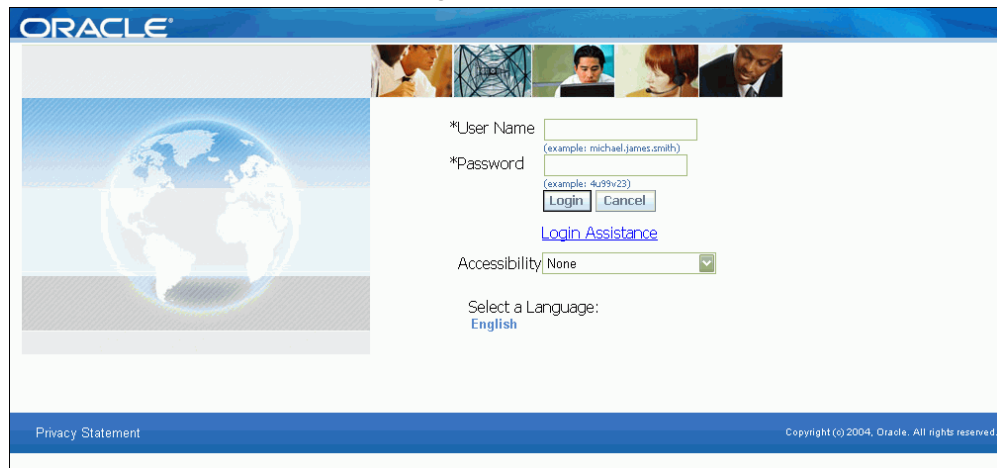
To access Patch Wizard, log in to Oracle Applications Manager (OAM) and choose Patch Wizard from the Navigator pane.

1. Enter the following URL in your browser:

`http://<HTTP hostname>.<domain>:<HTTP port>/OA_HTML/AppsLogin`

The Welcome page appears.

Oracle E-Business Suite Welcome Page



The screenshot shows the Oracle E-Business Suite Welcome Page. At the top, there is a blue header with the Oracle logo. Below the header, there is a navigation pane on the right with several icons. The main content area features a large globe graphic on the left and a login form on the right. The login form includes fields for *User Name (with example: michael.james.smith) and *Password (with example: 4099v23), a Login button, a Cancel button, a Login Assistance link, an Accessibility dropdown menu set to None, and a language selection option for English. The footer contains a Privacy Statement link and a copyright notice for Oracle (c) 2004.

Enter your user name and password, and click Login. The Oracle E-Business Suite Home page appears. Click the System Administration link in the Navigator pane. Another Navigator pane appears to the right.

2. Go to the Patch Wizard main page.

Scroll down to the Oracle Applications Manager section of the right-hand Navigator pane, and click Patch Wizard to go to the main page.

Patch Wizard Main Page

Use the main page to access all features of Patch Wizard and to view the results of your requests for recommended patches. The Select Feature drop-down list at the top of the page provides access to the Applied Patches, File History, Timing Reports, and Register

Flagged Files features.

Patch Wizard Main Page

ORACLE Applications Manager

Support Cart Setup Home Logout Help

Applications Dashboard | Site Map

Applications System: ADOAGC05 >

Patch Wizard : ADOAGC05

Select Feature Applied Patches Go

Last Updated : 2006/Oct/11 17:29:08

Patch Wizard Tasks

Task Name	Description	Tasks	Job Status
Patch Wizard Preferences	Set download, merge, and stage area preferences		
Define Patch Filters	Create custom patch filters		
Recommend/Analyze Patches	Submit requests for patch advice or analysis		
Download Patches	Submit requests to download patches		
Aggregate Patch Impact	Aggregate Patch Impact		

Filter Criteria

Filter Name contains

Completion Date is (yyyy/MMM/dd)

Go

Recommended Patches Results

Previously submitted Filter Names/ Patch Lists that do not appear in the Recommended Patches Results section have been purged according to the frequency setting in Purge Concurrent Request. Change the frequency setting in Applications Dashboard > Critical Activities if needed.

Filter Name/ Patch List	Total (Applied & Unapplied)	Unapplied	Requested By	Completion Date	Run Status	Request Set	Details
Recommended Patches and New Codelevels	1	1	SYSADMIN	2006/Oct/11	Normal	280040	

Patch Wizard Tasks Table

The Patch Wizard Tasks table lists the tasks available in Patch Wizard. The table contains the following columns of information for each task:

- Task Name: Name of the Patch Wizard task.
- Description: Describes the task.
- Tasks: Link to the page associated with the Patch Wizard task.
- Job Status: Link to the request submission status of the task where you can review warnings or errors for your request.

From the Patch Wizard Tasks table, choose the Tasks icons to view:






- The Patch Wizard Preferences page
- From the Preferences page, set the staging directory, merge patch defaults, the languages and platform defaults for downloading patches, and whether to display or hide hidden patches. You must define the values on this page before you use any other Patch Wizard feature. You can also use this page to modify existing preferences at a later date.

- The Define Patch Filters page
Typically, you see only those patches that are recommended for your system. Use this page to set up filters that report only those patches that may affect your system.
- The Recommend/Analyze Patches page
Once you have selected values for a filter, submit the request to create a report of recommended patches based on that filter. You can also analyze specific patches by entering a comma-separated list of patch numbers.
- The Download Patches page
You can download patches, specify the language of patches to download, and merge patches from this page.

Recommended Patches Results

The Recommended Patches Results section of the Patch Wizard main page lists all in-progress and completed requests for patch recommendations, based on the information you entered in the Filter Criteria section of the main page. For example, you can view only the results that contain a certain text string in the filter name, or only the results of requests completed on a certain date.

Patch Wizard Main Page - Recommended Patches Results

Recommended Patches Results							
<small>ⓘ Previously submitted Filter Names/Patch Lists that do not appear in the Recommended Patches Results section have been purged according to the frequency setting in Purge Concurrent Request. Change the frequency setting in Applications Dashboard > Critical Activities if needed.</small>							
Filter Name/Patch List	Total (Applied & Unapplied)	Unapplied	Requested By	Completion Date	Run Status	Request Set	Details
Recommended Patches and New Codelevels	1	1	SYSADMIN	2006/Oct/11 02:51:41	Normal	280040	
Recommended Patches	0	0	SYSADMIN	2006/Oct/11 02:32:16	Normal	280014	
4502603	1	1	SYSADMIN	2006/Oct/09 00:09:09	Warning	279849	
4502400	1	1	SYSADMIN	2006/Oct/08 23:52:48	Warning	279809	
Recommended Patches and New Codelevels	1	1	SYSADMIN	2006/Oct/08	Normal	279792	

Clicking the icon in the Details column of a specific recommended patch request accesses the Recommended Patches Results page. After setting up and submitting a request, view the details of the recommended patches on this page.

Patch Wizard Preferences

The site-specific information you set on the Patch Wizard Preferences page applies to other functions of Patch Wizard, such as Recommend/Analyze, Download Patches, and Codelevels Summary. From the Patch Wizard main page, click the Tasks icon for Patch Wizard Preferences.

The top portion of the page contains these sections: Staging Directory, Merge Option Defaults, and Language and Platform Details.

Patch Wizard Preferences Page - Top

The screenshot shows the top portion of the Oracle Applications Manager Patch Wizard Preferences page. The header includes the Oracle logo and 'Applications Manager' text. Navigation links for 'Support Cart', 'Setup', 'Home', 'Logout', and 'Help' are visible. The page title is 'Patch Wizard Preferences : r121test'. Below this, it shows the last update time and the user's Oracle MetaLink ID. The 'Staging Directory' section contains a text input field with the value '/slot/ems1227/appmgr/stage'. The 'Merge Option Defaults' section has a checkbox for 'Automatically merge downloaded patches' and radio buttons for 'One merged patch: US and non-US', 'Two merged patches: US; non-US', and 'Multiple merged patches: US; language1; language2;...'. The 'Language and Platform Defaults' section features a list of available languages (Albanian, Arabic, Brazilian Portuguese, Canadian French, Croatian, Czech, Danish, Dutch, Finnish, French) and a 'Selected Languages' list. Below this, there are radio buttons for platform options: Linux x86, Linux x86-64, Sun Solaris OS (SPARC 64-bit), Microsoft Windows (32-bit), AIX/SL Based Systems (64-bit), HP-UX PA-RISC (64-bit), and HP-UX Itanium.

The bottom portion of the page contains the In Use Products Defaults and the Display Option Defaults sections.

Patch Wizard Preferences Page - Bottom

The screenshot shows the bottom portion of the Oracle Applications Manager Patch Wizard Preferences page. The 'In Use Products Defaults' section has a checkbox for 'Patch recommendations for In Use products only'. The 'Display Option Defaults' section has a checkbox for 'Show Hidden Patches (with the check mark in the Hide Patch column)'. At the bottom, there are 'OK' and 'Cancel' buttons, and an 'Add to Support Cart' button. The footer contains 'Support Cart Setup Home Logout Help' and copyright information for Oracle Corporation.

Staging Directory

Depending on the product type (AD or non-AD), Patch Wizard downloads patches to a subdirectory under the staging directory. If the patch that ADOAM is trying to download is an AD product patch, it will be downloaded to the "ad" subdirectory under the stage directory. Non-AD Patches will be downloaded to "nonad" subdirectory.

The staging directory is also used by Patch Wizard to create temporary files and subdirectories for patch recommendation requests. These temporary files and directories are deleted after processing.

Note: Oracle recommends you use the same staging directory each time you run Patch Wizard.

Merge Option Defaults

Merging patches reduces patch application time by eliminating redundant tasks (such as responding to prompts and linking executables) for each individual patch.

Note: AD product patches cannot be merged with other product patches.

Patch Wizard screens allow you to merge AD and non-AD patches by specifying the appropriate options as discussed below. You can choose to automatically merge patches that you download.

The following merge options can be set:

- **Metalink Credentials:** This link accesses the OAM Update Metalink Credentials page, which allows you to set your My Oracle Support user ID and password. The recommend, analyze, and download patches features typically require that your My Oracle Support user ID and password are set in the OAM Update Metalink Credentials page. The recommend and analyze features can also be used when Patch Wizard does not have access to a direct Internet connection. If this is the case, leave the Metalink user ID and password empty.

The patches will be downloaded to the relevant subdirectory for the patch's product. If the patch that ADOAM is trying to download is an AD product's patch, it will be downloaded to "ad" subdirectory under the stage directory. Otherwise, it will be downloaded to the "nonad" subdirectory.

- **Automatically merge downloaded patches:** This check box tells Patch Wizard to automatically merge all downloaded patches.
- **Merging Strategy:** The following merging strategies when are available when downloading translation patches:
 - **One merged patch:** If the list of patches contains a single AD product patch or a single non-AD product patch, a single merged patch containing US and non-US patches will be generated. If the list of patches contains both AD and non-AD product patches, two merged patches will be generated: one for the merge of all the AD product patches, and the other for the merge of all non-AD product patches.

For example, if "mrg_2009072781642" is the merged patch name provided in the Downloads page, the merged AD patch directory name will be mrg_2009072781642_A, and the merged non-AD patch directory name will be mrg_2009072781642_N.

- **Two merged patches:** If the list of US patches includes both AD and non-AD product patches, the patches will be merged separately for AD and non-AD product patches. One merged patch will contain all US patches, and a second merged patch will contain all non-US patches.

For example, if "mrg_2009072781642" is the merged patch name provided in the Downloads page, the merged AD US patch directory name will be mrg_2009072781642_US_A, and the merged non-AD US patch directory name will be mrg_2009072781642_US_N.

- **Multiple merged patches:** If the list of US patches includes a combination of AD products, non-AD products, AD language patches, and non-AD language patches, the patches will be merged separately for each of these categories.

For example, if "mrg_2009072781642" is the merged patch name provided in the Downloads page, the merged AD US patch directory name will be mrg_2009072781642_US_A, the merged non-AD US patch directory name will be mrg_2009072781642_US_N, the merged AD NLS patch directory name will be mrg_2009072781642_<lang_abbr>_A, and the merged non-AD NLS patch directory name will be mrg_2009072781642_<lang_abbr>_N.

Language and Platform Details

You can select the languages (one or more) of patches that Patch Wizard will recommend and download. You can also select the platform of the patches you want recommended and downloaded.

In Use Products Defaults

Selecting the Patch Recommendation for In Use Products Only check box directs Patch Wizard to display on the Recommended Patch Results page only patches for the products marked as in use (active) in your system.

Both Patch Wizard and AutoPatch look at the In Use flag to determine what products you are using. To review the products currently marked as being in use in your system, click the Codelevels Summary Page link. See: Codelevels Introduced by the Patch, page 4-17.

Display Option Defaults

Hidden patches are patches that you choose not to see in your reports. For example, if Patch Wizard recommends patches for products you do not need for your system, you can choose to hide these patches.

However, checking the Show Hidden Patches box in the Patch Wizard Preferences page overrides the hidden patch setting, and all patches, even hidden patches, are reported.

After you have made all your selections on the Patch Wizard Preferences page, click OK to save them or click Cancel to discard.

Define Patch Filters

The Patch Information Bundle file contains information for all recommended patches for all products. If Patch Wizard were to compare patches in the patch information database against all metadata in the Patch Information Bundle file, the number of recommended patches in the report might be too large to be useful for an individual system. Patch Wizards provides filters so that only those patch types and products in the metadata that apply to your system are included in the comparison.

From the main page, click the Tasks icon for Define Patch Filters to see all filters created for the current system. Patch Wizard provides three pre-seeded filters. In addition, you can create your own custom filters.

Define Patch Filters Page

The screenshot shows the Oracle Applications Manager interface for the 'Define Patch Filters' page. The page title is 'Define Patch Filters : ADOAGC05'. Below the title, there is a message: 'The Oracle Patch Filters (Recommended Patches, New Codelevels, Recommended Patches and New Codelevels) cannot be edited.' A 'Create New' button is located to the right of this message. Below the message is a table with the following data:

Select Patch Filter Name	Type	Description	Updated By	Updated Date
<input checked="" type="radio"/> New Codelevels	Oracle	New Codelevels	INITIAL SETUP	2006/Sep/26 05:30:01
<input type="radio"/> Recommended Patches and New Codelevels	Oracle	Current Recommended Patches and New Codelevels	INITIAL SETUP	2006/Sep/26 05:30:01
<input type="radio"/> Recommended Patches	Oracle	Recommended Patches for Current Codelevel	INITIAL SETUP	2006/Sep/26 05:30:01
<input type="radio"/> create new testing	Custom	This is test filter	SYSADMIN	2006/Sep/26 15:44:52
<input type="radio"/> ATG	Custom	ApplicationsTech	SYSADMIN	2006/Sep/25 04:15:17

Below the table, there is an 'Add to Support Cart' button. The page footer contains the text: 'Copyright 2001, 2006 Oracle Corporation. All Rights Reserved. About Oracle Applications Manager Version 2.3.1'.

Note that the pre-seeded filters are marked "Oracle" in the Type column, and filters you create are marked "Custom." You cannot edit or delete the pre-seeded filters, but you can use any of them as a template to create a new filter.

The pre-seeded filters are:

- **New Codelevels:** Determines recommended patches for release update packs (RUPs), product family RUPs, and pre-upgrade patches. These patches update versions of a product, family, or the entire Oracle E-Business Suite system.
- **Recommended Patches:** Determines recommended patches for the current codelevel.

- Recommended Patches and New Codelevels: Determines recommended patches for both the current and new codelevels.

Creating a New Custom Patch Filter

From the Define Patch Filters page, click Create New to create a new custom filter.

Create Patch Filters Page

The screenshot shows the Oracle Applications Manager interface for creating a patch filter. The page title is "Create Patch Filters : ADOAGC03". It includes a breadcrumb trail: "Applications System:ADOAGC03 > Patch Wizard > Define Patch Filters >". The main section is titled "Patch Filter" and contains a form with two required fields: "Name" and "Description". Below the form is a table titled "Licensed Product Families" with three columns: "Product Family", "Recommended Patches", and "New Codelevel". The table lists various product families with checkboxes for selection.

Product Family	Recommended Patches	New Codelevel
Advanced Planning	<input type="checkbox"/>	<input type="checkbox"/>
Applications Technology	<input type="checkbox"/>	<input type="checkbox"/>
Communications	<input type="checkbox"/>	<input type="checkbox"/>
Contracts Suite	<input type="checkbox"/>	<input type="checkbox"/>
Discrete Manufacturing	<input type="checkbox"/>	<input type="checkbox"/>
E-Business Intelligence	<input type="checkbox"/>	<input type="checkbox"/>
Exchange Suite	<input type="checkbox"/>	<input type="checkbox"/>
Financial Globalizations Suite	<input type="checkbox"/>	<input type="checkbox"/>
Financial Payables Suite	<input type="checkbox"/>	<input type="checkbox"/>
Financial Receivables Suite	<input type="checkbox"/>	<input type="checkbox"/>
Financial Services Applications	<input type="checkbox"/>	<input type="checkbox"/>
Financials	<input type="checkbox"/>	<input type="checkbox"/>

On the Create Patch Filters page, enter a unique name and a description for each new custom filter. All licensed product families are listed in the Licensed Product Families section. Non-licensed product families are listed at the bottom of the page. There are two columns for each product family: Recommended Patches and New Codelevel. Select a box for each product family you want to include in the new filter.

By choosing the Create Like button on the Define Patch Filters page, you indicate you want to use an existing filter as a template to create a new filter. The system supplies the filter criteria from the existing filter for the new filter. Edit the criteria by checking or clearing the boxes. Click Continue to create the new filter.

Recommend/Analyze Patches

After setting up the Patch Wizard staging area (and optionally creating custom filters) on the main page, you can submit requests for processing. Click the Tasks icon for Recommend/Analyze Patches.

Recommend Patches Page - Top

The Options section of the Recommend Patches page contains the following actions:

- **Create Recommendation**

This generates recommendations based on the patch filter you selected. Choose one of the pre-seeded filters or any custom filter you created on the Define Patch Filters page. Patch Wizard uses the filter and compares the patch information database against the patch metadata to recommend which patches you should apply.
- **Analyze Specific Patches**

This generates recommendations for specific patches. After downloading specific patches from My Oracle Support and placing them in the staging area, you can analyze these patches to determine their affect on your system by entering the patch numbers in this section. (Enter either bug numbers (for example, 1234567) or full patch names (for example, 1234567_R12.AD.A).) Check the "Analyze Aggregate Patch Impact" box to analyze Aggregate Patch Impact.

You can enter a date and time in the Schedule section of this page to run the request at a later time. The default setting is to run the job immediately. You can also schedule automatic recurring requests by entering the information in the Recurrence section.

Recommend Patches Page - Bottom

Schedule

ⓘ If no date or an earlier date is specified, the request will be scheduled to run immediately.

Date: 18-Mar-2009 (dd-MMM-yyyy)

Time: 00:00

Recurrence

Never Repeat

Repeat

Every: [] Days

End Date: [] (dd-MMM-yyyy)

End Time: 00:00

OK Cancel

Add to Support Cart

Click OK to submit the request.

When the request is submitted, Patch Wizard looks in the specific directories under the stage directory for the patches. If the patches are not present, or if there are newer versions available, Patch Wizard downloads them from My Oracle Support before creating recommendations or analyzing patches.

Note: If you want to create recommendations or analyze specific patches without using an Internet connection, refer to *Creating Patch Recommendations Without an Internet Connection*, page 3-53 or *Analyzing Patches Without an Internet Connection*, page 3-64.

Each time you submit a request to analyze specific patches or to recommend patches, Patch Wizard creates a Request Set ID. The Request Set ID is shown in the Results section of the Patch Wizard main page. To check the status of your request, click on the Request Set ID corresponding to your request in the Recommended Patches Results section.

Download Patches

You use the Download Patches page to request a download of specific patches from My Oracle Support. From the main page, click the Download Patches tasks icon. From the Download Patches page, enter the patch numbers in the input field, separated with commas. Enter either bug numbers (for example, 1234567) or full patch names (for example, 1234567_R12.AD.A). You can also choose to analyze the patches while downloading, or analyze and compute aggregate patch impact while downloading.

Download Patches Page - Top

Support Cart Setup Home Logout Help

Applications Dashboard Site Map
Applications System:r121test > Patch Wizard >

Download Patches : r121test

Last Updated : 18-Mar-2009 14:48:20
Staging Directory : /slot/ems1227/appmgr/stage
Oracle MetaLink User ID : murali.kumanduri@oracle.com

OK Cancel

Patch Selection

i You must set up your Metalink credentials before downloading patches.

Patch List

(Enter Patch numbers, separated by commas ex: 1234567, 8765432_R12.AD.A, 8888888.AD.B. Maximum number of allowed characters is 175)

Options

Download only
(Download the exact list of patches above)

Download and Analyze
(Download only patches listed above that have not been applied and their prerequisite patches)

Download, Analyze and Aggregate Patch Impact
(Download only patches listed above that have not been applied and their prerequisite patches, analyze them, and compute aggregate patch impact)

Merge Options

Automatically merge downloaded patches
(Merge happens only if all patches are downloaded successfully)

Merged Patch Name

Merging Strategy

One merged patch: US and non-US

Two merged patches: US; non-US

Multiple merged patches: US; language1;language2;...

The Merge Options section is where you define how patches should be merged after downloading. The defaults for merging are set on the Patch Wizard Preferences page. If you choose to automatically merge patches while downloading, you can modify the merged patch name and specify the merging strategy in this section. The default merged patch name is "merged_YYYYMMDDhhmmss", where "hh" is in 24-hour format.

The Languages and Platform section allows you to select the languages and platform of the downloaded patches. The defaults for languages and platform are set on the Patch Wizard Preferences page. However, you can modify the information on this page. When you provide information in this section, Patch Wizard downloads only patches that match the languages and platform you select.

Download Patches Page - Bottom

The screenshot displays a web interface for configuring patch downloads. It is divided into two main sections: 'Languages and Platform' and 'Schedule'.

Languages and Platform:

- Select Languages and Platform for downloading patches.**
- Available Languages:** A list box containing Albanian, Arabic, Brazilian Portuguese, Canadian French, Croatian, Czech, Danish, Dutch, Finnish, and French.
- Selected Languages:** An empty list box.
- Buttons:** Move, Move All, Remove, and Remove All.
- Platform:** Radio buttons for Linux x86, Linux x86-64, Sun Solaris OS (SPARC 64-bit), Microsoft Windows (32-bit), AIX/SL Based Systems (64-bit), HP-UX PA-RISC (64-bit), and HP-UX Itanium.

Schedule:

- If no date or an earlier date is specified, the request will be scheduled to run immediately.**
- Date:** A date picker set to 18-Mar-2009.
- Time:** Two dropdown menus set to 00:00.

At the bottom right, there are 'OK' and 'Cancel' buttons, and an 'Add to Support Cart' button. The footer contains 'Support Cart Setup Home Logout Help' and 'Copyright 2001, 2009 Oracle Corporation. All Rights Reserved. About Oracle Applications Manager Version 2.3.1'.

You can also provide information in the Schedule section about downloads you want to perform at a later date.

Recommended Patches Results

From the main page, click the Details icon associated with a patch request in the Results section to access the Recommended Patches Results page. This page presents a set of recommended patches based on the results of the selected Recommend Patches request that you submitted.

The first section lists the recommended patches.

Recommended Patches Results Page

ORACLE Applications Manager Support Cart Setup Home Logout Help

Applications Dashboard | Site Map
 Applications System:ADDAGC03 > Patch Wizard >

Recommended Patches Results : ADDAGC03
 Last Updated : 10-Aug-2007 20:49:55

Patch Filter/Patch ID **Recommended Patches**
 Requested By **SYSADMIN**
 Completion Date **10-Aug-2007 20:38:43**
 View Aggregate Patch Impact **Aggregate Impact**

Show Hidden Patches (with the check mark in the Hide Patch column)

! If the Show Hidden Patches checkbox is not selected, the number of patches displayed may be less than the number listed on the Patch Wizard page.
 Only patches selected on the current page can be downloaded.

Recommended Patches Results

Select Patch and ... Previous 1-25 of 27 Next 2

Select All | Select None

Select Patch	Product	Prerequisites	Codelevel Introduced	Status	PAA Reason Recommended	Patch Description	Hide Patch	Included in Aggregate Patch Impact	Impact
<input type="checkbox"/> 4502962.R12.ad		0	No	Applied	Yes High Priority Patch	Minipack 4502962	<input type="checkbox"/> No	No	
<input type="checkbox"/> 5582560.A.xla	xla	0	No	Unapplied	Yes High Priority Patch	R12: LEDGER UPGRADE NOT	<input type="checkbox"/> No	No	

The second section lists patches that introduce new codelevels.

Recommended Patches Results - Patches that Introduce New Codelevels

Patches that Introduce New Codelevels

Select Patch and ...

Select All | Select None

Select Patch	Product	Prerequisites	Codelevel Introduced	Status	PAA Reason Recommended	Patch Description	Hide Patch	Included in Aggregate Patch Impact	Impact
<input type="checkbox"/> 6435000.R12.au		0	Yes	Unapplied	Yes Maintpack	ORACLE E-BUSINESS SUITE 12.0.4 RELEASE UPDATE PACK (RUP4)	<input type="checkbox"/> No	No	
<input type="checkbox"/> 6493602.A	fin_pf	0	No	Unapplied	Yes Familypack	Oracle Financials Release Update Pack 4 for 12.0 (R12.FIN_FF.A.DELTA.4)	<input type="checkbox"/> No	No	
<input type="checkbox"/> 6494646.A	hr_pf	0	No	Unapplied	Yes Familypack	R12.HR_FF.A.DELTA.4	<input type="checkbox"/> No	No	
<input type="checkbox"/> 6497749.A	prc_pf	0	No	Unapplied	Yes Familypack	Oracle Procurement Release Update Pack 4 for 12.0 (R12.PRC_FF.A.DELTA.4)	<input type="checkbox"/> No	No	
<input type="checkbox"/> 6507355.A	cc_pf	0	No	Unapplied	Yes Familypack	Oracle CRM Release Update Pack 4 for 12.0 (R12.CC_FF.A.delta.4)	<input type="checkbox"/> No	No	
<input type="checkbox"/> 6508131.A	scp_pf	0	No	Unapplied	Yes Familypack	Oracle Supply Chain Planning Release Update Pack 4 for 12.0 (R12.SCP_FF.A.DELTA.4)	<input type="checkbox"/> No	No	
<input type="checkbox"/> 6508212.A	scm_pf	0	No	Unapplied	Yes Familypack	Oracle SCM Release Update Pack 4 for 12.0 (R12.SCM_FF.A.DELTA.4)	<input type="checkbox"/> No	No	

Support Cart Setup Home Logout Help

Copyright 2001, 2006 Oracle Corporation. All Rights Reserved.
 About Oracle Applications Manager Version 2.3.1

Both sections display the following columns:

- **Select:** Select this check box and click the Download button if you want the patch number sent to the Download Patches page for submission. You can select any number of patches.
- **Patch:** The patch number of the recommended patch.
- **Product:** The product to which patch applies.

- Prerequisites: Specifies the codelevel required before you can apply this patch.
- Codelevel Introduced: Indicates whether the patch introduces a new codelevel for that product. If a new codelevel is introduced, the Yes indicator contains a link the Codelevels Introduced by the Patch page.
- Status: Indicates whether the patch is applied, unapplied, missing, or obsolete.
- PAA: This indicates whether there are manual steps you have to perform if you apply the patch.
- Reason Recommended: The reason the patch is recommended, for example, it is a high-priority patch, or part of a maintenance release pack (RUP) or a product family pack (RUP).
- Patch Description: Describes the patch.
- Hide Patch: Select this check box to hide the patch from the list of recommended patches. Use this feature to hide patches that you do not want to apply to your system. To hide or show selected patches, use the Show Hidden Patches check box at the top of the page and click Redisplay Data to refresh the page. The default values are set in the Patch Wizard Preferences page.
- Included in Aggregate Patch Impact: Yes or No.
- Impact: (For unapplied or missing patches only) Click this icon to access the Patch Impact Analysis page. If you submitted a specific patch to analyze, click the Impact icon on the Recommended Patches Request page to view this report.

Click the Download button to transfer the selected patch number(s) to the Download Patches page for submission.

Codelevels Introduced by the Patch

If new codelevels are introduced by the patch, the Yes indicator in the Codelevel Introduced column of the Recommended Patches Results page is a link that takes you to the Codelevels Introduced by the Patch page.

Codelevels Introduced by the Patch Page

The screenshot shows the Oracle Applications Manager interface. At the top, it says "ORACLE Applications Manager" with navigation links for "Support Cart", "Setup", "Home", "Logout", and "Help". Below this is the "Applications Dashboard" and "Site Map". The main content area is titled "Codelevels Introduced by the Patch : ADOAGC03" and "Last Updated : 15-Jun-2007 00:15:22".

Patch: 5082400

Product **au** Reason Recommended **Maintpack**
Description **12.0.1 : RUP1 FOR ORACLE APPLICATIONS RELEASE 12**

ⓘ This table lists all the Codelevels introduced by the above patch. To see the Codelevels of all products and product families available in your system, see the [Codelevels Summary](#) page

Abbreviation	Name	Type	Codeline	New Codeline	Codelevel Introduced
pj_pf	Projects Suite	product_family		Yes	A
pjb	Project Billing	product		Yes	A
pjc	Project Costing	product		Yes	A
pjf	Project Foundation	product		Yes	A
pjl	Project Collaboration	product		Yes	A
pjr	Project Resource Management	product		Yes	A
pjt	Project Management	product		Yes	A

[Add to Support Cart](#)

Support Cart Setup Home Logout Help

Copyright 2001, 2006 Oracle Corporation. All Rights Reserved.
About Oracle Applications Manager Version 2.3.1

This page contains the following information.

Patch Information

This section includes the following information:

- Product: Product name associated with the patch.
- Description: Describes the selected patch.
- Reason Recommended: The reason for which the patch is recommended.

Codelevels Information

This section includes the following information:

- Abbreviation: The abbreviation for the product, product family, or feature to which this patch applies.
- Name: The full name of the product, product family, or feature to which this patch applies.
- Type: Indicates whether this patch applies to a product, product family, or feature.
- Codeline: Indicates the codeline of the product, product family, or feature in your current system.
- New Codeline: Indicates whether this patch introduces a new codeline.

- Codelevel Introduced: The new codelevel that this patch introduces for the product, product family, or feature.

To view the codelevels of all the products and product families available and in use in your system, click the Codelevels Summary link in the Patch Information section.

Codelevels Summary Page

ORACLE Applications Manager

Support Cart Setup Home Logout Help

Applications Dashboard | Site Map

Applications System:ADOAGC01 > Patch Wizard > Patch Wizard Preferences >

Codelevels Summary : ADOAGC01

Last Updated : : 05-Feb-2008 20:35:25

Filter Name contains Go

Patch Wizard and AutoPatch look at the In Use flag to determine what products you are using. Make sure you do not uncheck any products that are active in your system or are required for system operation.

Short Name	Name	Type	Codeline	Codelevel	In Use
ad	Applications DBA	product	A	a.3	<input checked="" type="checkbox"/>
ahl	Complex Maintenance, Repair & Overhaul	product	A	A	<input checked="" type="checkbox"/>
ak	Common Modules	product	A	A	<input checked="" type="checkbox"/>
alr	Alert	product	A	A	<input checked="" type="checkbox"/>
ame	Approval Management Engine	product	A	A	<input checked="" type="checkbox"/>
aml	Leads Management	product	A	A	<input checked="" type="checkbox"/>
ams	Marketing	product	A	A	<input checked="" type="checkbox"/>
amv	Marketing Encyclopedia System	product	A	A	<input checked="" type="checkbox"/>
amw	Internal Controls Manager	product	A	A	<input checked="" type="checkbox"/>
ap	Payables	product	A	A	<input checked="" type="checkbox"/>
ar	Receivables	product	A	A	<input checked="" type="checkbox"/>
as	Sales Foundation	product	A	A	<input checked="" type="checkbox"/>
asf	Sales Online	product	A	A	<input checked="" type="checkbox"/>
asg	CRM Gateway for Mobile Devices	product	A	A	<input checked="" type="checkbox"/>
asl	Sales Offline	product	A	A	<input checked="" type="checkbox"/>

Copyright © 2001, 2008 Oracle Corporation. All Rights Reserved. About Oracle Applications Manager Version 2.3.1

This page summarizes available products. For each product, it indicates the current Codeline and Codelevel and whether it is in use (active) in your system. Patch Wizard and AutoPatch look at the In Use flag to determine active products.

Caution: If you use this page to update your In Use products list, be sure you do not uncheck any products that are active in your system, or are required for system operation.

Patch Impact Analysis

From the Recommended Patches Results page, click an icon in the Impact column to view the Patch Impact Analysis page for that patch.

ORACLE Applications Manager Support Cart Setup Home Logout Help

Applications Dashboard | Site Map
 Applications System: ADOAGC03 > Patch Wizard > Recommended Patches Results >

Patch Impact Analysis for Patch 4630372-R12: ADOAGC03

Patch Description **R12 izu: R12.IZU.A**

Patch Readme

Total Files in Patch **4**

Files to install **4 (100.00%)**

Direct Impact Summary	Indirect Impact Summary
Applications Patched 2	Unchanged Files Affected 0 JSPs
File Types Installed 2	Menu Navigation Trees Affected 0 Responsibilities, 0 Paths
New Files Introduced 4	Diagnostics Tests to Re-Run 0 Test(s)
Existing Files Changed 0	
Flagged Files Changed 0	
Existing Files Unchanged 0	
Non-US Language Patches Required 0	

TIP Analysis on Unchanged Files Affected only available for JSPs
TIP Click on the Prerequisite Patches link to toggle between Aggregate and Individual Impact Analysis
TIP Aggregate Impact Analysis only for patches with metadata uploaded from InfoBundle.zip
TIP Click on Patch ID in the Aggregate Impact Analysis Table to view individual Impact Analysis for Pre-reqs
TIP Non-US Language Patches are considered required when there are non-US languages installed and the base patch contains new or changed files of translated file types.

[Add to Support Cart](#)

Support Cart Setup Home Logout Help
 Copyright 2001, 2006 Oracle Corporation. All Rights Reserved.
 About Oracle Applications Manager Version 2.3.1

This page displays a list of summary information about which files are new, which files are changed, and which files are ignored when you apply the patch. Prerequisite patches and the readme file for this patch are also shown. Each of the summary items is a link to more complete information.

The key information on this page is separated into these sections.

General Patch Information

General patch information includes:

- Patch Description: Describes the patch.
- Patch Readme: Click this icon to see the readme file for the patch.
- Total Files in Patch: The total number of files in the patch. Click the number link to access the Patch Impact Details page, which lists each file in the patch.
- Files to Install: The number of files the patch will install.

Summary Information

There are two types of summary information: Direct Impact and Indirect Impact. Each summary item is a link to a page that lists the details for the summary count. For example, if you click the number of Existing Files Changed, the details might look similar to this:

ORACLE Applications Manager

Support Cart Setup Home Logout Help

Applications Dashboard | Site Map

Applications System: ADOAGC01 > Patch Wizard > Recommended Patches Results > Patch Impact Analysis >

Patch Impact File Details for Patch 6435000-R12: ADOAGC01

Apply Filter

Patch Description: R12 au: ORACLE E-BUSINESS SUITE 12.0.4 RELEASE UPDATE PACK (RUP4)

App Short Name:

File Name:

Directory:

Impact Type:

Object Type:

View Only Flagged Files

View Only Branched Files

Application	Directory	File Name	Impact Type	Version in APPL_TOP	Version in Patch	Objects Affected	Flagged Files
[AD] Applications DBA	java/clone/uttl	CloneCleaner.class	Changed File	120.5	120.5.12000000.1	N/A	No
[AD] Applications DBA	java/oam/bobj/history	CodeLevels.class	Changed File	120.3	120.3.12000000.1	N/A	No
[AD] Applications DBA	java/oam/bobj/history	CodeLevelsComponent.class	Changed File	120.1	120.1.12000000.1	N/A	No
[AD] Applications DBA	java/advisor	PAProcessExec.class	Changed File	120.6	120.6.12000000.1	N/A	No
[AD] Applications DBA	java/oam/bobj/advisor	RecommCodeLevels.class	Changed File	120.4	120.4.12000000.1	N/A	No
[AD] Applications DBA	java/oam/bobj/advisor	RecommCodeLevelsComparator.class	Changed File	120.2	120.2.12000000.1	N/A	No
[AD] Applications DBA	java/oam/bobj/advisor	RecommCodeLevelsComponent.class	Changed File	120.1	120.1.12000000.1	N/A	No
[AD] Applications DBA	java/oam/bobj/timing	TimingReports.class	Changed File	120.7	120.7.12000000.1	N/A	No
[AD] Applications DBA	java/oam/bobj/timing	TimingReportsComparator.class	Changed File	120.4	120.4.12000000.1	N/A	No
[AD] Applications DBA	java/oam/bobj/timing	TimingReportsComponent.class	Changed File	120.5	120.5.12000000.1	N/A	No
[AD] Applications DBA	java/oam/bobj/advisor	ViewAdvice.class	Changed File	120.13	120.13.12000000.1	N/A	No
[AD] Applications DBA	java/oam/bobj/advisor	ViewAdviceComparator.class	Changed File	120.2	120.2.12000000.1	N/A	No
[AD] Applications DBA	java/oam/bobj/advisor	ViewAdviceComponent.class	Changed File	120.4	120.4.12000000.1	N/A	No
[AD] Applications DBA	java/oam/bobj/advisor	ViewPAResultsComparator.class	Changed File	120.3	120.3.12000000.1	N/A	No
[AD] Applications DBA	java/oam/bobj/advisor	ViewPAResultsComponent.class	Changed File	120.7	120.7.12000000.1	N/A	No

Previous 1-15 of 27246 Next 15

TIP Use (%) as wildcard in filters

TIP Use Application Short Name in filter

Apply Filter

Support Cart Setup Home Logout Help

Copyright 2001, 2008 Oracle Corporation. All Rights Reserved.
About Oracle Applications Manager Version 2.3.1

Direct patch impact includes:

- **Applications Patched:** The number of products that will have files updated. Click the number link to see details of each product affected, and how.
- **File Types Installed:** The number of different file types in the patch. Click the number link to see the file types and how they impact the system.
- **New Files Introduced:** The number of new files that will be introduced by the patch. Click the number link to details about each new file introduced.
- **Existing Files Changed:** The number of existing files in the system that will be changed by the patch. Click the number link to see the existing files changed and the new version numbers.
- **Flagged Files Changed:** The number of custom files that will be changed by this patch. Click the number link to identify the custom files changed by this patch.
- **Existing Files Unchanged:** The number of files unchanged because the version in the patch is older than the version in the system. Click the number link to see the files in the patch that are of the same or of earlier versions than those currently in the system.
- **Non-US Language Patches Required:** If the patch supports multiple languages, click the number link to identify the other languages available.

Indirect summary information includes:

- **Unchanged Files Affected:** The number of system files with dependencies on patched files.
- **Menu Navigation Trees Affected:** The number of menu navigation trees that will be updated by the patch.

Register Flagged Files

With the Register Flagged Files tool, you can record any files in which you have made customizations. In previous releases of Oracle E-Business Suite, the `applcust.txt` file contained the records for all customized files. You had to maintain your custom files records in this file. In this release, information about customized files is still written to the `applcust.txt` file. However, with the Register Flagged Files tool you can download the files and maintain them in a web-based interface.

The Register Flagged Files tool displays the following information about customized files:

- Product abbreviation
- Directory where the files are located
- Name of modified file
- Comments

The Register Flagged Files Interface

The Register Flagged Files tool is a Web-based utility in Oracle Applications Manager. From the Register Flagged Files home page, you can import, export, add, delete, and view records of customized files.

Accessing Register Flagged Files

To access the Register Flagged Files tool, log in to Oracle Applications Manager (OAM) and choose Register Flagged Files from the Site Map.

Step 1: Log in to Oracle Applications Manager

Follow the instructions in *Accessing Patch Wizard*, page 4-4 to access OAM. From the Applications Dashboard, click the Site Map tab.

Step 2: Go the Register Flagged Files home page

On the Site Map page, Register Flagged Files is included on the Maintenance tab under the Patching and Utilities heading. Click the Register Flagged Files link to go to the home page.

Step 3: Select filter criteria

From the Register Flagged Files home page, you can search the records of customized files by product abbreviation, directory, file name, or a combination of product abbreviation, and directory or file name.

Register Flagged Files Home Page

This section describes the Register Flagged Files home page.

Register Flagged Files Home Page

ORACLE Applications Manager

Support Cart Setup Home Logout Help

Applications Dashboard | Site Map
Applications System:ADOAGC03 >

Register Flagged Files : ADOAGC03

Last Updated : 14-Aug-2006 16:50:50

Select Feature Register Flagged Files Go

Import Export Add Cancel Apply

Filter Criteria

Product Abbreviation contains

Directory contains

Go Clear

Select and ... Delete

Select All | Select None

Select	Product Abbreviation	Directory	File Name	Comments
<input type="checkbox"/>	FND	3rdparty/wintertree	accent.tlx	You can put your local directory name in here to note the location of the customized files.
<input type="checkbox"/>	FND	3rdparty/wintertree/N	sscenb2.clx	The directory of the customized files is : xx/3rdparty/wintertree/N
<input type="checkbox"/>	FND	3rdparty/wintertree/NL	sscedu.tlx	

Add to Support Cart Import Export Add Cancel Apply

Support Cart Setup Home Logout Help

Copyright 2001, 2006 Oracle Corporation. All Rights Reserved.
About Oracle Applications Manager Version 2.3.1

Use the following buttons to:

- **Import:** Import a list of customized files from the applcust.txt file in csv format. Use this option to import a local applcust.txt file to the current system or from another system to the current one.
- **Export:** Export a list of customized files in csv format. The default name of the exported file is oamreport.csv. You can use this function to export a list of customized files for import into another system.
- **Add:** Add a record of a customized file.
- **Cancel:** Return to the Applications Dashboard home page.
- **Apply:** Apply and save any changes made to the Comments field of the list of customized files.

You can filter results either by product abbreviation, by directory/file name, or by a combination of product abbreviation, and directory or file name.

- **Product Abbreviation:** To search by product abbreviation, enter the abbreviation of the product for which the customization is made.
- **Directory/File Name:** You can filter the results by directory or file name. Enter the directory or file name for which the customization is made.

List of Customized Files

The list of customized files appears at the bottom of the Register Flagged Files home page. Each line item represents a customized file.

The details provided for each line item are:

- **Select:** Select and delete the corresponding customized file. You can select and delete one file or multiple files at a time.
- **Product Abbreviation:** The abbreviated name of the Oracle E-Business Suite product family for which there is a customization.
- **Directory:** Directory path of the customized file.
- **File Name:** Name of the modified file.
- **Comments:** Use this area to add any comments associated with the customization. Oracle recommends using this area to record the exact location of the customized file.

Adding a Flagged File

Click the Add button on the Register Flagged Files main page to access the Add Flagged Files page. From this page, you can add customized files. Use the Filter Criteria section to search for files you want to add.

Add Flagged Files Page

ORACLE Applications Manager

Support Cart Setup Home Logout Help

Applications Dashboard | Site Map

Applications System: ADOAGC03 > Register Flagged Files >

Search Results Selected Data

Add Flagged Files : ADOAGC03

Last Updated : 14-Aug-2006 16:55:31

Apply Cancel

Filter Criteria

At least one filter must have a value.

Product Abbreviation contains gl

Filename contains glcon

Go Clear

Search Results

Select and ... Add

Select All | Select None

Select	Product Abbreviation	Directory	File Name
<input type="checkbox"/>	GL	admin/driver	glcon.drv
<input type="checkbox"/>	GL	patch/115/import/US	glconreg.ltd
<input type="checkbox"/>	GL	patch/115/odf	glcon.odf

Selected Data

Product Abbreviation	Directory	File Name	Comments	Delete
GL	patch/115/import/US	glconreg.ltd	The local directory is xx/patch/115/import/US	

You can filter results either by product abbreviation, by directory/file name, or by a combination of product abbreviation, and directory or file name.

- **Product Abbreviation:** To search by product abbreviation, enter the abbreviation of the product for which the customization is made.
- **Directory/File Name:** You can filter the results by directory or file name. Enter the directory or file name for which the customization is made.

From the Search Results section, select a file, then click the Add button. The files you select appear in the Selected Data section. Use this section to add any comments you have for each file. Then click the Apply button to confirm your selection.

Importing a Flagged File

Click the Import button on the Register Flagged Files main page to access the Import Flagged File page. From this page, you can import one file or a list of customized files in csv format.

You can import an existing applcust.txt file or a file from another system to the current one. This feature saves you the time of using the Add button to add flagged files individually.

Import Flagged File Page

ORACLE Applications Manager

Support Cart Setup Home Logout Help

Applications Dashboard | Site Map

Applications System:ADOAGC03 > Register Flagged Files >

Import Flagged File : ADOAGC03

i The imported file must be in csv format.

Select File:

Support Cart Setup Home Logout Help

Copyright 2001, 2008 Oracle Corporation. All Rights Reserved.
About Oracle Applications Manager Version 2.3.1

In the Select File field, enter the name of the file you want to import or click the Browse button to navigate to the file. Then click Import. Click Cancel to return to the Register Flagged Files main page.

Patch Reporting Utilities

Applied Patches

With the Applied Patches reporting tool, you can view information about the patches applied to your system. This patch history includes information such as:

- Patch number
- Driver file name
- Platform
- APPL_TOP on which the patch was applied
- Contents and language of the patch
- Files changed or copied
- Bug fixes included in each driver file
- Whether the fix was applied successfully, or reason it was not applied
- Timing information (start time, end time, elapsed time during application, restart time)

How Patch Information Is Stored

AutoPatch stores patch information in the database automatically each time it successfully applies a patch. However, if the patch is not applied successfully, or when you run AutoPatch in pre-install mode, patch history is not written directly to the database, but instead is written to these *patch information files*:

- `javaupdates<YYYYMMDDhhmiss>.txt`, which contains information about changes

to Java files

- adpsv<YYYYMMDDhhmiss>.txt, which contains information about changes to all files *except* Java files

Note: In the file name, *hh* is in 24-hour format.

Both files are located in the <APPL_TOP>/admin/<SID> directory. Each time you run AutoPatch, it checks this directory for the existence of the patch information files. If it finds them, it automatically uploads the information they contain to the patch history database. If the upload is successful, AutoPatch then deletes the files from the directory. The AutoPatch log file records whether the upload was successful or unsuccessful.

AutoPatch Modes

The way you run AutoPatch affects the way it stores patch history information. When you apply a patch in test mode (using *apply=no* on the command line), AutoPatch does not write to the patch information files, and it does not upload patch history information to the database. When you apply a patch in pre-install mode (using *preinstall=y* on the command line), AutoPatch writes patch history information to the patch information files, and it uploads the contents of these files to the database the next time it runs. See: AutoPatch Modes, page 2-31.

Note: Running AutoPatch interactively or non-interactively does not affect the way information is stored in the database.

The Applied Patches Interface

The Applied Patches reporting tool is a Web-based utility in Oracle Applications Manager. The Simple Search page serves as a home page.

Simple Search page

From this page, you can perform a simple search or access the Advanced Search page. You can use either of these pages to query the database for applied patches (the default) or to see a history of changed files. The results of either type of query appear at the bottom of the search page.

Patch Details page

In the search results for both applied patches or file history, there is a Details column. Clicking any link in this column accesses the Patch Details page. From this page, you can go to the Timing Details page, the Files Copied page, the Bug Fixes page, or the Action Summary page.

Note: The discussion of each page contains more detail. The OAM help feature also contains information about the Applied Patches utility.

Accessing Applied Patches Information

To query the patch history database for information about patches applied to your system and the files affected, log in to Oracle Applications Manager and choose Applied Patches from the Site Map.

Step 1: Log in to Oracle Applications Manager

Follow the instructions in Accessing Patch Wizard, page 4-4 to access OAM. From the Applications Dashboard, click the Site Map tab.

Step 2: Go the Simple Search page

From the Site Map, Applied Patches is on the Maintenance tab under the Patching and Utilities heading. Click the Applied Patches link to go to the Simple Search page.

Step 3: Select search criteria

From the Simple Search page, you can perform a query for applied patches or file history. Or, you can go to the Advanced Search page to perform a more detailed search.

Note: See detailed descriptions of individual pages in this chapter. See also Analyzing Applied Patches, page 3-54.

Applied Patches Search Pages

This section describes queries for applied patches.

Simple Search

You can perform a Simple Search from this page by entering the required information in the input fields.

Applied Patches Page - Simple Search

ORACLE Applications Manager

Support Cart Setup Home Logout Help

Applications Dashboard Site Map

Applications System: ADOAGC04 >

Applied Patches : ADOAGC04 Select Feature Applied Patches Go

Last Updated : 08-Jun-2006 18:46:09

Simple Search

At least one field must be completed. Advanced Search

Querying by a specific patch ID will return all patches that are included in the specified patch ID.

Patch

(Enter either a 7-digit patch number or merge patch name)

Applied Within Last Days

Applied From Date To Date

(dd-MMM-yyyy) (dd-MMM-yyyy)

Language

Go Reset

A Bug Fix resolves a specific issue and a patch may contain one or more Bug Fixes.

Previous 1-25 of 36 Next 11

Patch Name	Patch Description	Merged Patches	APPL_TOP Name	Language	Completion Date	Details
4502962.R12		No	ADOAGC04_ap6191rt	US	30-May-2006 03:46:17	
4502962.AD.1.0		No	ADOAGC04_ap6191rt	US	30-May-2006 03:46:17	
5106018.R12		No	ADOAGC04_ap6191rt	US	29-May-2006 08:28:01	
7260000.A		No	ADOAGC04_ap6191rt	US	26-May-2006 11:33:27	
5251362.AD.1.0		No	ADOAGC04_ap6191rt	US	26-May-2006 11:31:21	

There are four fields in the Simple Search section:

- Patch: Enter the patch number in this field.
- Applied Within Last <number> Days: This field allows you to restrict the search to a specific timeframe. The default is 60 days.
- Applied From Date <begin date> To Date <end date>: This field allows you to search for patches that were applied during a specified period of time. Click the calendar icon to select the date or enter the date directly in the field. Some examples for the use of this field are:
 - Enter only the begin date. This search returns all patches applied from the begin date through today's date.
 - Enter only the end date. This search returns all patches applied up to the end date.
 - Enter the begin date and the end date. This search returns all patches applied between the begin date and the end date.
- Language: This drop-down list allows you to select the language of a patch to be queried. You can select only one language in this field. To select multiple languages, go to the Applied Patches Advanced Search page.

You must enter a value in at least one of the fields. If you do not, an error page reminds you to go back and enter a value. To submit the query, click the Go button. The Reset button clears the entered search criteria.

Advanced Search

Click the Advanced Search button to see the Applied Patches Advanced Search page, then enter the search criteria information.

Applied Patches Page - Advanced Search

The screenshot shows the Oracle Applications Manager interface for the 'Applied Patches' section. The page title is 'Applied Patches : ADOAGC04'. The 'Advanced Search' section is active, showing various search criteria fields. The 'Applications System Name' is set to 'ADOAGC04'. The 'APPL_TOP' section has 'Name' selected. The 'Product' field is empty. The 'Patch' field is empty. The 'Applied Within Last' field is set to '60' days. The 'Applied From Date' and 'To Date' fields are empty. The 'Language' section shows a list of available languages and a selected languages list.

ORACLE Applications Manager Support Cart Setup Home Logout Help

Applications Dashboard Site Map
Applications System: ADOAGC04 >

Applied Patches : ADOAGC04 Select Feature Applied Patches Go

Last Updated : 08-Jun-2006 18:47:23

Advanced Search Simple Search

*** Indicates required field.**

Querying by a specific patch ID will return all patches that are included in the specified patch ID.

* Applications System Name ADOAGC04

APPL_TOP
 Name
 All APPL_TOP Server Types
 Specific Server Type
 Form Concurrent Web Administration

Product
Patch

Applied Within Last 60 Days
Applied From Date To Date

Language

Available Languages
AR - Arabic
CS - Czech
D - German
DK - Danish
E - Spanish
EL - Greek
ESA - Latin American Spanish
F - French
FRC - Canadian French
HR - Croatian

Selected Languages

Go Reset

On the Advanced Search page, there are additional search criteria to narrow the results of a query:

- Applications System Name (required): Defaults to the name of your Oracle E-Business Suite system. If you have migrated applied patches information from another system, and want to search those records, enter the name of that system.
- APPL_TOP: Select Name and enter the name of the APPL_TOP where the patches were applied.
- Product: Enter the product short name of the product that owns the patch in this field. The product short names for gl, ap, and fa are SQLGL, SQLAP, and OFA respectively. For all other products, the short name is the uppercase equivalent to the product abbreviation. For example, "AD" or "INV". This field is not case sensitive.
- Patch: Enter the patch number in this field.
- Applied Within Last <number> Days: Restricts the timeframe during which the patches were applied.

- Applied From Date <begin date> To Date <end date>: Narrows the search to a specified period of time. Click the calendar icon to select the date or enter the date directly in the field.
- Language: Select the language of a patch to be queried. Select one language or multiple languages in the Available Languages box and click the Move button.

Search Results

After a search, the results appear at the bottom of the search page. If the results section contains multiple pages of retrieved information, use the Previous and Next links or the drop-down list to navigate from page to page. The retrieved patch information is presented in increments of 25 line items per page. Each line item represents an applied patch.

The details provided for each patch are:

- Patch Name: Name of the patch.
- Patch Description: Describes the patch.
- Merged Patches: Lists patches that have been merged.
- APPL_TOP Name: Name of the APPL_TOP where the patches were applied.
- Language: Patch language.
- Completion Date: Date and time the patch application was completed.
- Details: Provides access to the Patch Details report.

Click a Details icon in the report to open the Patch Details report, which provides details for a specific patch. From the Patch Details report, you can drill down and access reports showing timing details for the patch, all files copied to the file system by this patch, all bug fixes that were applied by this patch, and all actions taken by the patch driver.

Note: For more information on the Patch Details report, see: Patch Details, page 5-9.

File History Search Pages

To search for files that have been updated by a patch, click the File History option in the Select Feature drop-down list on the Applied Patches search pages.

Simple Search

You can perform a Simple Search from this page by entering the required information in the input fields.

File History Page - Simple Search

The screenshot shows the Oracle Applications Manager interface. The top navigation bar includes 'Support Cart', 'Setup', 'Home', 'Logout', and 'Help'. The main content area is titled 'File History : ADOAGC04'. Below this, there is a 'Simple Search' section with a form. The form includes a 'File Name' field, an 'Applied Within Last' dropdown set to '30 Days', 'Changed From Date' and 'To Date' fields with calendar icons, and a 'Language' dropdown. A 'Go' button is at the bottom of the form. Below the form is a table of search results. The table has the following columns: APPL_TOP Name, Product, Directory, File, Version, Changed Date, Patch Details, and Action. The table contains several rows of data, including file names like 'adpamlsd.sql' and 'adpaupg.sql'.

APPL_TOP Name	Product	Directory	File	Version	Changed Date	Patch Details	Action
ADOAGC04_ap6191rt	AD	patch/115/sql	adpamlsd.sql	120.6	17-May-2006 15:18:35	94464216	
ADOAGC04_ap6191rt	AD	patch/115/sql	adpamlsd.sql	120.4	16-May-2006 16:28:22	4461237	
ADOAGC04_ap6191rt	AD	patch/115/sql	adpaupg.sql	120.4	16-May-2006 16:28:22	4461237	
ADOAGC04_ap6191rt	AD	patch/115/sql	adpfamsd.sql	120.1	17-May-2006 15:18:35	94464216	

The following fields are in the Simple Search section:

- **File Name (required):** Enter the name of a file in this field. Do not include a directory path. This field is case-sensitive and accepts a % wildcard symbol in combination with literal characters.
- **Applied Within Last <number> days:** Enter the number of days to include in the search. The default is 60 days.
- **Changed From Date <begin date> To Date <end date>:** Search for files that were updated during a specified period of time. Click the calendar icon to select the date or enter the date directly in the field. Some examples for the use of this field are.
 - Enter only the begin date. This search returns file history information from the begin date through today's date.
 - Enter only the end date. This search returns file history information up to the end date.
 - Enter the begin date and the end date. This search returns file history information between the begin date and the end date.
- **Language:** Select the language of a file to be queried. You can select only one language in this field. To select multiple languages, go to the File History Advanced

Search page.

To submit the query, click the Go button. The Reset button clears the entered search criteria.

Note: If you have not entered a value in the File Name field, a message prompts you to go back and complete the field.

Advanced Search

Click the Advanced Search button. Then enter the search criteria information on the Advanced Search page.

File History Page - Advanced Search

The screenshot shows the Oracle Applications Manager interface for the File History page. The page title is "ORACLE Applications Manager" and the breadcrumb is "Applications System: ADOAGC04 > File History : ADOAGC04". The page includes a "Select Feature" dropdown set to "File History" and a "Go" button. Below this, there is a "Last Updated" timestamp of "08-Jun-2006 18:58:46". The "Advanced Search" section contains several fields: "Applications System Name" (ADOAGC04), "APPL_TOP Name", "File Name", "Latest Version Only" (radio buttons for Yes and No, with No selected), "Applied Within Last" (60 Days), "Changed From Date" (calendar icon), and "To Date" (calendar icon). A "Language" section has two lists: "Available Languages" (AR - Arabic, CS - Czech, D - German, DK - Danish, E - Spanish, EL - Greek, ESA - Latin American Spanish, F - French, FRC - Canadian French, HR - Croatian) and "Selected Languages". Navigation buttons include "Go", "Reset", and "Simple Search".

There are additional search criteria on the Advanced Search page to narrow the results of a query:

- Applications System Name (required): Defaults to the name of your Oracle E-Business Suite. If you have migrated file history information from another system, and want to search those records, enter the name of that system.
- APPL_TOP name: Name of the APPL_TOP containing the file.
- File Name (required): Enter the name of a file in this field. Do not include a directory path. This field is case-sensitive and accepts a % wildcard symbol in combination with literal characters.
- Latest Version Only: The options are Yes or No. Yes returns information for only

the latest version of the file. No returns information for all versions of the selected file.

- Applied Within Last <number> days: Enter the number of days.
- Changed From Date <begin date> To Date <end date>: Search for file history information spanning a specified period of time. Click the calendar icon to select the date or enter the date directly in the field.
- Language: Select the language of a patch to be queried. Select one language or multiple languages in the Available Languages box and click the Move button.

Search Results

After a search, the results appear at the bottom of the page. Each line item represents a file that was changed due to its inclusion in a patch. The details provided for a file are:

- APPL_TOP Name: This is the name of the APPL_TOP containing the files.
- Product: Name of the product that owns the file.
- Directory: Directory path where the file is located.
- File: Name of the file.
- Version: Version number of the file.
- Changed Date: Date this version of the file was updated by a patch.
- Patch Details: Click on the patch number to see the Patch Details report for the patch in which the file was included.
- Action: Click on the icon to see the Action Summary report for the action that updated the file.

If a file has never been patched, the message "The above criteria resulted in no rows" appears in the APPL_TOP Name column. If the number of files retrieved exceeds 200, the report lists only the first 200 files. Use the filter to reduce the number of files in the report.

Patch Details

From the Applied Patches page, click the Details icon in a selected row from the results section, or from the File History page, click the patch number link in the Patch Details column, to open the Patch Details report. This report provides details for a specific patch. The patch summary information is carried over and appears at the top of the Patch Details report.

Patch Details Report

The screenshot shows the Oracle Applications Manager interface. At the top, it says 'ORACLE Applications Manager' with navigation links for 'Support Cart', 'Setup', 'Home', 'Logout', and 'Help'. Below this is the 'Applications Dashboard' and 'Site Map'. The main content area is titled 'Patch Details : 6016104.A : ADOAGC03'. It displays the following information:

- Last Updated : 11-May-2007 18:17:41
- APPL_TOP Name : ap6277rt
- Patch : 6016104.A
- Merged Patches : No
- Language : US
- Drivers Applied : 1
- Completion Date : 03-May-2007 03:41:39
- Patch Description : R12.0.2 - CONSOLIDATED PATCH FOR ALL ADOAM BUG FIXES

Below this information is a note: 'If the same patch is applied multiple times, this page displays all of them.' There are four tabs: 'Select Driver File and view...', 'Timing Details', 'Files Copied', 'Bug Fixes', and 'Action Summary'. The 'Select Driver File and view...' tab is active, showing a table with the following data:

Select	Driver File	Start Date	End Date	AutoPatch Options	Platform	Patch Top	Codelevel Introduced
<input checked="" type="radio"/>	u6016104.drv	03-May-2007 03:34:19	03-May-2007 03:41:39	hotpatch	GENERIC	/ SLOTS/ slot03/ appmgr/ patch/ 6016104	

At the bottom right of the table area is a button labeled 'Add to Support Cart'. The footer of the page contains the text: 'Support Cart Setup Home Logout Help' and 'Copyright 2001, 2008 Oracle Corporation. All Rights Reserved. About Oracle Applications Manager Version 2.3.1'.

This report contains the following information:

- **Select:** This option button determines which driver file details are presented in the Timing Details report, Files Copied report, the Bug Fixes report, or the Action Summary report.
- **Driver File:** Name of the driver file.
- **Start Date:** Date and time the application of the driver file began.
- **End Date:** Date and time the application of the driver file was complete.
- **AutoPatch Options:** Displays any command line options used to run the driver file.
- **Platform:** Platform of the driver file.
- **Patch Top:** Location of the driver when it was run.
- **Codelevel Introduced:** Link to the Codelevel Introduced report for the patch.

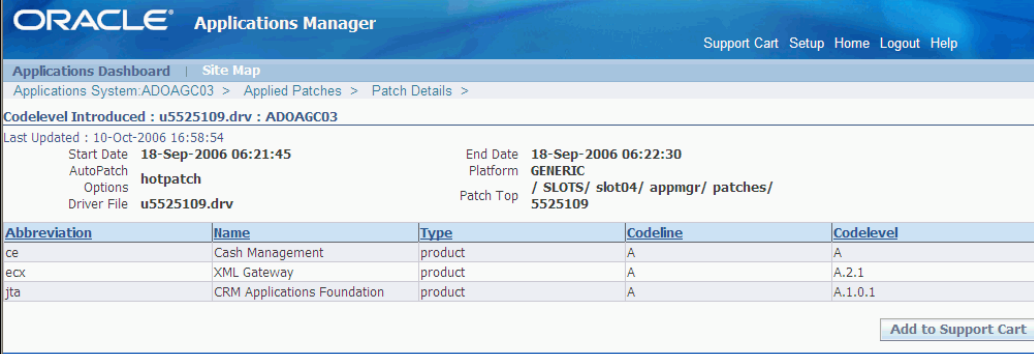
To see additional details for a patch, click one of the following buttons on the report:

- **Timing Details:** Takes you to the AutoPatch Timing Details report.
- **Files Copied:** Takes you to the Files Copied report.
- **Bug Fixes:** Takes you to the Bug Fixes report.
- **Action Summary:** Takes you to the Action Summary report.

Codelevel Introduced

From the Patch Details page, click the Codelevel Introduced icon to access the Codelevel Introduced report.

Codelevel Introduced Report



ORACLE Applications Manager Support Cart Setup Home Logout Help

Applications Dashboard | Site Map
Applications System: ADOAGC03 > Applied Patches > Patch Details >

Codelevel Introduced : u5525109.drv : ADOAGC03

Last Updated : 10-Oct-2006 16:58:54

Start Date	18-Sep-2006 06:21:45	End Date	18-Sep-2006 06:22:30
AutoPatch	hotpatch	Platform	GENERIC
Options	hotpatch	Patch Top	/ SLOTS/ slot04/ appmgr/ patches/ 5525109
Driver File	u5525109.drv		

Abbreviation	Name	Type	Codeline	Codelevel
ce	Cash Management	product	A	A
ecx	XML Gateway	product	A	A.2.1
jta	CRM Applications Foundation	product	A	A.1.0.1

[Add to Support Cart](#)

Support Cart Setup Home Logout Help

Copyright 2001, 2006 Oracle Corporation. All Rights Reserved.

This report contains the following information about the codelevel introduced:

- **Abbreviation:** The abbreviation for the product, product family, or feature to which this patch applies.
- **Name:** The full name of the product, product family, or feature to which this patch applies.
- **Type:** Indicates whether this patch applies to a product, product family, or feature.
- **Codeline:** Indicates the codeline of the current product, product family, or feature in the patch. (For example, codeline A for Release 12.0, codeline B for Release 12.1, and so on.)
- **Codelevel:** Indicates the codelevel of the current product, product family, or feature in the patch. (For example, codelevel A.1 for RUP1, A.2 for RUP 2, and so on.)

You can sort each of these columns by clicking the column title at the top of the report.

Timing Details

The AutoPatch Timing Details can also be accessed through the Timing Reports link from the Maintenance tab on the OAM Site Map.

See Timing Reports, page 5-15 for more information on Timing Details.

Files Copied

The Files Copied report lists all files copied to the file system as a result of the actions in the selected driver file. You access this report by selecting a driver file in the Patch Details report and clicking the Files Copied button.

Files Copied Report

The screenshot shows the Oracle Applications Manager interface for the 'Files Copied Report'. The report title is 'Files Copied : u5251362.drv : ADOAGC04'. It displays the start and end dates of the patch, the platform (LINUX), and the patch top path. A filter section allows users to search for files by product and directory. The main table lists the copied files with columns for Product, Directory, File, and Version.

Product	Directory	File	Version
AD	lib	adphist.o	120.14

This report contains the following information about the files copied:

- Product: Short name for the product that owns the file.
- Directory: Directory path where the file was copied.
- File: Name of the file.
- Version: Version number of the copied file.

You can sort each of these columns by clicking the column title at the top of the report. If there are no files copied in the patch, no rows are displayed. If the number of files copied exceeds 200, the report lists only the first 200 files. Use the filter to reduce the number of files in the report.

Bug Fixes

The Bug Fixes report lists all bug fixes included in the selected driver file. Select a driver file in the Patch Details report and click the Bug Fixes button.

Bug Fixes Report

The screenshot shows the Oracle Applications Manager interface for a Bug Fixes Report. The header includes the Oracle logo and 'Applications Manager'. Navigation links for 'Support Cart', 'Setup', 'Home', 'Logout', and 'Help' are present. The breadcrumb trail is 'Applications Dashboard > Site Map > Applications System: ADOAGC04 > Applied Patches > Patch Details > Bug Fixes : u5251362.drv : ADOAGC04'. The report details include: 'Last Updated : 08-Jun-2006 19:06:19', 'Start Date: 26-May-2006 10:31:16', 'AutoPatch Options: hotpatch', 'Driver File: u5251362.drv', 'End Date: 26-May-2006 10:34:19', 'Platform: LINUX', and 'Patch Top: / SLOTS/ slot03/ appmgr/ patches/ 5251362'. A filter section shows 'Bug Fix' selected in a dropdown, followed by 'contains' and an empty search box with a 'Go' button. Below this is a table with columns: Bug Fix, Product, Applied, and Remarks. The table contains two rows: one for bug fix 5251362 with product 'ad' and 'Applied' status 'Y', and another for 4654046 with product 'ad' and 'Applied' status 'Y'. An 'Add to Support Cart' button is at the bottom right. The footer contains 'Support Cart Setup Home Logout Help' and 'Copyright 2001, 2006 Oracle Corporation. All Rights Reserved. About Oracle Applications Manager Version 2.3.1'.

This report contains the following information about bug fixes:

- Bug Fix: Number of the bug fixed as a result of the selected driver file. Some items in this column are links. Clicking a linked item accesses the Action Summary report.
- Product: Short name for the product for which the bug was fixed.
- Applied: Indicates whether the bug fix was applied.
- Remarks: If the bug fix was not applied, the reason is stated here.

You can sort each of these columns by clicking the column title at the top of the report. If there are no bug fixes in the patch, no rows are displayed. If the number of bug fixes exceeds 200, the report lists only the first 200. Use the filter to reduce the number of items in the report.

Action Summary

The Action Summary report provides summary information for the actions of a selected driver file. Each line item represents a performed action. You access this report either by selecting a driver file in the Patch Details report and clicking the Action Summary button, by clicking the Action icon in the File History search results, or by clicking a bug fix number in the Bug Fix column of the Bug Fixes report.

Action Summary Report

ORACLE Applications Manager

Support Cart Setup Home Logout Help

Applications Dashboard | Site Map

Applications System: ADOAGC04 > Applied Patches > Patch Details >

Action Summary : u5251362.drv : ADOAGC04

Last Updated : 08-Jun-2006 19:07:23

Start Date: 26-May-2006 10:31:16
AutoPatch Options: hotpatch
Driver File: u5251362.drv

End Date: 26-May-2006 10:34:19
Platform: LINUX
Patch Top: / SLOTS/ slot03/ appmgr/ patches/ 5251362

Filter: Bug Fix [v] contains [v] [Go]

The details icon displays additional information about database updates.

Previous 1-25 of 189 Next 25

Product	Directory	File	Action	Phase	Run	Bug Fix	Details
FND	include	afugai.h	copy		N	5251362	[Details]
AD	bin	adadmin	link		Y	5251362	[Details]
AD	bin	adident	link		Y	5251362	[Details]
AD	bin	adjkey	link		Y	5251362	[Details]
AD	bin	admrpoch	link		Y	5251362	[Details]

The Action Summary report contains the following summary information:

- Product: Short name for the product that owns the file referenced by the action.
- Directory: Directory path for the file referenced by the action.
- File: Name of the file referenced by the action.
- Action: Type of action performed on the updated file.
- Phase: Phase in which the action occurred.
- Run: Signifies whether the action was executed.
- Bug Fix: Number of the bug fixed as a result of the selected driver file.
- Details: This link is active if AutoPatch performed database actions, usually SQL or EXEC actions where Run = y. Click this link to access the Action Details report.

You can sort each of these columns by clicking the column title at the top of the report. If the number of actions exceeds 200, the report lists only the first 200. Use the filter to reduce the number of items in the report.

Action Details

To access this report, click the Details icon in a selected row of the Action Summary report. The Action Summary information is carried over and presented at the top of the report.

Action Details Report

The screenshot shows the Oracle Applications Manager interface. At the top, it says "ORACLE Applications Manager" with navigation links for "Support Cart", "Setup", "Home", and "Logout Help". Below this is a navigation bar with "Applications Dashboard" and "Site Map". The main content area is titled "Action Details : jtfhcldc.sql : x6Sanity". It shows the action was last updated on 28-Jul-2006 at 18:58:33. Key statistics include 5082400 bug fixes, a phase of "first", product "JTF", and directory "patch/ 115/ sql". Below this, there's a section for "jtfhcldc.sql" with "Before Version" and "After Version" both listed as "N/A". A table follows with columns: Arguments, Command Modifier, Check Object, Elapsed Time, Start Time, Restart Time, End Time, and Restarted?. The table contains one row with values: N/A, sql, JTF_CAL_ADDRESSES &un_jtf &pw_jtf, 0 sec, 12-Jun-2006 12:44:38, 12-Jun-2006 12:44:38, and N. At the bottom right of the table area is an "Add to Support Cart" button. The footer of the page includes "Support Cart Setup Home Logout Help" and "Copyright 2001, 2009 Oracle Corporation. All Rights Reserved. About Oracle Applications Manager/Version 2-3-1".

This report contains the following information about action details:

- Arguments: Specific argument for SQL and EXEC commands.
- Command Modifier: SQL or EXEC command modifier in the database section of the driver.
- Check Object: Name of the database object to check for, along with name and password of the schema where AutoPatch looks for the checked object.
- Elapsed Time: Time required to complete the action.
- Start Time: Date and time the action began.
- Restart Time: Date and time the action was restarted.
- End Time: Date and time the action was complete.
- Restarted?: States whether the action was restarted.

N/A in the report represents action details that are not specified. For example, in the Arguments field, N/A means no additional arguments were specified.

Timing Reports

The Timing Reports utility provides the job history of applied patches. It captures statistics about, and job timing information for, adop and AD Administration maintenance sessions that run parallel workers. Both adop and AD Administration store information about a processing session in database tables. You can access this information, either during the session or after it is complete, through the OAM interface.

Note: You can also access job timing information during a current session or for a completed session by running `adtimrpt.sql` from the command line. This script creates the `adt<session_id>.lst` report. For more information, see: AD Job Timing Report, page 8-4 in the Maintenance section of this book.

During a parallel session, AD utilities assign processing jobs to workers. For jobs that affect the database, job actions are grouped in phases to reduce dependencies between jobs - workers do not have to wait for another worker to complete a dependent job before completing their assigned task. See: Using Parallel Processing, page 7-55 in the Maintenance section of this book.

The Timing Reports utility lists processing tasks and provide details about the elapsed time for phases, jobs, and sessions. The information includes:

- Jobs run successfully on the first try
- Failed jobs that were restarted and then run successfully
- Failed jobs that were skipped
- Long-running jobs
- Summary information for each parallel phase
- Time taken to run a job
- Overall elapsed time for each session

The Timing Reports Interface

The Timing Reports interface consists of a main page, a Timing Details page, and a View Log Files page which provides links to reports about specific maintenance session information.

Main Page

From the Timing Reports main page, you can view a list of all in-progress, stopped, aborted, and completed maintenance sessions. Click the Details icon to access the Timing Details page or click the Log Files icon to access the View Log Files page.

Timing Details Page

There are two types of Timing Details reports - those associated with an adop session and those associated with an AD Administration session.

Note: The discussion of each page contains more detail. The OAM help feature also contains information about the Timing Reports.

View Log Files Page

This page contains a list of log files generated for the corresponding maintenance session.

Accessing Timing Reports

To access the Timing Reports main page, log in to Oracle Applications Manager and choose Timing Reports from the Site Map.

Step 1: Log in to Oracle Applications Manager

Follow the instructions in Accessing Patch Wizard, page 4-4 to access OAM. From the Applications Dashboard, click the Site Map tab.

Step 2: Access Timing Reports

From the Site Map, Timing Reports is on the Maintenance tab under the Patching and Utilities heading. Click the Timing Reports link to go to the main page.

Step 3: Filter the results

There is a filter at the top of the page that allows you to narrow the contents of the list. You can filter based on the following status of the tasks: Any tasks, In-progress tasks, Completed tasks, Stopped tasks, or Aborted tasks. You can also filter by Task Name, Status, Start Date, and Run Time. Click Go to activate the filter.

Note: See descriptions of individual pages in this chapter for details. See also AD Administration, page 7-62 in the Maintenance section of this book.

Timing Reports Main Page

The Timing Reports main page shows information for each maintenance session.

Timing Reports Main Page

Task Name	Status	Start Date	Run Time	Last Update	Details	Log Files
AD Administration - Check for missing files	✓	08-Aug-2006 06:02:05	1 min, 38 sec	08-Aug-2006 06:03:43		
AutoPatch - u3456789.driv	✓	08-Aug-2006 05:03:51	2 min, 16 sec	08-Aug-2006 05:06:07		
AutoPatch - u5400663.driv	✓	02-Aug-2006 07:13:43	38 sec	02-Aug-2006 07:14:21		
AutoPatch - u5414900.driv	✓	02-Aug-2006 07:09:32	2 min, 59 sec	02-Aug-2006 07:12:31		
AutoPatch - u94464216.driv	✓	01-Aug-2006 23:24:59	7 min, 24 sec	01-Aug-2006 23:32:23		
AD Administration - Compile APPS schema	✓	17-Jul-2006 18:15:51	42 sec	17-Jul-2006 18:16:33		
AD Administration - Disable Maintenance Mode	✓	17-Jul-2006 18:15:46	1 sec	17-Jul-2006 18:15:47		
AD Administration - Compile APPS schema	✓	17-Jul-2006 09:43:40	50 sec	17-Jul-2006 09:44:30		
AutoPatch - c2647958.driv	✓	17-Jul-2006 09:40:42	1 min, 1 sec	17-Jul-2006 09:41:43		
AutoPatch - c2647958.driv	✗	17-Jul-2006 09:35:37	54 sec	17-Jul-2006 09:36:31		
AD Administration - Enable Maintenance Mode	✓	17-Jul-2006 09:23:11	1 sec	17-Jul-2006 09:23:12		

- **Task Name:** Name and brief description of the maintenance session.
- **Status:** Status of the timing report. A clock icon means the session is still in-progress, an exclamation icon means the session has stopped, an X icon means the session was aborted (that is, the AD utility was restarted with the gf option), and a check mark means the session has completed.
- **Start Date:** Date and time the maintenance session began.
- **Run Time:** Time required to complete the maintenance session.
- **Last Update:** Time the timing information was last updated.
- **Details:** Access the Timing Details for the maintenance session.
- **Log Files:** Access the log files of the maintenance session.

adop Timing Details

Click the Details icon of a selected row (with an adop task name) in the Timing Reports list to open the adop Timing Details report. This report provides details for a specific session of adop.

adop Timing Details Report

ORACLE Applications Manager

Support Cart Setup Home Logout Help

Applications Dashboard Site Map

Applications System: ADOAGC04 > Timing Reports >

AutoPatch Timing Details : u5251362.drv : ADOAGC04

Last Updated : 08-Jun-2006 19:19:02

Filter: Elapsed Time greater than Go
(Enter number of seconds)

Timing Details

[Expand All](#) | [Collapse All](#)

Focus Task Name	Elapsed Time	Start Date	End Date
AutoPatch			
Run a single patch driver file	1 min, 38 sec	26-May-2006 11:25:53	26-May-2006 11:27:31
Copy portion steps	42 sec	26-May-2006 11:25:55	26-May-2006 11:26:37
Relink executables	32 sec	26-May-2006 11:26:04	26-May-2006 11:26:36
Database portion steps	41 sec	26-May-2006 11:26:37	26-May-2006 11:27:18
Run SQL scripts and EXEC commands	5 sec	26-May-2006 11:26:40	26-May-2006 11:26:45
Running database update commands	5 sec	26-May-2006 11:26:40	26-May-2006 11:26:45
Running SQL and EXEC commands in parallel	5 sec	26-May-2006 11:26:40	26-May-2006 11:26:45
<u>Running parallel SQL and EXEC commands</u>	5 sec	26-May-2006 11:26:40	26-May-2006 11:26:45
Compile invalid objects in DB	30 sec	26-May-2006 11:26:45	26-May-2006 11:27:15
Run ST parallel compile (APPS)	14 sec	26-May-2006 11:26:45	26-May-2006 11:26:59
Steps after generate portion	13 sec	26-May-2006 11:27:18	26-May-2006 11:27:31
Save Patch History	13 sec	26-May-2006 11:27:18	26-May-2006 11:27:31
Save Patch History to database	13 sec	26-May-2006 11:27:18	26-May-2006 11:27:31

Run Information

The adop Timing Details report lists every task performed in a maintenance session. The Timing Details section contains the following information for each task:

- **Focus:** Select the circle icon next to a task to see just the sub-tasks within it.
- **Task Name:** Name of the task. Click the plus-sign icon to expand or contract the sub-tasks within the task. The underlined Task Names are links to the Job Timing report for that particular task.
- **Elapsed Time:** Time required to complete the task. This field is not applicable for stopped or in-progress tasks.
- **Start Date:** Date and time the task began.
- **End Date:** Date and time the task was complete. This field is not applicable for stopped or in-progress tasks.

Use the filter to adjust the list of tasks based on their elapsed time. The default list shows all tasks with elapsed time of greater than 4 seconds. Use the Expand All link to see all sub-tasks and the Collapse All to see just the top-level task.

When you access the adop Timing Details report for a stopped or in-progress task, the page defaults to display the most recently performed sub-tasks. For in-progress tasks, you can use the Refresh icon to get the latest running tasks. The Refresh icon is a picture of a page with a blue circular arrow.

Run Information

Additional adop task information is available by clicking the plus-sign icon for the Run Information section at the bottom of the page. The subsections in Run Information are General, Timing Summary, and Files Installed on this APPL_TOP.

adop Timing Details Report - Run Information

Run Information		Timing Summary	
General	Utility Name: AutoPatch	Start Date:	26-May-2006
	Task: u5251362.drv	End Date:	26-May-2006
	Log File: /slot03/appmgr/ADOAGC04appl/admin/ADOAGC04/log/u5251362.log	Total Run Time:	1 min, 42 sec
	Driver File: /SLOTS/slot03/appmgr/patches/5251362/u5251362.drv		
	Patch Top: /SLOTS/slot03/appmgr/patches/5251362	Files Installed on this APPL_TOP	
	Options: hotpatch	Administration:	Yes
	Platform: LINUX	Java and HTML:	Yes
	Applications System Name: ADOAGC04	Forms:	Yes
	Oracle Database: ADOAGC04	Concurrent Processing:	Yes
	Oracle Home: /slot03/appmgr/ADOAGC04ora/8.0.6		
	APPL_TOP Name: ADOAGC04_ap6191rt		
	APPL_TOP Directory: /slot03/appmgr/ADOAGC04appl		

General

This subsection contains the following information:

- Utility Name: Name of the utility used to perform the task.
- Task: Task performed.
- Log File: Name and location of the log file.
- Driver File: Name and location of the patch driver file.
- Patch Top: Location of the patch driver files.
- Options: Command options used when running adop.
- Platform: Platform of the system.
- Applications System Name: Name of the Applications system on which the task was performed.
- Oracle Database: Name of the database.
- Oracle Home: Directory path to the Oracle home used to link the executables.
- APPL_TOP Name: Name of the APPL_TOP.
- APPL_TOP Directory: APPL_TOP directory path.

Timing Summary

This subsection contains the following information:

- **Start Date:** Date and time the task began.
- **End Date:** Date and time the task was complete. This field does not apply for stopped or in-progress tasks.
- **Total Run Time:** Time required to complete the task. This field does not apply for stopped or in-progress tasks.

Files Installed on this APPL_TOP

This subsection contains the following information:

- **Java and HTML:** States whether the APPL_TOP on which the task was performed supports HTTP (Web) services.
- **Forms:** States whether the APPL_TOP on which the task was performed supports forms services.
- **Concurrent Processing:** States whether the APPL_TOP on which the task was performed supports concurrent processing services.

AD Administration Timing Details

Click the Details icon of a selected row (with an AD Administration task name) in the Timing Reports list to open the AD Administration Timing Details report. This report provides details for a specific session of AD Administration.

AD Administration Timing Details Report

ORACLE Applications Manager

Support Cart Setup Home Logout Help

Applications Dashboard | Site Map

Applications System: ADOAGC04 > Timing Reports >

AD Administration Timing Details : Compile APPS schema : ADOAGC04

Last Updated : 08-Jun-2006 19:22:12

Filter: Elapsed Time greater than 4 Go
(Enter number of seconds)

Timing Details

Expand All | Collapse All

Focus Task Name	Elapsed Time	Start Date	End Date
AD Administration			
Compile APPS schema	43 sec	16-May-2006 12:57:28	16-May-2006 12:58:11
Compile APPS	41 sec	16-May-2006 12:57:30	16-May-2006 12:58:11
Run ST parallel compile (APPS)	20 sec	16-May-2006 12:57:30	16-May-2006 12:57:50
Run AD parallel compile (APPS)	21 sec	16-May-2006 12:57:50	16-May-2006 12:58:11
Parallel Compile - Compile	7 sec	16-May-2006 12:57:55	16-May-2006 12:58:02

Run Information

General		Timing Summary
Utility Name	AD Administration	Start Date
Task	Compile APPS schema	12:57:28
Log File	/ slot03/ appmgr/ ADOAGC04appl/ admin/ ADOAGC04/ log/ adadmin.log	End Date
Driver File	N/A	16-May-2006
Patch Top	N/A	12:58:11
Options	N/A	Total Run Time
Platform	LINUX	43 sec
Applications System Name	ADOAGC04	Files Installed on this APPL_TOP
Oracle Database	ADOAGC04	Administration
Oracle Home	/ slot03/ appmgr/ ADOAGC04ora/ 8.0.6	Java and HTML
APPL_TOP Name	ADOAGC04_ap6191rt	Forms
APPL_TOP Directory	/ slot03/ appmgr/ ADOAGC04appl	Concurrent Processing
		Yes
		Yes
		Yes
		Yes

The Timing Details and Run Information sections contain the same types of information for each task as the adop Timing Details report. See: adop Timing Details, page 5-18.

In-Progress Timing Details

You can use the In-Progress Timing Details page to monitor the job while it is running. Click the Refresh icon to view the steps that are in progress.

In-Progress Timing Details Page

Run Information

In Progress AD Utility : u94464216.drv : ADDAGC05

Last Updated : 10-Aug-2006 13:58:48

Filter: Elapsed Time greater than 4 Go
(Enter number of seconds)

AutoPatch Timing Details

Expand All | Collapse All

Focus Task Name	Elapsed Time	Start Date	End Date	Number of jobs in this task	Number of jobs completed
AutoPatch					
AutoPatch startup after aimini	9 sec	10-Aug-2006 13:57:46	10-Aug-2006 13:57:55	N/A	N/A
Run a single patch driver file		10-Aug-2006 13:58:18			0
Copy portion steps	25 sec	10-Aug-2006 13:58:22	10-Aug-2006 13:58:47	N/A	N/A
Read file driver files to get list of valid files	9 sec	10-Aug-2006 13:58:22	10-Aug-2006 13:58:31	N/A	N/A
Maintain Oracle Applications Java files	12 sec	10-Aug-2006 13:58:35	10-Aug-2006 13:58:47	N/A	N/A
Perform jcopy actions	12 sec	10-Aug-2006 13:58:35	10-Aug-2006 13:58:47	N/A	N/A
Run adjcopy.class	11 sec	10-Aug-2006 13:58:36	10-Aug-2006 13:58:47	N/A	N/A
Database portion steps		10-Aug-2006 13:58:47			0
Get initial list of invalid objects in DB		10-Aug-2006 13:58:47			0

Last Updated : 10-Aug-2006 13:58:48

Run Information Return to Top

Support Cart Setup Home Logout Help

Copyright 2001, 2006 Oracle Corporation. All Rights Reserved.

In the In-Progress AD Utility section, you can filter the results by elapsed time.

The adop Timing Details section contains the following information for each task:

- Focus: Select the circle icon next to a task to see just the sub-tasks within it.
- Task Name: Name of the task. Click the plus-sign icon to expand or contract the sub-tasks within the task. The underlined Task Names are links to the Job Timing report for that particular task.
- Elapsed Time: Time required to complete the task. This field is not applicable for stopped or in-progress tasks.
- Start Date: Date and time the task began.
- End Date: Date and time the task was complete. This field does not apply for stopped or in-progress tasks.
- Number of jobs in this task: Number of jobs contained within each level of the task.
- Number of jobs completed: Number of jobs completed within this level of the task.

Job Timing

The underlined Task Names in the adop Timing Details report and the AD Administration Timing Details report link to the Job Timing report for that particular task. This report provides timing information for each job within the selected task and

allows you to drill down into each task to view any exception reports.

Job Timing Report

Job Timing Summary

Jobs that ran successfully	3	Total Elapsed Time	12 sec
Exceptions	0	Total Job Time	10 sec
Total Number of Jobs	3	Total Number of Workers	1

Filter Criteria

Phase contains (Case sensitive)

Run Time greater than (Enter number of seconds)

Restarted jobs may have Run Time not equal to the difference between Start Time and End Time.

Phase	Product	Directory	File	Action	Start Time	End Time	Run Time	Restarted?
pls	ad	patch/ 115/ sql	adpamis.pls	package	28-Mar-2010 02:22:21	28-Mar-2010 02:22:22	1 sec	N
plb	ad	patch/ 115/ sql	adpamisb.pls	package	28-Mar-2010 02:22:23	28-Mar-2010 02:22:23	0 sec	N
daa+51	fnd	patch/ 115/ import/ US	afoamadmmenu.ldt	bin	28-Mar-2010 02:22:24	28-Mar-2010 02:22:33	9 sec	N

Click a Task Name on the Timing Details report to open the Job Timing report. The Job Timing Summary information appears at the top of the Job Timing report and the details appear at the bottom.

The summary information includes:

- Jobs that ran successfully: Number of successful jobs.
- Exceptions: Number of jobs that were not completed successfully. If exceptions exist, it is a hyperlink to the Exceptionreport.
- Total Number of Jobs: Number of jobs within the task.
- Total Elapsed Time: Time required to complete the task.
- Total Job Time: Time required to complete the jobs within the task.
- Total Number of Workers: Number of workers used to perform the task.

The Job Timing Details section contains the following information for each job:

- Phase: Database processing phase.
- Product: Abbreviation for the product being updated.

- Directory: Directory path of the file run by the job.
- File: File used to perform the job.
- Action: Action of the job.
- Start Time: Date and time the job began.
- End Time: Date and time the job completed.
- Run Time: Total time of the job.
- Restarted?: States whether the job was restarted.

The filters at the top of the Details section allow you to adjust the list of jobs based on the property and run time of jobs. You can filter based on the following properties of the jobs: Phase, Product, Directory, File, Action, or Restarted. Click Go to activate the filter.

Click the Phase Info button to open the Phase Information report.

Phase Information

Clicking the Phase Info button provides timing information by phase for a task selected in either the adop Timing Details report or the AD Administration Timing Details report.

Phase Information Report

The screenshot shows the Oracle Applications Manager interface. At the top, it says 'ORACLE Applications Manager' with navigation links for 'Support Cart', 'Setup', 'Home', 'Logout', and 'Help'. Below this is a breadcrumb trail: 'Applications Dashboard > Site Map > Applications System:ADOAGC04 > Timing Reports > AD Administration Timing Details > Job Timing Report >'. The main heading is 'Phase Information : Run AD parallel compile (APPS) : ADOAGC04'. Below the heading, it shows 'Last Updated : 08-Jun-2006 19:25:22' and 'Driver File N/A'. A table displays the phase information for the task 'Run AD parallel compile (APPS)'. The table has columns for Phase, Start Time, Elapsed Time, Jobs, Total Job Time, Restarted?, and Skipped. The data rows are: 'Parallel Compile - Setup' (1 sec, 1 job, 1 sec, N, 0), 'Parallel Compile - Compile' (7 sec, 8 jobs, 42 sec, N, 0), 'Parallel Compile - Cleanup' (1 sec, 1 job, 1 sec, N, 0), and 'Parallel Compile - Error Check' (3 sec, 1 job, 3 sec, N, 0). At the bottom right of the table area is an 'Add to Support Cart' button. The footer of the page contains 'Support Cart Setup Home Logout Help' and 'Copyright 2001, 2006 Oracle Corporation. All Rights Reserved. About Oracle Applications Manager Version 2.3.1'.

Phase	Start Time	Elapsed Time	Jobs	Total Job Time	Restarted?	Skipped
Parallel Compile - Setup	16-May-2006 12:57:54	1 sec	1	1 sec	N	0
Parallel Compile - Compile	16-May-2006 12:57:55	7 sec	8	42 sec	N	0
Parallel Compile - Cleanup	16-May-2006 12:58:02	1 sec	1	1 sec	N	0
Parallel Compile - Error Check	16-May-2006 12:58:03	3 sec	1	3 sec	N	0

The general information presented at the top of the Phase Information report are:

- Driver File: Name of the driver file.
- Task Name: Name of the task performed.

The Phase Information details include:

- Phase: Database processing phase.

- Start Time: Date and time the phase began.
- Elapsed Time: Time required to complete the phase.
- Jobs: Number of jobs in the phase.
- Total Job Time: Time required to complete the jobs within the phase.
- Restarted?: States whether any jobs within the phase was restarted.
- Skipped: Number of jobs within the phase that were skipped.

Product Information

Clicking the Product Info button displays timing information for all products, with the aggregate timing information for database tasks being shown in the Total Job Time column for each product.

Overall Product Summary

The screenshot shows the Oracle Applications Manager interface. At the top, there is a blue header with the Oracle logo and 'Applications Manager'. Below the header, there are navigation links: 'Support Cart', 'Setup', 'Home', 'Logout', and 'Help'. The main content area has a breadcrumb trail: 'Applications Dashboard > Site Map > Applications System:r1206 > Applied Patches > Patch Details > AutoPatch Timing Details > Job Timing Report >'. Below the breadcrumb, it says 'Last Updated : 05-Apr-2010 10:30:56' and 'Driver File u6688850.drv'. The main table has a task name 'Running parallel SQL and EXEC commands'. The table has four columns: Product, Phase, Jobs, and Total Job Time. There are three rows of data: fnd (9 sec), ad (0 sec), and ad (1 sec). At the bottom right, there is a button labeled 'Add to Support Cart'.

Product	Phase	Jobs	Total Job Time
fnd	daa+51	1	9 sec
ad	plb	1	0 sec
ad	pls	1	1 sec

Clicking the link for an entry in the Product column displays full status details for a specific product. The example shown below is for FND.

Specific Product Details

ORACLE Applications Manager

Support Cart Setup Home Logout Help

Applications Dashboard | Site Map

Products Licensed with Base Version 12.0.0: r1206
Last Updated: 05-Apr-2010 10:31:59

Summary

Status	Count
Licensed	196
Shared	5

List of Products

Filter: Product Abbreviation is fnd
If searching by 'status', please input 'licensed', 'shared' or 'not licensed' as the search keyword.

Select a Product and View... Patch Information

Select	Product Abbreviation	Product Name	Patch set	Status
<input type="radio"/>	FND	Application Object Library	R12.FND.B.1	Licensed

Exceptions

Clicking the Exceptions number in the Job Timing report opens the Exception report. This report is available only for jobs that have an Exceptions value greater than zero in the Job Timing report. It provides a list of exceptions encountered during the maintenance session.

Exception Report

ORACLE Applications Manager

Support Cart Setup Home Logout Help

Applications Dashboard | Site Map

Applications System: ADOACC04 > Timing Reports > AutoPatch Timing Details > Job Timing Report >

Exception Report : Running parallel SQL and EXEC commands : ADOACC04
Last Updated : 10-Aug-2006 19:55:19

Driver File: **u4440000.drv** Jobs failed, restarted manually, then run successfully: **0**

Task Name: **Running parallel SQL and EXEC commands** Jobs failed and skipped: **3**

Phase	Status	Product	Directory	File	Run Time	Restarted?
last+1	Skipped	java	oracle/apps/xdo/oa/util	XLIFFLoader.class	7 sec	N
last+1	Skipped	java	oracle/apps/xdo/oa/util	XLIFFLoader.class	7 sec	N
last+1	Skipped	java	oracle/apps/xdo/oa/util	XLIFFLoader.class	8 sec	N

TIP Skipped : Jobs failed and skipped
Restarted : Jobs failed, restarted automatically, then run successfully

Support Cart Setup Home Logout Help

Copyright 2001, 2006 Oracle Corporation. All Rights Reserved.

The general information includes:

- Driver File: Name of the driver file being run when the exception occurred.
- Task Name: Task being performed when the exception occurred.

- Jobs Failed, then restarted successfully: Number of jobs that initially failed but were restarted successfully.
- Jobs Failed and skipped: Number of failed jobs that were skipped.

The Exception details include:

- Phase: Database processing phase.
- Status: Status of the exception.
- Product: Owner of the file with the exception.
- Directory: Location of the file.
- File: File being processed when the exception occurred.
- Run Time: Total time the process ran.
- Restarted?: States whether the job with the exception was restarted.

View Log Files

Clicking the Log Files icon of a selected row opens the View Log Files page. This page lists all the log files generated for a specific maintenance session. You can view log files of completed jobs or jobs that are in progress.

View Log Files Page

ORACLE Applications Manager

Support Cart Setup Home Logout Help

Applications Dashboard Site Map

Applications System: ADOAGC04 >

View Log Files : ADOAGC04

Last Updated : 08-Jun-2006 20:38:12

Indicates required field.

Task Name: **AutoPatch - u4502962.drv** Start Date: **29-May-2006 09:18:43**
Status: **Completed** Last Updated: **30-May-2006 03:47:00**
Log Directory: Run Time: **18 hr, 28 min**

Log Files

Select a Log File and...

Select Log File	Log Type
<input checked="" type="radio"/> adpatch.log	Main Log File
<input type="radio"/> adpatch.lgi	Other Log File
<input type="radio"/> adlibout.log	Other Log File
<input type="radio"/> adlibin.log	Other Log File
<input type="radio"/> adrelink.log	Other Log File
<input type="radio"/> adwork001.log	Worker Log File
<input type="radio"/> adwork002.log	Worker Log File
<input type="radio"/> adwork003.log	Worker Log File
<input type="radio"/> adwork004.log	Worker Log File
<input type="radio"/> adwork005.log	Worker Log File
<input type="radio"/> adwork006.log	Worker Log File

The View Log Files page contains the following information:

- **Task Name:** Name of the task for which these log files have been generated.
- **Status:** Status of the task. Valid status types are: In-progress, Completed, Stopped, or Aborted.
- **Log Directory:** Location of the listed log files. It is defined by the user when the task is run. This is a required field.
- **Start Date:** Date and time the task was started.
- **Last Updated:** Date and time the task was completed.
- **Run Time:** Time required to complete the task.

The following buttons apply to the Log Files section:

- **View:** Use this button to view a log file after you have selected the radio button of the corresponding log file.
- **Download:** Use this button to download a log file after you have selected the radio button of the corresponding log file.

The Log Files section contains the following information:

- **Select:** Use this radio button to select the corresponding log file for viewing or downloading.

- Log File: Name of the log file.
- Log Type: Log files can be one of these types:
 - Main - primary log files for patch applications and AD Administration tasks.
 - Worker - generated by the processes that run in parallel, using a number of workers.
 - Other - informational files and log files created by the copy portion of a universal (u) patch driver, for example.

View Log Details

Details of each log file are displayed on the View Log Details page. To access this page, select the radio button next to one of the log files and click the View button. The page displays up to 500 lines. It automatically defaults to the last page in the log file.

View Log Details Page

The View Log Details page contains the following information:

- Number of lines per page: Use this to specify the number of lines per page you want to display. The maximum number of lines per page you can display is 500.
- View specific page number: Use this to view a specific page. Enter the page number in the field and click the Go button.

You can use the following buttons to navigate to specific portions of the log file:

- First: Go to the first page.
- Previous: Go to the previous page.
- Next: Go to the next page.
- Last: Go to the last page.

Contents of ...

This section displays the contents of the log file. Click the Go to bottom link to navigate to the last page of the log file. From the last page of the file, click the Return to top link to go back to the first page.

Software Updates

Software Updates is a portal from which you can view all the patching-related activities of your system. From the Software Updates main page, you can access information such as:

- patches that have or have not been applied
- latest three patch recommendation requests from the Patch Wizard page
- latest eight jobs run from the Timing Reports page
- links to patching related pages

The Software Updates Interface

The Software Updates page is a Web-based utility in Oracle Applications Manager. From this page, you can get an overview of all patching-related information.

Accessing Software Updates

To view patching-related activities for your system, log in to Oracle Applications Manager and click the Software Updates tab.

Step 1: Log in to Oracle Applications Manager

Follow the instructions in Accessing Patch Wizard, page 2-2 to access OAM.

Step 2: Click on the Software Updates tab

From the Applications Dashboard, click the Software Updates tab.

Software Updates Page

This section describes the Software Updates page.

Software Updates Page

Applied Patch Check

Use this function to determine if a patch has been or has not been applied.

Patch

(Enter Patch numbers, separated by commas)

Applied	Not Applied

Patch Recommendation Requests

Filter Name / Patch List	Total	Unapplied	Status	Details
No records found.				

Maintenance Activities

Task Name	Status	Last Updated	Details	Log Files
AutoPatch - u6655114.drv	✓	19-Dec-2007 18:22:56		

The top of the page indicates the version of your Oracle E-Business Suite system.

Applied Patch Check

Use the Applied Patch Check field to check if a patch or a series of patches have been applied to your system. Enter a patch ID or a series of IDs separated by commas to check if the patch or patches have been applied to your system. The table below the field contains two columns: Applied and Not Applied. Your queried patch ID appears in the corresponding column depending on whether it has been applied or not.

Patch Recommendation Requests

This section lists the latest three patch recommendation requests from the Patch Wizard main page. From this section, you can view the following information: Filter Name/Patch List, Total (Applied and Unapplied), Unapplied, Status, and Details.

Click the Full List button in the Patch Recommendation Results section to go directly to the Patch Wizard main page. Click the Details icon to go directly to the Recommended Patches Results of the associated patch. See: Patch Wizard, page 2-2.

Maintenance Activities

This section lists the latest eight jobs from the Timing Reports page. From this section, you can view the following information: Task Name, Status, Last Updated, Details, and Log Files.

Click the Full List button in the Maintenance Activities section to go directly to the Timing Reports main page. Click the Details icon to go directly to the Timing Details report of the associated task, or click the Log Files icon to go directly to the View Log Files page of the associated task. See: Timing Reports, page 5-15.

Related Links

This section lists the links related to patching activities for your system.

Part 2

General Maintenance

Basic DBA Tasks

Overview of Oracle E-Business Suite DBA Duties

Oracle E-Business Suite database administration (DBA) combines the efforts of an Oracle E-Business Suite System Administrator and an ORACLE database administrator.

ORACLE Schemas

Upon installation of Oracle E-Business Suite, a number of schemas (sometimes called ORACLE schemas) are present in the database. You do not need to create these schemas; however, you should change the default passwords.

These schemas come from different sources and can be described as being of the following types:

1. Schemas that exist in every Oracle database (whether used by Oracle E-Business Suite or not) [for example, SYS, SYSTEM].
2. A small set of schemas used by shared components of Oracle E-Business Suite (for example, APPLSYSPUB, APPLSYS, APPS).
3. A large set of schemas provided by the individual products of Oracle E-Business Suite (for example, ABM, AHL,..., ZSA, ZX).
4. A set of schemas that belong to optional database features or third party products these fall into three subtypes:
 1. Used by and patched with Oracle E-Business Suite (for example, CTXSYS).
 2. Used by Oracle E-Business Suite but patched only with the RDBMS (for example, MDSYS, ORDSYS).
 3. Not used by Oracle E-Business Suite (for example, SCOTT).

At no time do any of the schemas provided with Oracle E-Business Suite relate to a particular Oracle E-Business Suite user.

All types of schemas are used during runtime operations of Oracle E-Business Suite and the schemas of type 2, 3 & 4.1 are accessed during initial installation and patching.

For schema passwords, Oracle E-Business Suite concerns itself with mainly three passwords for its schemas:

1. A password for APPLSYSPUB (also known as the GATEWAY user). The default password is 'PUB'.
2. A password shared between APPLSYS and APPS (also known as FNDNAM). The default password is 'APPS'.
3. A password for all of the product-specific base schemas (type 3). The default password for these schemas is same as the schema name.

Important: You should change these passwords upon installation.

Note that the Oracle database schemas and passwords connect to the ORACLE database, while application usernames and passwords access Oracle E-Business Suite.

Registering an ORACLE Schema

The installation process automatically registers Oracle E-Business Suite ORACLE schemas, so you only need to register any additional ORACLE schemas that you need using the ORACLE Users window.

You must register an ORACLE schema with Oracle E-Business Suite if:

- you create a custom application using Oracle Application Object Library
- you want to associate an additional ORACLE schema with an Oracle E-Business Suite product

Important: Before you can register an ORACLE schema, your database administrator must first create an ORACLE schema that connects to the ORACLE database. You then use the ORACLE Users window to register your ORACLE schema.

Reregistering ORACLE schemas

You should also reregister ORACLE schemas associated with custom applications built using Oracle Application Object Library each time you upgrade Oracle Application Object Library

Initialization Code

You can add in custom initialization SQL code to be executed when a database session starts up or when it is re-initialized. You specify this code using a profile option.

The code is executed by FND_GLOBAL.INITIALIZE and APPS_INITIALIZE immediately after initializing global variables, profiles, and the contents of client_info on session startup.

Profile Option Initialization SQL Statement - Custom

Using the profile option Initialization SQL Statement - Custom, you can add site-specific initialization code, such as optimizer settings. This profile value must be a valid SQL statement, or a PL/SQL block for more than one statement, that is to be executed once at the startup of every database session.

This profile option can be set at any level by the System Administrator, and is reserved for use by customers.

Profile Option Initialization SQL Statement - Oracle

This profile option is used by Oracle E-Business Suite. This profile option and its value settings are delivered as seed data, and must not be modified.

Related Topics

ORACLE Users (Schemas) Window, page 6-15

Applications Window, page 6-18

Oracle E-Business Suite and Cost-Based Optimization, page 19-1

Overview of Oracle E-Business Suite DBA Duties, page 6-1

ORACLE Users Window, page 6-15

Resource Consumer Groups in Oracle E-Business Suite

The Database Resource Manager introduced in Oracle8i is used to allocate and manage resources among database users and applications.

Resource consumer groups and resource plans provide a method for specifying how to partition processing resources among different users. A resource consumer group defines a set of users who have similar resource usage requirements. An overall resource plan specifies how resources are distributed among the different resource consumer groups.

Oracle E-Business Suite allows the system administrator to assign individual Oracle E-Business Suite users to resource consumer groups. In addition, concurrent programs and concurrent managers can be assigned to resource consumer groups.

Note: These resource consumer groups apply to CPU resources only.

For additional information, see the Oracle database documentation.

Assigning Resource Consumer Groups

The system administrator can assign a user to a resource consumer group by setting the value of the user profile option FND:Resource Consumer Group for that particular user. The user can see this profile option but cannot update it.

The system administrator can assign a concurrent program to a resource consumer group in the Parameters window of the Define Concurrent Program form. See: Concurrent Programs Parameters Window, *Oracle E-Business Suite Setup Guide*.

The system administrator can assign a concurrent manager to a resource consumer group in the Define Concurrent Manager form. See: Concurrent Managers Window, *Oracle E-Business Suite Setup Guide*.

Hierarchy of Resource Consumer Group Assignments

Conflicts can arise between the resource consumer groups associated with a single session. For example, a concurrent manager assigned to one resource consumer group may run a concurrent program assigned to another. A similar situation arises when a user performs a transaction managed by a transaction manager that has a different resource consumer group than the user. To resolve such conflicts, Oracle E-Business Suite uses a hierarchy.

In the case of a concurrent program, the system first checks to see if the program has an assigned resource consumer group and if so, uses that. If not, the system checks the concurrent manager running the program and uses its resource consumer group. If the concurrent manager is not assigned to a resource consumer group the system uses the default group "Default_Consumer_Group".

In the case of a transaction manager running a transaction program, the system once again checks the resource consumer group assigned to the program, if any, and if there is none, checks the transaction manager. If the transaction manager has no assigned resource consumer group the system then checks the profile option value for the user whose session began the transaction. If there is no resource consumer group defined the system uses the default resource consumer group.

For a user running a form, the system first checks the profile option value for that user and uses that if it is defined. Otherwise the system uses the default resource consumer group.

Oracle E-Business Suite Password Management

Oracle E-Business Suite provides two command line utilities, FNDCPASS and

AFPASSWD, for setting Oracle E-Business Suite schema passwords. These utilities change the password registered in Oracle E-Business Suite tables and the schema password in the database. The utilities can also be used to change user passwords.

Several important considerations apply:

- The Oracle E-Business Suite system must be shut down before a schema password is changed.
- If Oracle E-Business Suite user passwords are being changed, the relevant users should not be logged in.
- The FND_USER and FND_ORACLE_USERID tables should be backed up before any passwords are changed.

Oracle E-Business Suite releases that use an Oracle 11g database may optionally employ case-sensitive database passwords, allowing mixed case to be used. How to utilize this feature is described later in this section.

FNDCPASS Utility

The FNDCPASS utility can be used to change various types of passwords.

To change the APPS and APPLSYS (type 2) schema password:

Use this command to change passwords for schemas that are used by shared components of Oracle E-Business Suite.

```
FNDCPASS <logon> 0 Y <SYSTEM username>/<SYSTEM password> SYSTEM \  
<username> <new_password>
```

Use the above command with the following arguments. When specifying the SYSTEM token, FNDCPASS expects the next arguments to be the APPLSYS username and the new password.

logon The Oracle username.

Note: You can provide just the Oracle username, and the system will prompt you for the password. Alternatively you can provide the <username>/<password> pair.

system/password The username and password for the SYSTEM DBA account.

username The APPLSYS username. For example, 'applsys'.

new_password The new password.

This command does the following:

1. Validates APPLSYS.
2. Re-registers password in Oracle E-Business Suite.
3. Changes the APPLSYS and all APPS passwords (for multi-APPS schema installations) to the same password.

Because everything with a Privilege Level [set to any of ('E', 'U', 'D')] in the FND_ORACLE_USERID table must always have the same password, FNDCPASS updates these passwords as well as APPLSYS's password. For example, the APPS password will be updated when the APPLSYS password is changed.

4. ALTER USER is executed to change the ORACLE password for the above ORACLE users.

For example, the following command changes the APPLSYS password to 'WELCOME'.

```
FNDCPASS <APPS username> 0 Y <SYSTEM username>/<SYSTEM password> SYSTEM
APPLSYS WELCOME
```

Important Additional Instructions to Update WLS Datastore

Whenever you use FNDCPASS to change the APPLSYS password, you must also perform the following actions.

Important: These steps must be carried out on the *run* file system.

1. Shut down the application tier services using the `$INST_TOP/admin/scripts/adstpall.sh` script.
2. Change the APPLSYS password, as described in the previous section.
3. Start AdminServer using the `$INST_TOP/admin/scripts/adadminsrvctl.sh` script. Do not start any other application tier services.
4. Change the "apps" password in WLS Datasource as follows:
 1. Log in to WLS Administration Console.
 2. Click *Lock & Edit* in Change Center.
 3. In the Domain Structure tree, expand Services, then select Data Sources.
 4. On the "Summary of JDBC Data Sources" page, select EBSDataSource.
 5. On the "Settings for EBSDataSource" page, select the Connection Pool tab.
 6. Enter the new password in the "Password" field.

Note: You can provide just the Oracle username, and the system will prompt you for the password. Alternatively you can provide the <username>/<password> pair.

<SYSTEM username>/<SYSTEM password>	The username and password for the SYSTEM DBA account.
username	The Oracle username. For example, 'GL'.
new_password	The new password.

For example, the following command changes the GL user password to 'GL1'.

```
FNDCPASS <APPS username> 0 Y <SYSTEM username>/<SYSTEM password> ORACLE  
GL GL1
```

Note: The system prompts the user for the APPS password in the above example.

To change all ORACLE (type 3) schema passwords:

Use this command to change the passwords of all schemas provided by Oracle Application products.

```
FNDCPASS <logon> 0 Y <SYSTEM username>/<SYSTEM password> ALLORACLE \  
<new_password>
```

Use the above command with the following arguments. When specifying the ALLORACLE token, FNDCPASS expects the next argument to be the new password.

logon	The Oracle username.
--------------	----------------------

Note: You can provide just the Oracle username, and the system will prompt you for the password. Alternatively you can provide the <username>/<password> pair.

<SYSTEM username>/<SYSTEM password>	The username and password for the SYSTEM DBA account.
new_password	The new password.

For example, the following command changes all ORACLE schema passwords to "WELCOME":

```
FNDCPASS <APPS username> 0 Y <SYSTEM username>/<SYSTEM password>  
ALLORACLE WELCOME
```


Note: The system prompts the user for the APPS password in the above example.

To change an Oracle E-Business Suite user's password:

You can use this command to change an individual Oracle E-Business Suite user's password.

```
FNDCPASS <logon> 0 Y <SYSTEM username>/<SYSTEM password> USER \  
<username> <new_password>
```

Use the above command with the following arguments. When specifying the USER token, FNDCPASS expects the next arguments to be an Oracle E-Business Suite username and the new password.

logon The Oracle username.

Note: You can provide just the Oracle username, and the system will prompt you for the password. Alternatively you can provide the <username>/<password> pair.

<SYSTEM username>/<SYSTEM password> The username and password for the System DBA account.

username The Oracle E-Business Suite username. For example, 'VISION'.

new_password The new password.

For example, if you were changing the password for the user VISION to 'WELCOME', you would use the following command:

```
FNDCPASS <APPS username> 0 Y <SYSTEM username>/<SYSTEM password> USER  
VISION WELCOME
```

Note: The system prompts the user for the APPS password in the above example.

FNDCPASS prompting for password if not provided

FNDCPASS prompts the user for the APPS user password if it is not given on the command line.

Here is an example in which the APPS password provided on the command line:

```
FNDCPASS APPS/APPS 0 Y <SYSTEM username>/<SYSTEM password> USER  
operations welcome
```

You can choose not to give the APPS password in the same command, as in the following example.

```
FNDCPASS APPS 0 Y <SYSTEM username>/<SYSTEM password> USER operations
welcome
ORACLE Password:
```

Here the APPS password is not provided on the command line, but instead you are prompted for it.

FNDCPASS Example

The following is an example of using the FNDCPASS utility to change the password, where <username> is the Oracle schema name.

```
FNDCPASS <APPS username>/<APPS password> 0 Y \
<SYSTEM username>/<SYSTEM password> \
ORACLE <username> <new_password>
```

Changing the APPLSYS and APPS passwords is a special case. Here, all application tier services must first be shut down using the command `$INST_TOP/admin/scripts/adstpall.sh`. In the command, ORACLE must be replaced with SYSTEM: SYSTEM mode changes both the APPLSYS and APPS passwords, and thereby keeps them synchronized.

Note: FNDCPASS requires the next argument after SYSTEM to be the APPLSYS username.

```
FNDCPASS <APPS username>/<APPS password> 0 Y \
<SYSTEM username>/<SYSTEM password> \
SYSTEM APPLSYS <new_password>
```

You will then need to run AutoConfig (`adautoconfig.sh`) using <new_password> as the APPS password, and finally restart application tier services using the command `$INST_TOP/admin/scripts/adstrtal.sh`.

Tip: For assistance in resolving any issues, refer to My Oracle Support Knowledge Document 1306938.1, *FNDCPASS Troubleshooting Guide For Login and Changing Applications Passwords*.

AFPASSWD Utility

AFPASSWD is an enhanced version of FNDCPASS, and includes the following features:

- AFPASSWD only prompts for passwords required for the current operation, allowing separation of duties between applications administrators and database administrators. This also improves interoperability with Oracle Database Vault. In contrast, the FNDCPASS utility currently requires specification of the APPS and the SYSTEM usernames and corresponding passwords, preventing separation of duties between applications administrators and database administrators.

- When changing a password with AFPASSWD, the user is prompted to enter the new password twice to confirm.

FNDCPASS will continue to be shipped with Oracle E-Business Suite, and customers can migrate to the AFPASSWD utility at their discretion.

Important: The FNDCPASS utility must still be used to migrate the password hashing scheme, as described in My Oracle Support Document 457166.1, *FNDCPASS Utility New Feature: Enhance Security With Non-Reversible Hash Password*.

AFPASSWD Usage

The AFPASSWD command is used with the relevant command line options to perform the desired action.

```
AFPASSWD [-c <APPSUSER>[@<TWO_TASK>]] [-f <FNDUSER>]
```

```
AFPASSWD [-c <APPSUSER>[@<TWO_TASK>]] [-o <DBUSER>]
```

```
AFPASSWD [-c <APPSUSER>[@<TWO_TASK>]] [-a]
```

```
AFPASSWD [-c <APPSUSER>[@<TWO_TASK>]] [-l <ORACLEUSER> [<TRUE>] | [<FALSE>]]
```

```
AFPASSWD [-c <APPSUSER>[@<TWO_TASK>]] [-L [<TRUE>] | [<FALSE>]]
```

```
AFPASSWD [-c <APPSUSER>[@<TWO_TASK>]] [-s <APPLSYS>]
```

These options have the following functions:

- **-c {APPSUSER}[@{TWO_TASK}]** - Specifies the connection string to use, the Applications user, and/or the value of TWO_TASK. This option can be use in combination with others. If it is not specified, default values from the environment will be used.

Note: The password will be prompted for, and is not to be provided in the connection string.

- **-f {FNDUSER}** - Changes the password for an Applications user. A username that contains spaces must be enclosed in double quotation marks; for example, "JOHN SMITH".
- **-o {DBUSER}** - Changes the password for an Oracle E-Business Suite database user.

Note: This only applies to users listed in the FND_ORACLE_USERID table, not database users in general.

- **-a** - Changes all Oracle (ALLORACLE) passwords (except the passwords of APPS, APPLSYS, APPLSYSPUB) to the same password, in the same way as the ALLORACLE mode does in FNDCPASS.
- **-l** - Locks individual {ORACLE_USER} users (except required schemas). {TRUE} = LOCK, {FALSE} = UNLOCK.
- **-L** - Locks all Oracle (ALLORACLE) users (except required schemas). {TRUE} = LOCK, {FALSE} = UNLOCK.
- **-s {APPLSYS}** - Changes the password for the APPLSYS user and the APPS user. This requires the execution of AutoConfig on all tiers.
- **-h** - Displays help.

Using Case-Sensitive Database Passwords

To help meet increasing and often mandatory requirements for complex passwords, Oracle E-Business Suite now supports the use of Oracle Database 11g case-sensitive passwords. This is in contrast to the traditional Oracle Application Library behavior of storing and validating all database passwords as uppercase, regardless of the case in which they are entered.

Case-sensitive database passwords can be employed with any Oracle E-Business release that uses an Oracle 11g database. Using mixed case enables more secure application schema passwords to be specified.

Overview

There are two possible situations:

Case sensitivity disabled (default) - For new database accounts or changed database passwords, Oracle automatically records the case in which the password was originally specified and stores it as a hash value in the data dictionary table that holds user information. However, new or changed database account passwords will continue to be case-insensitive unless and until the mixed-case feature is explicitly enabled.

Case sensitivity enabled - After the feature is enabled, database passwords created or changed since the upgrade to Oracle 11g will need to be entered in the case specified originally. Only database passwords that remain unchanged in the Oracle 11g database will continue to be case-insensitive. The database stores a case-sensitive version of the password created or changed in Oracle 11g, whether the mixed-case feature is enabled or not. The case-sensitive version of the password is therefore ready for immediate use as soon as the feature is enabled.

The case sensitivity capability for Oracle E-Business Suite database passwords is analogous to the way the SIGNON_PASSWORD_CASE profile is used to determine how new or changed Oracle E-Business Suite user passwords will be stored.

Enabling Case-Sensitive Database Passwords

Case sensitivity is controlled by the setting of the Oracle 11g database initialization parameter `SEC_CASE_SENSITIVE_LOGON`. The default for Oracle E-Business Suite databases is `FALSE`, which means that new, existing (pre-11g), and changed database passwords will all remain case-insensitive.

The case sensitivity feature is enabled as follows:

1. Set the Oracle 11g Database initialization parameter `SEC_CASE_SENSITIVE_LOGON` to `TRUE`.

Note: The default for Oracle E-Business Suite databases is `FALSE`, which means that new, existing (pre-11g), and changed database passwords will remain case-insensitive.

2. Shut down and restart the database. New and changed database passwords will now be case-sensitive, with the case of the password being preserved as it was entered.

Bear in mind the following important points:

- As mentioned above, existing (pre-11g) database passwords will remain case-insensitive until changed, after which they will become case-sensitive.
- 11g passwords that were changed prior to the activation of case sensitivity will be case-sensitive after conversion, but they must be entered in upper case as this is the way they were stored. This includes the `APPLSYS` and `APPS` passwords.
- After conversion, it is advisable to run `FNDCPASS` or `AFPASSWD` to change the `APPLSYS` and `APPS` passwords and any other Oracle E-Business Suite database passwords. This will ensure that the database passwords are stored in the expected case.

Warning: The first time `FNDCPASS` or `AFPASSWD` are run after conversion to use case-sensitive passwords, the `APPS` password must be entered in uppercase or the connection attempt will fail. After three failed attempts, the `APPS` account will be locked as per the default behavior of the 11g database user profile.

- Once the `APPS` password is changed, the case in which the password was entered will be the case of the password (unless it is changed again). An `APPS` password that is hardcoded (for example in a script) may no longer work after conversion, until updated to the new case-sensitive password.

3. After password case sensitivity has been enabled, a DBA should immediately change the passwords of administrative accounts such as SYS and SYSTEM, and may also wish to employ a password management policy (profile) to ensure system administrators change the Oracle E-Business Suite database passwords within a reasonable time.

If for some reason password case sensitivity needs to be disabled once again, this can be done by setting `SEC_CASE_SENSITIVE_LOGON = FALSE` and shutting down and restarting the database.

Warning: In such a case, all Oracle E-Business Suite database passwords (stored in the `FND_ORACLE_USERID` table) must be reset, or database authentication will fail. Disabling password case sensitivity is therefore not recommended.

Additional Considerations for Oracle E-Business Suite DBAs

Be aware of the following points:


- We recommend that the `APPLSYSPUB` password should be changed on all Release 12.x systems, using either via `AFPASSWD` or `FNDCPASS`. `AutoConfig` should be run after changing the password, to synchronize all the application tier files.

Important: The `APPLSYSPUB` password is an exception to the standardization of mixed case passwords, and must always be in upper case. This is true even if case-sensitive passwords have been enabled in your Oracle11g database (`SEC_CASE_SENSITIVE_LOGON=true`).

For more information, refer to My Oracle Support Knowledge Document 403537.1, *Secure Configuration Guide for Oracle E-Business Suite Release 12*.

- Passwords with special characters or multibyte characters are not currently supported with Oracle E-Business Suite.

ORACLE Users Window



Use this window to register an ORACLE username with Oracle E-Business Suite. An ORACLE username grants access privileges to the ORACLE database.

Before registering an ORACLE username, create the ORACLE username that matches your application needs (this function is usually performed by a database administrator). The ORACLE username must include the *create session* privilege.

The installation process always registers your ORACLE username, so you need not register it unless you create a custom application using Oracle Application Object Library, or if you wish to associate an additional ORACLE username with Oracle E-Business Suite.

If you register an ORACLE username as a "restricted" ORACLE username, you submit a concurrent request to set up read-only privileges to the Oracle Application Object Library tables. An "enabled" ORACLE username has all privileges to those tables. A "disabled" ORACLE username has no privileges to those tables.

If you do not register and enable your ORACLE username or if you disable a registered ORACLE username, your user cannot use Oracle Application Object Library features such as menus and flexfields.

You should not change the registration of any ORACLE usernames that the installation process registers, other than changing the passwords.

If you are registering a change to an existing ORACLE password, make the password change in the database immediately AFTER you register the password change in Oracle E-Business Suite. Until you register the password changes in Oracle E-Business Suite

and implement them in the database, responsibilities using this ORACLE username cannot connect to the database.

Your password must follow the guidelines for creating passwords discussed in the Oracle database documentation. Remember that if you use non-character values in your password, you may need to use quotation marks around your password when changing it in the database.

Note: Use FNDCPASS to change the password, not the ORACLE Users window>

Warning: If you are changing the password to the *appls* ORACLE username, which contains the Oracle Application Object Library tables, you must *not* change the passwords to any other ORACLE usernames at the same time.

As soon as you change and save the password, you should immediately log out of the Oracle E-Business Suite, make the *appls* password change in the database, and then sign on again before you do anything else. You should also ensure that no other users are logged on to the Oracle E-Business Suite while you are changing the *appls* password.

Important: For passwords for the APPS accounts, the *appls* password must be identical to the password for the APPS accounts (APPS, APPS2, APPS3). The uniform passwords enable the different sets of books to operate correctly.

ORACLE Users Block

This block contains the following:

Password

Enter the password of your ORACLE username. Your password is not displayed. If you are registering a change to an existing ORACLE password, make the password change in the database immediately AFTER you register the password change in Oracle E-Business Suite.

Until you register the password changes in Oracle E-Business Suite and implement them in the database, responsibilities using this ORACLE username cannot connect to the database.

Warning: If you are changing the password to the *appls* ORACLE

username, which contains the Oracle Application Object Library tables), you must *not* change the passwords to any other ORACLE usernames at the same time.

As soon as you change and save the password, you should immediately log out of the Oracle E-Business Suite, make the *appls* password change in the database, and then sign on again before you do anything else. You should also ensure that no other users are logged on to the Oracle E-Business Suite while you are changing the *appls* password.

Using FNDCPASS is the preferred method for changing passwords. Note that FNDCPASS automatically updates all instances of the password for you.

Privilege

Enter the type of privilege to the Oracle Application Object Library database tables that you want this ORACLE username to have. The Oracle Application Object Library tables contain information for Oracle Application Object Library features such as menus, help text, and flexfields. If you do not have access to these tables, you cannot use these features.

The default value for this field is Enabled.

- Enabled - An enabled ORACLE username has full privileges (insert, query, update, and delete) to the Oracle Application Object Library database tables.
- Restricted - A restricted ORACLE username has only query privileges to the Oracle Application Object Library database tables. This ORACLE username can view Oracle Application Object Library data, but cannot insert, update, or delete information.
- Disabled - A disabled ORACLE username has no privileges to the Oracle Application Object Library database tables. This ORACLE username cannot insert, query, update, or delete Oracle Application Object Library information and cannot use Oracle Application Object Library features.

Two additional privilege types appear, associated with ORACLE usernames configured at installation. However, these privilege types cannot be selected from your list of values.


- Public - The installation process registered an ORACLE username with the Public privilege, allowing all users to access the Application Sign-On Security form where they must enter a valid Oracle E-Business Suite username and password.
- Appls - The installation process registered the Oracle Application Object Library ORACLE username with the Appls privilege.

Install Group

The value of the installation group associated with your ORACLE username. Install group numbers should be consecutive whole numbers, where 1 represents the first set of books (or first set of product installations), 2 is the second set of books, 3 is the third set of books, and so on. Install group number 0 represents products that need only single installations.

Important: Since the installation process does not affect ORACLE usernames (also known as "schemas") for custom applications, this value is for your reference only and is currently not used.

Applications Window



Use this window only if you are creating a custom application to isolate custom code and/or data from shipped Oracle E-Business Suite applications.

When you define a custom application, you supply several pieces of information to Oracle E-Business Suite. You must register your application name, application short name, application basepath, and application description with Oracle Application Object Library. Oracle Application Object Library uses this information to identify application objects such as responsibilities and forms as belonging to your application. This identification with your custom application allows Oracle E-Business Suite to preserve your application objects and customizations during upgrades. The application basepath tells Oracle Application Object Library where to find the files associated with your

custom application.

You can use your custom application to name your custom menus, concurrent programs, custom responsibilities, and many other custom components. For some objects, the application part of the name only ensures uniqueness across Oracle E-Business Suite. For other components, the application you choose has an effect on the functionality of your custom object.

Prerequisites

- If you are creating a custom application: Define an environment variable that translates to your application's basepath (see *Oracle E-Business Suite Concepts* for your operating system).
- If you are creating a custom application: Set up a directory structure for your application (see *Oracle E-Business Suite Concepts* for your operating system)
- If you are using this window with Oracle Alert: If your application resides in a database other than the database where Oracle Alert resides, you must create a database link.

Applications Block

When you register a custom application, you provide the information Oracle uses to identify it whenever you reference it. Although you can change the application short name of an application, doing so may cause a change in the application code where you hardcode your application short name. For example, if you pass program arguments through the menu that have application name hardcoded, you will also have to update them.

Important: You should not change the name of any application that you did not develop, as you cannot be sure of the consequences. You should never change the name of any Oracle E-Business Suite application, because these applications may contain hardcoded references to the application name.

Application

This user-friendly name appears in lists seen by application users.

Short Name

Oracle E-Business Suite uses the application short name as an internal key; for example, when identifying forms, menus, concurrent programs and other application components. The short name is stored in hidden fields while the name displays for users.

Your short name should not include spaces. You use an application short name when you request a concurrent process from a form, and when you invoke a subroutine from a menu.

Tip: Although your short name can be up to 50 characters, we recommend that you use only four or five characters for ease in maintaining your application and in calling routines that use your short name. To reduce the risk that your custom application short name could conflict with a future Oracle E-Business Suite short name, we recommend that your custom application short name begins with "XX".

Basepath

Enter the name of an environment variable that represents the top directory of your application's directory tree. Oracle E-Business Suite searches specific directories beneath the basepath for your application's files and scripts.

In general, your application's basepath should be unique so that separate applications do not write to the same directories.

However, you may define custom applications that will be used only for naming your custom responsibilities, menus and other data components. In this case, you can use the basepath of the Oracle application that uses the same forms as your application. For example, if you are defining a Custom_GL application, you could use the GL_TOP basepath for your custom application. In this case, however, you should not create custom components in the directory structure, such as custom forms and reports, because they will be difficult to isolate for maintenance and upgrading.

See: *Oracle E-Business Suite Concepts*

Network Test Window

The screenshot shows the Oracle Applications Network Test window. The window title is "Oracle Applications" and the application title is "Network Test". The menu bar includes "File", "Edit", "View", "Folder", "Tools", "Window", and "Help". The toolbar contains various icons for file operations and testing. The main area is divided into several sections:

- Latency:** Input fields for "Trials" (5) and "Iterations" (100).
- Bandwidth:** Input fields for "Trials" (5) and "Iterations" (10).
- Notes:** A text area for entering notes.
- Buttons:** "Clear Old Test Data" and "Run Test".
- Results:** A section with a "Test Date" field and a "Batch" checkbox.
- Latency Results:** A table with columns for "Trial", "Test", "LAN", and "WAN".
- Bandwidth Results:** A table with columns for "Trial", "Test", "LAN", and "WAN".
- Legend:** "LAN = Client at HQ (Redwood Shores) and Server at HQ on LAN." and "WAN = Client at HQ and server in Orlando (T1 line)."
- Status Bar:** "<OSC>" and "<DBG>" indicators.

You can use the Network Test form to monitor the latency and bandwidth of the network for forms applications, or to help create a baseline for use in comparing response times from within the application. This information enables you to make comparisons between locations, or at different times of day at the same location. The form shows the time taken to perform one or more Oracle Forms round trips, and the throughput used.

The latency shown on the form represents a combination of the round trip time needed to traverse the physical network (including any devices), and the Forms overhead to process a packet. The network test form is designed to more closely measure the network latency and bandwidth of an actual forms user. Note that the results are not expected to match the times returned by `ping`, `tracert`, or other diagnostic network commands.

To test the network latency, a short sequence of packets is sent from the client application to the application server, then on to the database server, and back to the client. You need to specify the number of sequences (iterations) you want to send, and the number of times you want to send each set of iterations (trials). The default setting is 5 trials of 100 iterations each. The average latency is the total time for all round trips in a trial, divided by the number of iterations.

The bandwidth test (or more accurately, throughput test), examines the data transfer rate, and shows how many bytes per second your network transferred between the client, application server and database server.

Running a Test

Click the Run Test button to perform the test.

You can provide notes to indicate the conditions for each test you run.

Evaluating the Test Results

If one test result varies significantly from the other trials, discard that information.

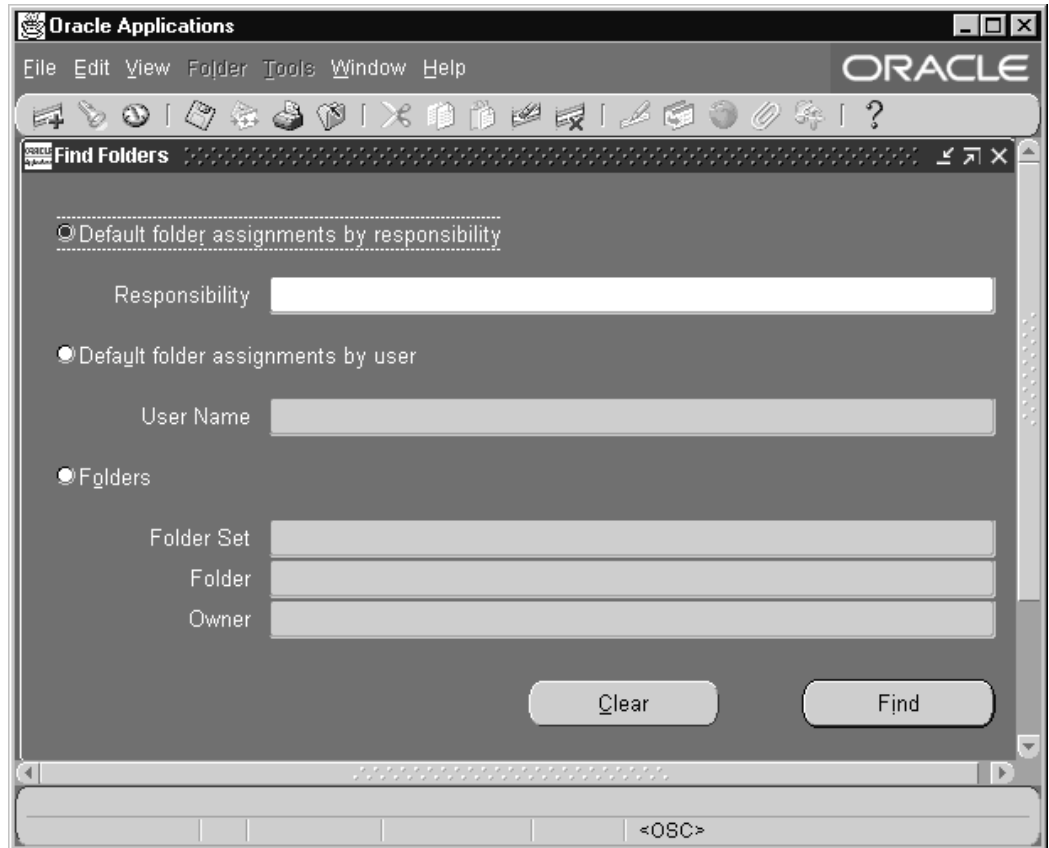
Use the Clear Old Test Data button to purge previous test results from your database.

The results of both the latency and throughput tests are displayed in the Results block.

- *Latency Results* displays the minimum, average, and maximum round trip time for a single round trip from a PC client to the server.
- The *Bandwidth* window shows the throughput results, and displays the minimum, average, and maximum data rate in bytes per second.

For comparison, the sample data fields show the results of tests completed at Oracle Headquarters in Redwood Shores, California.

Administering Folders



Administer folders by assigning default folder definitions either to a specific user or to a responsibility. Manage folder definitions by assigning them to new owners, determining which folder definitions should be public (accessible by anyone), and setting the AutoQuery behavior of the folders.

You can do different tasks depending on how you search for folders or folder assignments in the Find Default Folders window.

You can assign a default folder to a user or responsibility in "restricted mode" such that all folder functionality is disabled at runtime for the user. In this way you can, for example, prevent users from seeing specific fields, or control which records they can query. This behavior is controlled by the Behavior Mode poplist value, set when the folder is assigned.

You must have default folders before you perform these steps.

To Assign a Folder to a Responsibility:

Follow these steps to assign a folder to a responsibility:

1. Navigate to the Find Default Folders window. Use "Default folder assignments by responsibility" to view the responsibilities for which to assign default folders.
2. You can assign default folders for each responsibility. When users of this responsibility navigate to this folder block, they see the default folder you specify, unless it is overridden by a user-level default.

From the Folder field, enter the name of the default folder. The name of the folder set to which the folder belongs is filled in automatically.

If you do not know the name of the folder, enter the folder set first, then view the folders that belong to that set.

After you save a default folder definition for a folder set, that folder set no longer appears in the list of values.

Folder Set: Every folder set is associated with a particular folder block, and a user or responsibility can have one default folder within each folder set. The folder set name generally describes the records shown in the block; some blocks may have multiple sets of folders associated with them.

To Assign a Folder to a User:

Follow these steps to assign a folder to a user:

1. Navigate to the Find Default Folders window. Use "Default folder assignments by user" to view a list of eligible users.
2. You can assign default folders for each responsibility. When users navigate to this folder block, they see the default folder you specify.

From the Folder field, enter the name of the default folder. The name of the folder set to which the folder belongs is filled in automatically.

If you do not know the name of the folder, enter the folder set first, then view the folders that belong to that set.

After you save a default folder definition for a folder set, that folder set no longer appears in the list of values.

Folder Set: Every folder set is associated with a particular folder block, and a user or responsibility can have one default folder within each folder set. The folder set name generally describes the records shown in the block; some blocks may have multiple sets of folders associated with them.

Source Type: Either User or Responsibility. Records entered in this window use the source type of User. If one of the current user's responsibilities has default folders defined, the default folders are listed with a source type of Responsibility.

User defaults override Responsibility defaults. You cannot delete Responsibility default folders in this window.

Responsibility: The responsibility which uses this default folder definition.

To Assign Ownership of a Folder:

Follow these steps to assign ownership of a folder:

1. Navigate to the Find Default Folders window. Use "Folders" to view general information about folders.
2. Select the folder(s) that requires a change of ownership.
3. Choose "Change Owner" and enter the new owner for the selected folders, or change the value in the Owner field to change the owner of a single folder.

Folder Set: Every folder set is associated with a particular folder block, and a user or responsibility can have one default folder within each folder set. The folder set name generally describes the records shown in the block; some blocks may have multiple sets of folders associated with them.

Public: Whether this folder definition is public; whether users besides the owner can use it. Use this field to determine whether to make folder definitions generally available.

Anyone's Default: Whether this folder definition is used as a default by a user or a responsibility. If it is a default definition, use Default Assignments to view the users and responsibilities for which it is the default folder definition.

Default Assignments: The users and responsibilities that use this folder definition as a default.

To Delete a Folder Definition

Follow these steps to delete a folder definition:

1. Navigate to the Find Default Folders window. Use "Folders" to view general information about folders.
2. If you queried up multiple folders, select the folder(s) to delete.
3. Delete the folder. Deleting folders deletes the folder definition along with any user and responsibility default assignments for the folder.

To Create and Assign a Folder in "Restricted Mode"

Use the steps below to create and assign a folder in a "restricted mode". When user opens a folder in restricted mode, all folder functionality is disabled.

1. Run the folder form and navigate to the folder block.
Hide or show fields as you wish. Take care in choosing the appropriate fields, as the

fields that are hidden will not be accessible for users or responsibilities of this folder block after it is assigned to them as a default folder in restricted mode.

2. Save the folder.
3. Assign the folder as a default folder to a responsibility or user.
4. Set the value of the new Behavior Mode poplist to "Restrict fields and folder functions".

A default folder can have one of the following values for Behavior Mode:

- No restrictions - End user can perform all folder functions.
- Restrict fields and folder functions - End user cannot perform any folder functions. This is "Restricted Mode".

Runtime Scenarios with Restricted Mode

Here are two scenarios with restricted mode:

End user runs folder form with restrictions

When the user opens the restricted default folder form, all folder functions are disabled. For example, the user cannot open any other folders, or move or resize fields.

Within a folder block, once a restricted default folder loads, all folder functionality will become disabled even if that block supports other folder objects.

System Administrator wants to change the default restricted folder

Once a default folder is assigned with the Behavior Mode "Restrict fields and folder functions" to any user or responsibility, it no longer appears in the list of available folders for opening by any user (even though this folder is defined as "Public").

To change this default folder, you should first assign the default folder to yourself. Then run the folder form and navigate to the folder block so that the default folder will load. You can then make modifications and save the folder. Even though the Behavior Mode is restricted, the folder functions can still be performed since you have become the owner of the folder.

Applications DBA System Maintenance Tasks and Tools

Overview

This chapter describes the various operations you will need to perform during your work as an Oracle E-Business Suite DBA. Depending on your organization's requirements, some of the tasks will need to be performed often, and others rarely or never.

In addition, the frequency with which many of the tasks will need to be performed is very likely to vary over the life cycle of your Oracle E-Business Suite system. For example, you will probably apply patches throughout the life of the system, but only add an NLS language when there is a specific business need.

Choosing the Correct File System For Maintenance Tasks

As described in the Patching part of this book, and in Chapters 2 and 4 of *Oracle E-Business Suite Concepts*, all patching in Release 12.2 is done while the Oracle E-Business Suite system is running and available for use. One of the key concepts involved is the employment of a dual file system: users are connected to the run file system, while patches are, when needed, applied to the patch file system. After all the relevant patches have been applied, the current online patching cycle concludes with the two file systems swapping their identities.

As pre-12.2 Oracle E-Business Suite releases only employed one file system, there was no choice about where maintenance activities had to be carried out from. Use of a dual file system in Release 12.2 has required the introduction of a *Configuration Change Detector*, which automatically detects when changes are made to one file system and replicates them to the other.

With technology stack components, however, you may need to make updates manually. In such cases, the question arises as to which file system should you perform the updates on.

Your choice should be made as follows:

- When a patching cycle is *active*, you should perform administration tasks on the *patch* file system. As part of the cutover phase that concludes a patching cycle, the administrative actions will be propagated to the other file system. So you can apply patches and perform general maintenance activities during the course of a patching cycle.

Important: If for some reason you need to perform administration tasks on the *run* file system while a patching cycle is in progress, you must either complete or abort the patching cycle before starting the administration tasks.

- When a patching cycle is *not active*, you should perform administration tasks on the *run* file system. You must then clone the run file system to the patch file system using the `adop fs_clone` command.

Important: It is strongly recommended that you manually run `fs_clone` on the run file system immediately after the maintenance tasks have been completed. If you fail to do so, `adop` will run it automatically in the prepare phase of the next patching cycle (which will mean that cycle may take longer).

Managing Files

This section contains information about maintenance tasks associated with Oracle E-Business Suite files.

Generating Product Files:

Requirement

I want to generate missing product files.

Discussion

Every Oracle E-Business Suite product contains generated files, such as form, report, message, and JAR (Java archive) files. Run AD Administration when you suspect generated files are missing. For example, if users are not able to use a certain General Ledger form, regenerating the form file may resolve the issue. You may also need to generate files after you license additional products.

Note: You do not have to shut down your system to generate files. However, users that are accessing the files being generated (for example, for Human Resources forms) must log off.

Actions

Perform the following steps:

1. Determine the file types that require generation.
2. Start AD Administration by setting the environment and then entering `adadmin` on the command line.

Note: For more information, see *Setting the Environment in Running AD Utilities*, page 7-52.

3. From the AD Administration Main menu, go to the Generate Applications Files menu and select the task for the type of files you want to generate, based on the following criteria:
 - When you choose one of the options for generating form or report files, you can select an individual file, a set of files, or all files of the selected type.
 - The "Generate product JAR files" option allows you to generate all JAR files for all products, or only JAR files that are out-of-date.

Important: If you are performing a new installation of Oracle E-Business Suite, you must create your own signature, and then force regeneration of all JAR files. This will avoid the occurrence of security warnings, for example when launching forms, that can result from the existence of multiple signatures.

- The "Generate message files" option generates all message files for all products.

Note: For more information, see *Generating Applications Files*, page 7-66.

4. Repeat the generation task on each APPL_TOP that contains the files (if the system contains multiple APPL_TOPs).
5. Review the AD Administration log file for warnings or errors.

Adding New Off-Cycle Products:

Requirement

I want to add a product that was released after the last release update pack.

Discussion

Products that are released in between maintenance releases are sometimes referred to as

off-cycle products. Since these new products do not appear in the OAM License Manager, you must add them to your product list by using AD Splicer. This utility splices the product into the list of existing products that are known to your system. This process makes the product available, so that you can register it as active, and thus, make it available to the AD maintenance utilities, such as AutoPatch.

Once you splice the product, you use AutoPatch to install all product-related files.

Note: For more information, see AD Splicer in this chapter.

Actions

Perform the following steps:

1. Download the initial product patch from My Oracle Support.

This patch contains information about the new product, AD Splicer control files required to add the product, and the associated product files.

2. Review the readme file.

Unzip the patch in the patch top directory. The patch readme file contains information on how to install the product. It may include manual steps to perform as part of this process.

Important: Do *not* apply the patch using AutoPatch.

3. Apply prerequisite patches (if any).

Follow the instructions about prerequisite patches in the patch readme file.

4. Create tablespaces (conditional).

If you initially installed your system with Rapid Install 11.5.10 or later, omit this step.

If your system was upgraded to Release 11.5.10 from a previous version of Release 11*i*, you may have chosen to continue using the OFA tablespace model. If so, create two tablespaces for each product, one for the product tables and another for the product indexes.

Note: For more information, see Tablespace Management in *Oracle E-Business Suite Concepts*.

5. Copy AD Splicer control files and product configuration file.

Copy <prod>prod.txt, <prod>terr.txt, and newprods.txt to APPL_TOP/admin.

Caution: If a newprods.txt already exists from a previous AD Splicer session, rename the existing file before copying the new newprods.txt file. If you need to edit this file, see AD Splicer in this chapter.

6. Add the off-cycle product to the list of products.

Log on as applmgr, set the environment, and run AD Splicer. It modifies the APPL_TOP and database, then performs the same registration function as OAM License Manager.

UNIX:

```
$ cd $APPL_TOP/admin
$ adsplice
```

Windows:

```
C:\>cd %APPL_TOP%\admin
C:\>adsplice
```

Run AD Splicer for each APPL_TOP and database combination so that the Applications utilities recognize the off-cycle products as active and valid.

7. Run the AD Configuration report (adutconf.sql). Review the list of registered products to verify that the product was spliced properly into the database.

Note: For more information, see AD Configuration Report , page 8-5in this book.

8. Log out and log in again so that the new environment file (UNIX) or environment subkey in the registry (Windows) is used to set up the environment.

Note: For more information, see Setting the Environment in Running AD Utilities, page 7-52.

9. Verify that <PROD>_TOP registry and environment variables are set correctly for the newly spliced off-cycle products.

10. Download and apply the patch that introduces the product functionality.

The documentation that instructed you to apply this patch using AD Splicer contains information about which patch you need to apply next.

Maintaining Snapshot Information:

Requirement

What is a snapshot, and how do I use it?

Discussion

Snapshots are current views of your system: they are created once, and then updated when appropriate to maintain a consistent view. There are two types of snapshot: APPL_TOP snapshots and global snapshots. An APPL_TOP snapshot lists patches and versions of files in the APPL_TOP. A global snapshot lists patches and latest versions of files in the entire Applications system (that is, across all APPL_TOPs).

Patch Wizard uses a global snapshot to determine which patches have already been applied. AutoPatch uses an APPL_TOP snapshot to determine if prerequisite patches have been applied to a particular APPL_TOP.

Note: For more information, see Maintain Snapshot Information in Maintaining Applications Files, page 7-69.

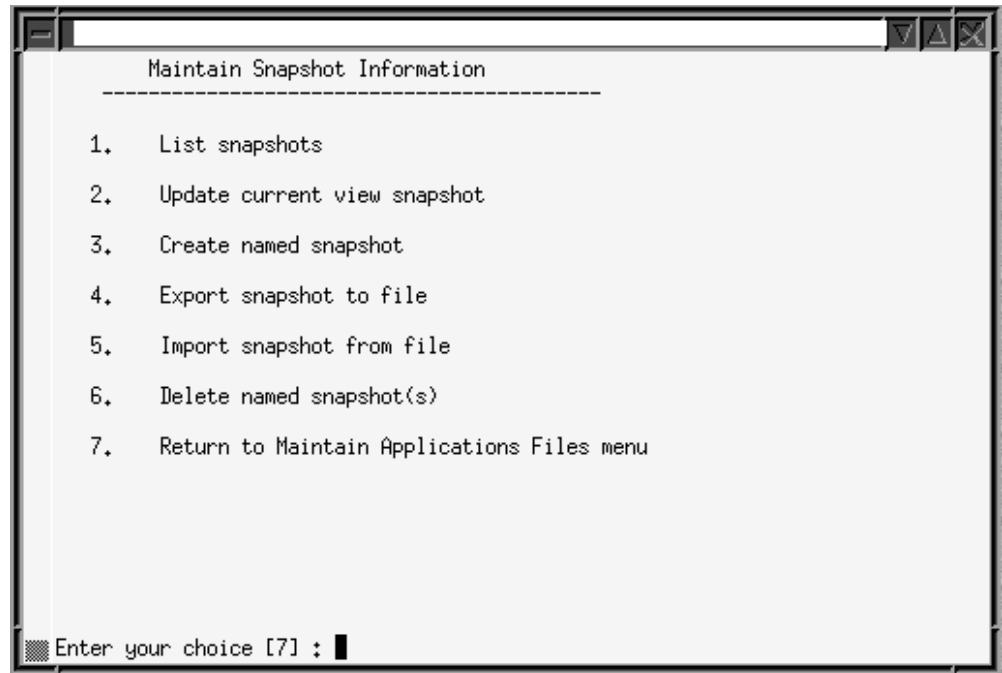
Actions

If you need to perform any of the Maintain Snapshot tasks, select an option from the Maintain Snapshot Information submenu.

1. Access the Maintain Snapshot Information menu.

From the AD Administration Main Menu, choose Maintain Applications Files. Then choose Maintain Snapshot Information.

Maintain Snapshot Information Menu



2. Choose an option.

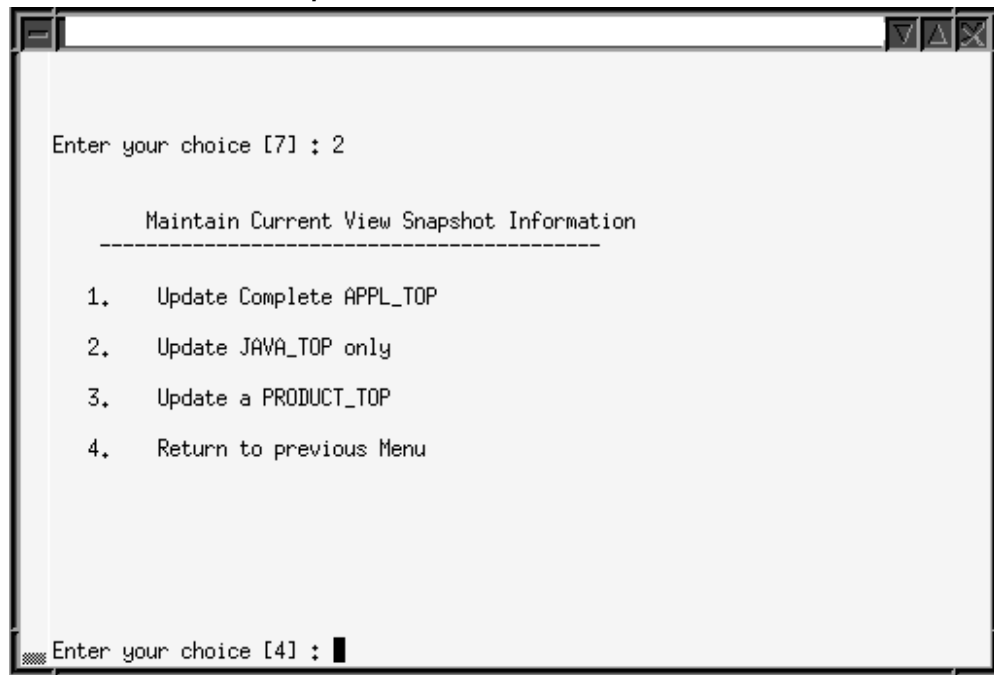
From this menu, you can:

- List snapshots (stored in the system)
- Update current view snapshot (full or partial APPL_TOP and global)
- Create named snapshot (select a current view snapshot to copy and name)
- Export snapshot to file (select one to export to a text file)
- Export snapshot to file (select one to export to a text file)
- Delete named snapshot (select a snapshot for deletion)

In addition to the existing snapshots tasks, you can choose to synchronize selected files - a *partial snapshot* - instead of synchronizing all files for the entire APPL_TOP. You would use this option when you have copied only a few files to the APPL_TOP.

1. Select the Update Current View Snapshot option.
2. From the snapshot submenu, select one of the following options:

Maintain Current View Snapshot Information Menu



- Update Complete APPL_TOP.
This is the original functionality of the Update Current View Snapshot option. It synchronizes all the files in your APPL_TOP.
- Update JAVA_TOP only.
Synchronizes only the files in the JAVA_TOP. At the prompt, enter the path to the JAVA_TOP subdirectory where the files were copied. If the files were copied to more than one directory, press Enter. AD Administration scans the entire JAVA_TOP and updates the information in both the current view and the global view snapshots.
- Update a <PRODUCT>_TOP.
Synchronizes only the files in a specific <PRODUCT>_TOP. Enter the product abbreviation, then provide the subdirectory information at the prompt.

Enter the path to a single subdirectory in the <PRODUCT>_TOP. If the files were copied to more than one directory in the <PRODUCT>_TOP, press Enter. AD Administration scans the entire <PRODUCT>_TOP and updates the information in both the current and the global view snapshots.

During a new installation, Rapid Install automatically creates a current snapshot as a baseline. Then, each time you run AutoPatch, it automatically creates a new (updated) snapshot so that the information is current as of the application of the

patch.

Tip: You can update snapshot information using the AD Administration task any time you think it necessary. However, the process can be time-consuming.

Relinking Product Executables:

Requirement

How do I relink product executables?

Discussion

Relinking executable programs with the Oracle server product libraries keeps them functioning properly. When you need to relink programs, run the AD Administration "Relink Applications Programs" task.

Note: For more information, see Relink Applications Programs in Maintaining Applications Files, page 7-69.

Actions

Perform the following steps:

1. Start AD Administration.

Set the environment and enter `adadmin` on the command line.

Note: For more information, see Setting the Environment in Running AD Utilities, page 7-52.

2. Shut down services.

When relinking files on a concurrent processing server, shut down the concurrent managers. When relinking files on a Forms node, shut down the Forms services.

Note: For more information, see Stopping and Starting Application Tier Services, page 7-22 in this chapter.

3. Relink programs.

From the AD Administration Main menu, go to the Maintain Applications Files menu. Then choose the "Relink Applications programs" task. For each product, choose whether to link all executables or only specific ones.

Relinking AD Executables:

Requirement

How do I relink AD executables?

Discussion

You cannot use AD Administration to relink AD executables. Instead, you run AD Relink. With this command line utility, you can relink several AD utilities with a single command.

AD Relink requires the `force=` parameter. There is no default for this parameter. You must specify either "n" to relink the executable program only if the dependent libraries or object files are more recent than the current executable program, or "y" to relink regardless of the status of the libraries or object files.

An optional command line argument is `backup_mode`. Use it to indicate whether you want to back up executables. There are three possible values for `backup_mode`:

AD Relink backup_mode Values

Value	Effect
<code>backup_mode=none</code>	Do not back up any executables.
<code>backup_mode=all</code>	Back up all executables.
<code>backup_mode=file</code>	Back up files according to instructions in <code>adlinkbk.txt</code> (the default).

Actions

Perform the following steps:

1. Log on as `applmgr` and set the environment.

Note: For more information, see *Setting the Environment in Running AD Utilities*, page 7-52.

Windows users must run `%<APPL_TOP>%\relinkenv.cmd`, executed from a command window. Change directory to `%APPL_TOP%` and run `apps . sh` to set up all required environment variables. (Note there is a space between the dots in this command.)

```
C:\> . ./apps.sh
```

2. Relink files.

Run AD Relink using the appropriate command for your operating system.

UNIX:

```
$ adrelink.sh force={y | n} [<optional arguments>] <ad program name>
```

Windows:

Change directory to %APPL_TOP%\bin and relink the desired file using the following syntax:

```
C:\> sh adrelink.sh force={y | n} [<optional arguments>] <ad program name>
```

If you want to relink several AD utilities, list the programs on the command line, separating each with a space and enclosing it in quotations. For example, to relink both AD Controller (adctrl) and AD Administration (adadmin), enter:

UNIX:

```
$ adrelink.sh force=y "ad adctrl" "ad adadmin"
```

Windows:

```
C:\> sh adrelink.sh force=y "ad adctrl.exe" "ad adadmin.exe"
```

To create a backup file (for all executables), use the following syntax:

UNIX:

```
$ adrelink.sh force=y backup_mode=all
```

Compressing, Archiving, and Deleting Files:

Requirement

Which Oracle E-Business Suite files can be safely compressed, archived, or deleted?

Discussion

There are several types of files that can be compressed, archived, or deleted: log and output files, upgrade files, and AutoPatch backup files. However, Oracle recommends this action only if there is no other way to increase available disk space.

Caution: We strongly recommend creating a backup before you delete any files and keeping the backup readily available in case you need to restore files.

Actions

Perform the following tasks, using the commands specific to your operating system.

1. Compress, archive, or delete the following files, according to your operational requirements. The three categories can be treated independently for cleanup purposes.

- Log and output files.

You can compress, archive, or delete log and output files created by AD utilities. They are located in the following directories, where <SID> is the name of the database instance for the current Applications system: \$APPL_TOP/admin/<SID>/log and \$APPL_TOP/admin/<SID>/out (UNIX) or %APPL_TOP%\admin\<SID>\log and %APPL_TOP%\admin\<SID>\out (Windows).

Caution: Log files may contain passwords. Back up these files to a secure location. Do not delete the directories.

- Upgrade files.

After you complete and verify an upgrade, you can compress, archive, or delete the upgrade files located in \$APPL_TOP/admin/preupg (UNIX) or in %APPL_TOP%\admin\preupg (Windows).

Caution: Do *not* remove any files under <PROD>_TOP/admin. They are used by AD utilities such as AutoPatch and AD Administration.

- AutoPatch backup files.

After you run AutoPatch, you can compress, archive, or delete old files that have been backed up in the patch top subdirectory.

Caution: Verify that the patch was applied successfully and the patched functionalities are fully tested before you delete backup files.

Adding NLS Languages

Adding NLS Languages:

Requirement

I want to add an additional language to my existing system.

Discussion

You can add a new language to your Release 12.2 system at any time after your installation or upgrade.

Note: For more information, see My Oracle Support Knowledge Document 252422.1, *Requesting Translation Synchronization Patches*.

Actions

Perform the following steps:

1. From Oracle Applications Manager, go to License Manager and activate or change your base language to a new one.
2. From AD Administration, run Maintain Multi-lingual Tables (AD Administration Main Menu > Maintain Applications Database Entities Menu).
3. To complete your language installation, refer to *Oracle E-Business Suite NLS Release Notes* for your current release level. Choose the appropriate language installation method for the release level.

Requirement

I want to confirm that my NLS language software is current with the latest US patch levels.

Discussion

If your Oracle E-Business Suite system has active languages other than American English, you can bring them up to the current US Applications patch level by using the Translation Synchronization Patch utility. Alternatively, you can individually download and apply the NLS version of all US patches you have applied to your system. Use AD Merge Patch to create a single patch, and then apply it using AutoPatch.

Actions

Perform the following steps:

1. For details of how to use the Translation Synchronization Utility, follow the instructions in My Oracle Support Knowledge Document 252422.1, *Requesting Translation Synchronization Patches*.

Requirement

I want to check if there are translation updates other than those associated with US patches.

Discussion

There may be updates that enhance your translated software that are not associated with US patches, and therefore, are not included in the updates you received when you requested a Translation Synchronization patch. You can request these updates using the Translation Synchronization Patch utility by selecting the Get Latest Translations check box on the file manifest submission form.

Note: For more information, see My Oracle Support Knowledge Document 252422.1, *Requesting Translation Synchronization Patches*.

Actions

Perform the following steps:

1. Run the Translation Synchronization Patch utility (adgennls.pl).
2. Create a manifest using the form provided in My Oracle Support. When you submit the manifest, click the Get Latest Translations check box option to get translation updates that were made available since the initial Release 12 NLS, in addition to any NLS patches needed to synchronize your NLS patch level with the US patch level.
3. When you are notified that it is available, apply the Translation Synchronization Patch (TSP) for all languages you requested.

Requirement

I want to deactivate a language.

Discussion

Deactivating a language is not supported. Once they have been activated, you must maintain all languages in an NLS system.

Actions

None.

Maintaining the Database

This section contains information you can use to maintain your database and effectively manage system resources.

Using System Resources Efficiently:

Requirement

How do I keep optimization statistics up to date?

Discussion

Optimization is the process of choosing the most efficient way to execute a SQL statement. Oracle E-Business Suite Release 12 uses cost-based optimization. By analyzing the "cost" of using each resource, you can keep your system tuned for optimum performance. The optimizer uses actual table statistics to determine the most efficient access paths and join methods for executing SQL statements.

These statistics are gathered when you run the Gather Schema Statistics concurrent

program. It is important to run this program after an upgrade and, subsequently, on a regular basis to avoid performance degradation (we recommend once a month). The length of time the statistics in an instance are of any value depends on the amount DML that is done during a period of time. For completely static tables, once may be enough for the life of the table. For tables that are completely reloaded all the time, you must run Gather Schema Statistics more often. Tables loaded during a Data Pull in Demand Planning or Advanced Planning and Scheduling are good examples. OE/OM tables are also constantly updated.

Tip: Based on usage, identify the frequency for gathering all statistics, and the frequency that works best for gathering statistics only for specific products.

Actions

Perform the following steps.

1. Log in to Oracle E-Business Suite with the System Administrator responsibility.
2. Navigate to the Submit Request window (Request > Run).
3. Submit the Gather Schema Statistics program.

Set the schema name to ALL to gather statistics for all Oracle E-Business Suite schemas (having an entry in the FND_PRODUCT_INSTALLATIONS table). In addition to gathering index and table-level statistics, the procedure also gathers column-level histogram statistics for all columns listed in the FND_HISTOGRAM_COLS table.

Note: For more information, see Cost-Based Optimization in Oracle E-Business Suite in the *Oracle E-Business Suite Configuration Guide*.

Validating the APPS Schema:

Requirement

How do I verify the integrity of my APPS schema?

Discussion

AD Administration can run a SQL script (advrfapp.sql) against the APPS schema that checks for certain conditions that are undesirable, but will not produce fatal problems. The Validate APPS Schema task executes this script.

You can run this task at any time, but it is most effective if run:

- Immediately after an upgrade
- Before converting to Multi-Org

- After performing an export/import (migration)
- As a part of custom development in the APPS schema

Actions

Perform the following tasks, using the commands specific to your operating system.

1. Start AD Administration.

Set the environment and enter `adadmin` on the command line.

Note: For more information, see *Setting the Environment in Running AD Utilities*, page 7-52.

2. Validate APPS schema.

Select the "Validate APPS schema" task from the Maintain Applications Database Entities menu. Review the output file (`<APPS schema name>.lst`) for invalid database objects. It is located in `$APPL_TOP/admin/<SID>/out` (UNIX) or in `%APPL_TOP%\admin\<SID>\out` (Windows)

Note: For more information, see *Validate APPS Schema* in this chapter.

You can also run this task from SQL*Plus:

UNIX:

```
$ cd $APPL_TOP/admin<SID>/out
$ sqlplus <SYSTEM username>/<SYSTEM password> \
@$AD_TOP/admin/sql/advrfapp.sql \
<APPS schema name> <AOL schema name>
```

Windows:

Change directory to `%APPL_TOP%\bin` and relink the desired file using the following syntax:

```
C:\> cd %APPL_TOP%\admin\<SID>out
C:\> sqlplus <SYSTEM username>/<SYSTEM password>
@%AD_TOP%\admin\sql\advrfapp.sql <APPS schema name> <AOL schema name>
```

3. Resolve any issues.

The `<APPS schema name>.lst` file is divided into three sections:

- Issues you *must* fix (not specific to the APPS schema)
- Issues you *must* fix (specific to the APPS schema)
- Issues you may want to address (specific to the APPS schema)

Each section of the file contains instructions for resolving the issues that are listed.

Recreating Grants and Synonyms:

Requirement

How do I recreate grants and synonyms in the APPS schema?

Discussion

In order to maintain database objects, you should check the APPS schema for missing grants and synonyms. Using the AD Administration menu, you can run tasks to validate the APPS schema and then recreate any missing grants and synonyms.

Note: For more information, see Recreate Grants and Synonyms for APPS Schema in this chapter.

Actions

Perform the following steps:

1. Start AD Administration.

Set the environment and enter `adadmin` on the command line.

Note: For more information, see Setting the Environment in Running AD Utilities, page 7-52.

2. Recreate grants and synonyms.

From the Main AD Administration menu, go to the Maintain Applications Database Entities menu. Select the "Recreate grants and synonyms for APPS schema" task.

Compiling Invalid Objects:

Requirement

When should I compile invalid objects?

Discussion

The Oracle database automatically compiles invalid database objects the first time an object is used and during patch application. This action can take some time, so you may want to compile objects before the first use, at a time when you know the system usage is low.

You compile invalid objects with AD Administration. This task is most effective under the following circumstances:

- After custom packages are moved to the APPS schema and need to be compiled

- After applying patches that alter packages in the APPS schema
- After validating the APPS schema and identifying invalid objects

Actions

Perform the following tasks, using the commands specific to your operating system.

1. Start AD Administration.

Set the environment and enter `adadmin` on the command line.

Note: For more information, see *Setting the Environment in Running AD Utilities*, page 7-52.

2. Compile Applications schema.

From the Main AD Administration menu, go to the Compile/Reload Database Entities menu. Choose the "Compile APPS schema" task.

Pre-Allocating Space for Packages and Functions:

Requirement

How do I ensure that there is enough space in the System Global Area (SGA) for packages and functions?

Discussion

If SGA space is fragmented, there may not be enough for certain packages or functions. You can pre-allocate space in the SGA shared pool by *pinning* packages, functions, and sequences. The scripts described in this procedure work well as templates and can be used to create your own custom pinning scripts.

Important: Run these scripts when packages or sequences are patched (and the patch readme file tells you to do so), or any time after objects are invalidated, either because of patching or customizations.

The `ADXGNPIN.sql` script pins packages and functions in the APPS schema, while `ADXGNPNS.sql` pins sequences in the base product schemas. Both scripts take the name of a schema as an argument, or `%` for all schemas. `ADXGNPIN.sql` generates and invokes another SQL file, `ADXSPPIN.sql`. `ADXGNPNS.sql` generates and runs `ADXSPDNS.sql`.

Actions

Perform the following tasks, using the commands specific to your operating system.

1. Create the `appsutil/admin` directory in the `ORACLE_HOME` of the database server, if it does not already exist.

2. Copy ADXGNPIN.sql and ADXGNPNS.sql from the AD_TOP/sql directory of the administration server to this directory.
3. Set the environment to point to the ORACLE_HOME for the database server.
4. On the database server, go to the directory created in Step 1 and run ADXGNPIN.sql and ADXGNPNS.sql from SQL*Plus:

```
$ sqlplus <SYS username>/<SYS password>@ADXGNPIN.sql \
<APPS schema name>
$ sqlplus <SYS username>/<SYS password>@ADXGNPNS.sql \
<Base product schema name>
```

Listing Objects in the Shared Pool:

Requirement

How can I see a list of objects stored in the shared pool?

Discussion

You can run the ADXCKPIN.sql script to query for objects stored in the SGA shared pool. It shows the objects known to the SGA and the size that they consume. The output file is ADXCKPIN.lst.

Actions

1. Run the following commands:

```
$ cd $APPL_TOP/admin<SID>/out
$ sqlplus <SYSTEM username>/<SYSTEM password> \
@%AD_TOP%\sql\ADXCKPIN.sql
```

Performing Maintenance Tasks Non-Interactively

Unless otherwise noted, maintenance tasks described in this book are performed interactively: they require user intervention, primarily in the form of responding to prompts. However, you can schedule certain AD Administration and AD Controller tasks to run with little or no user intervention by running these utilities *non-interactively*: instead of responding to prompts each time you run the task, you specify a file that contains the information necessary to complete the task without user intervention. In such a case, there is no need to monitor the process in order to respond to prompts. The file used is referred to as a *defaults file*.

Scheduling Non-Interactive Maintenance:

Requirement

How do I schedule and run maintenance tasks non-interactively?

Discussion and Actions

To set up a non-interactive task requires the creation of a *defaults file*. AutoConfig automatically creates a defaults file (adalldefaults.txt) each time it runs. This file can be used as a template to create a customized defaults file.

Important: Running AutoConfig is now the recommended way to create a defaults file, in contrast to previous releases where adadmin could be used to do so.

Once the defaults file has been created and customized, you start the relevant maintenance utility from the command line, specifying the name of the defaults file, a log file name, and the number of parallel workers.

The same defaults file can be used to run different AD Administration commands: a single such file can contain all your choices for the different menu options. In order to choose which task the defaults file will run, you also add `menu_option=<menu choice>` to the utility start command. This overrides any menu-specific key stroke information stored in the defaults file initially, and allows you to use the defaults file for any of the AD Administration menu items. It also ensures that the menu option you intended for the defaults file is always valid, even if the menu items are renumbered or relocated in subsequent releases.

Note: For more information, see *Preparing for Non-Interactive Processing*, page 7-62.

Restarting a Failed Session:

Requirement

My non-interactive AD Administration session failed. How do I restart it?

Discussion

To restart a failed non-interactive session, you run AD Administration using the `restart=yes` parameter.

Actions

Perform the following tasks, using the commands specific to your operating system.

1. Determine the reason the session failed and fix the issue.
2. Run AD Administration from the command line.

Use the same parameters that you used to start the original non-interactive session, plus the `restart=yes` parameter. For example:

UNIX:

```
$ adadmin defaultsfile=$APPL_TOP/admin/testdb1/adadmindef.txt \  
logfile=adadmin_noninteractive.log workers=5 interactive=n \  
restart=y menu_option=CHECK_DUAL
```

Windows:

```
C:\> adadmin defaultsfile=%APPL_TOP%\admin\testdb1\adadmindef.txt \  
logfile=adadmin_noninteractive.log workers=5 interactive=n restart=y \  
menu_option=CHECK_DUAL
```

3. AD Administration runs the task. It does *not* prompt you to continue the previous (failed) session.

Distribute Processing With Distributed AD

Requirement:

How can I distribute tasks across my multi-node system?

Discussion:

Distributed AD is a special parallel processing feature that can be employed to decrease the time needed for patch application (and other tasks) by allocating the associated worker processes to multiple application tier nodes. AD Administration and AutoPatch (adop) run on one node, and direct workers running both on that node and on other nodes in the system.

Note: You must have a shared application tier file system to use Distributed AD.

Actions:

The distribution of workers is specified as follows:

```
workers=<total number of workers> localworkers=<number of workers on  
primary node>
```

The following two examples will illustrate this.

Example 1 - Distribute a total of eight workers across a two-node system

1. To begin, enter a command that will run an adop session with three workers on the primary node and five workers on secondary nodes:

```
$ adop phase=apply input_file=myinput.txt
```

The file `myinput.txt` will need to include the lines:

```
workers=8  
localworkers=3
```

2. Now start an AD Controller session on each of the secondary nodes that will run

workers, using the `distributed=y` argument.

```
$ adctrl distributed=y
```

3. To start workers 4 through to 8 on a secondary node, enter "4-8" in response to the prompt from AD Controller:

```
Enter the worker range: 4-8
```

Note: Workers must be specified in contiguous sequences such as 1-4 or 5-8. You cannot, for example, start workers 1, 3, 5, 7 on one node, and workers 2, 4, 6, and 8 on another.

Example 2- Distribute a total of twelve workers across a three-node system

1. To begin, enter a command that will run an adop session with four workers on the primary node and eight workers on secondary nodes:

```
$ adop phase=apply input_file=myinput.txt workers=12 localworkers=4
```

The file `myinput.txt` will need to include the lines:

```
workers=12  
localworkers=4
```

2. Now start an AD Controller session on the second node, specifying that workers 5-8 should run there:

```
$ adctrl distributed=y  
Enter the worker range: 5-8
```

Note: As in the previous example, workers must be specified in contiguous sequences such as 1-4 or 5-8.

3. Finally, start AD Controller on the third node, specifying that the last four workers (9-12) should run there:

```
$ adctrl distributed=y  
Enter the worker range: 9-12
```

Managing Application Tier Services

When running certain scripts or utilities, you may be directed to stop application tier service (server) processes manually. This section contains information about stopping and starting these processes.

Note: Scripts in this section may contain system-specific information. If you change the Rapid Install defaults, you may need to edit the scripts before rerunning them.

Starting and Stopping Application Tier Services:

Requirement

How do I start and stop application tier services?

Discussion

When Rapid Install sets up and configures the service and server processes, it stores a control script for each process in the `$INST_TOP/admin/scripts` directory.

Certain maintenance procedures require that you stop one or more services or servers manually, and restart them after you complete the procedure. By running the appropriate script on the command line, along with a stop or start argument, you can stop (or start) a single server process, several processes, or all processes. The following table lists the key scripts.

Application Tier Service Process Control Scripts

Service Process and Script Action	Implementation on UNIX	Implementation on Windows
HTTP (Web) Server: Used to start, stop, and check the status of HTTP (Web) server.	<code>adapcctl.sh</code>	<code>adapcctl.cmd</code>
Oracle Process Manager (opmn): Used to start, stop and check the status of opmn.	<code>adopmnctl.sh</code>	<code>adopmnctl.cmd</code>
Concurrent Processing: Used to start, stop, and check the status of concurrent managers.	<code>adcmctl.sh</code>	<code>adcmctl.cmd</code>
Forms (Socket): Used to start, stop and check the status of the Forms services in Socket Mode.	<code>adformsrvctl.sh</code>	<code>adformsrvctl.cmd</code>
Managed Server Control: Used to start, stop and check the status of the Managed Servers. The desired server name must be supplied on the command line: oacore_server1, forms-c4ws_server1, forms_server1, or oafm_server1.	<code>admanagedsrvctl.sh</code>	<code>admanagedsrvctl.cmd</code>

Service Process and Script Action	Implementation on UNIX	Implementation on Windows
WLS Admin Server: Used to start, stop, and check the status of the Admin Server.	adadminsrvctl.sh	adadminsrvctl.cmd
Node Manager: Used to start, stop, and check the status of Node Manager. Each node in a WLS domain has a Node Manager.	adnodemgrctl.sh	adnodemgrctl.cmd
Oracle TNS Listener: Used to start, stop, and check the status of the TNS Listener.	adalnctl.sh	adalnctl.cmd
Start all application tier server processes: Used to start all processes with one command.	adstrtal.sh	adstrtal.cmd
Stop all application tier server processes: Used to stop all processes with one command.	adstpall.sh	adstpall.cmd

Actions

Choose the procedure that meets your needs.

To start or stop a single application tier server process (UNIX)

Use a command of the following form:

```
<process script name> [stop | start]
```

Tip: Many of the relevant scripts also have other options, such as 'status'. Entering the script name alone will display a list of the available options.

1. Open a terminal window.
2. To stop the Concurrent Processing server (for example), run the adcmctl.sh script with the 'stop' option:

```
% adcmctl.sh stop
You are running adcmctl.sh version 120.19
Enter the APPS username: <APPS username>
Enter the APPS password: <APPS password>
```

To start or stop a single application tier server process (Windows)

On Windows, services can be started or stopped using the appropriate process control script (command file), or from the Services Control Panel.

Using Process Script

1. Open a command window.
2. To stop the Concurrent Processing server (for example), run the `adcmctl.cmd` script with the 'stop' option:

```
C:\> adcmctl.cmd stop
You are running adcmctl.cmd version 120.19
Enter the APPS username: <APPS username>
Enter the APPS password: <APPS password>
```

Using Services Control Panel

1. Go to Start > Administrative Tools and click Services.
2. Select the relevant service in the Services window.
3. Click Start or Stop, as required.

To start all application tier server processes (UNIX)

Use a command of the following format:

```
<process script name> [stop | start]
```

1. Open a terminal window.
2. Enter the command:

```
$ adstrtal.sh
You are running adstrtal.sh version 120.24
Enter the APPS username: <APPS username>
Enter the APPS password: <APPS password>
Enter the WebLogic Server password: <WLS password>
```

A more secure alternative, which does not require the APPS credentials to be supplied, is:

```
$ adstrtal.sh -secureapps
You are running adstrtal.sh version 120.24
Enter the Applications username: <Concurrent Manager operator
username>
Enter the Applications password: <Concurrent Manager operator
password>
Enter the WebLogic Server password: <WLS password>
```

To start all application tier server processes (Windows)

Use a command of the following format:

```
<process script name> [stop | start]
```

1. Open a command window.

2. Enter the command:

```
$ adstrtal.cmd
You are running adstrtal.cmd version 120.24
Enter the Applications username: <APPS username>
Enter the Applications password: <APPS password>
Enter the WebLogic Server password: <WLS password>
```

A more secure alternative, which does not require the APPS credentials to be supplied, is:

```
$ adstrtal.cmd -secureapps
You are running adstrtal.cmd version 120.24
Enter the Applications username: <Concurrent Manager operator
username>
Enter the Applications password: <Concurrent Manager operator
password>
Enter the WebLogic Server password: <WLS password>
```

To stop all application tier server process (UNIX)

Use a command of the following format:

```
<process script name> [stop | start]
```

1. Open a terminal window.

2. Enter the command:

```
$ adstpall.sh
You are running adstpall.sh version 120.24
Enter the Applications username: <APPS username>
Enter the Applications password: <APPS password>
Enter the WebLogic Server password: <WLS password>
```

A more secure alternative, which does not require the APPS credentials to be supplied, is:

```
$ adstpall.sh -secureapps
You are running adstpall.sh version 120.24
Enter the Applications username: <Concurrent Manager operator
username>
Enter the Applications password: <Concurrent Manager operator
password>
Enter the WebLogic Server password: <WLS password>
```

To stop all application tier server process (Windows)

Use a command of the following format:

```
<process script name> [stop | start]
```

1. Open a command window.

2. Enter the command:

```
$ adstpall.cmd
You are running adstpall.cmd version 120.24
Enter the Applications username: <APPS username>
Enter the Applications password: <APPS password>
Enter the WebLogic Server password: <WLS password>
```

An alternative, which does not require the APPS password at all, is:

```
$ adstpall.cmd -secureapps
You are running adstpall.cmd version 120.24
Enter the Applications username: <Concurrent Manager operator
username>
Enter the Applications password: <Concurrent Manager operator
password>
Enter the WebLogic Server password: <WLS password>
```

Starting and Stopping Database Tier Services:

Requirement

How do I start or stop the Oracle Net Services listener manually?

Discussion

When Rapid Install sets up and configures the server processes during installation, it stores a script for the Net Services listener process in the Oracle 11g database server `$ORACLE_HOME/apputil/scripts/<CONTEXT_NAME>` directory. You use this script to start or stop the Net Services listener process for the database.

Actions

To start or stop the Net Services listener (UNIX)

1. Open a terminal window.
2. Log in as the oracle user on the database server and navigate to the `$ORACLE_HOME/apputil/scripts/<CONTEXT_NAME>` directory.
3. Enter a command of the form:

```
$ addlnctl.sh [start|stop] <listener_name>
```

Tip: Many of the relevant scripts also have a 'status' option, which is often useful.

For example, to start the PROD listener, enter:

```
$ addlnctl.sh start PROD
```

Note: For more information, see the *Oracle Net Services Administrator's Guide*.

To start or stop the Net Services listener (Windows)

1. As the oracle user, open a command window and navigate to the %ORACLE_HOME%\appsutil\scripts\

2. Enter a command of the form:

```
C:\> addlnctl.cmd [start|stop] <listener_name>
```

For example, to start the PROD listener, enter:

```
C:\> addlnctl.cmd start PROD
```

Note: For more information, see the *Oracle Net Services Administrator's Guide*.

Requirement

How do I start or stop the Oracle database manually?

Discussion

When Rapid Install sets up and configures the server processes during installation, it creates a script for the database process in the Oracle 11g database server \$ORACLE_HOME/appsutil/scripts/<CONTEXT_NAME> directory. You use this script to start or stop the database on your database tier.

Actions

To start or stop the Oracle database (UNIX)

1. Log in as the oracle user on the database server.
2. Open a terminal window and navigate to the \$ORACLE_HOME/appsutil/scripts/<CONTEXT_NAME> directory.
3. Enter a command of the form:

```
$ addbctl.sh [start|stop] {immediate|abort|normal}
```

Tip: Many of the relevant scripts also have a 'status' option, which is often useful.

For example, to stop the database using the normal option, you would enter:

```
$ addbctl.sh stop normal
```

To start or stop the Oracle database (Windows)

1. Log in as the oracle user on the database server.
2. Open a command window and navigate to the %ORACLE_

HOME%\appsutil\scripts\

3. Enter a command of the form:

```
C:\> addlnctl.cmd [start|stop] <listener_name>
```

For example, to start the PROD listener, enter:

For example, to stop the database using the normal option, you would enter:

```
C:\> addbctl.cmd stop normal
```

Oracle E-Business Suite Maintenance Utilities

You use Oracle E-Business Suite system maintenance utilities to perform a variety of operations from installing and upgrading Oracle E-Business Suite systems, to updating configuration parameters, to maintaining and patching your database and file system, to producing system reports.

In this book, these utilities have been categorized by the way you access and use them. This may be from the command line, or via a Web-based interface.

Command Line Utilities

The tools generally referred to as Applications DBA (AD) utilities are started and run from the command line. They initiate processes that perform a variety of system maintenance tasks, such as applying and merging patches. As they run, the utilities prompt you for system-specific parameters necessary to perform the maintenance task. In addition, many of the utilities produce reports that contain information such as job timing and file versions.

The AD utilities have similar interfaces, operation, input, and report formats. Many also share the ability to accept arguments, flags, and options, which you can use to refine the actions they perform. You add the argument on the command line when you start the utility. For example, to specify the number of workers that AutoPatch should run in parallel when applying a patch, you enter the number of worker processes on the command line when you start AutoPatch. A list of commonly used command line arguments and flags, and a brief description of how to use them, begins later in this chapter.

The command line maintenance utilities are listed in the following table. Their operation is described further in this book, or, in the case of Rapid Install, in *Oracle E-Business Suite Installation Guide: Using Rapid Install*.

AD Command Line Utilities

AD Utility Name	Executable or Script	Description	Usage Restrictions
AD Administration	adadmin	Performs maintenance tasks for Oracle E-Business Suite.	If a patch edition exists, the tool can be invoked from the patch file system. If a patch edition does not exist, the tool can be invoked from the run file system. In all other cases, adadmin will display an error message.
AD Check Digest	adchkdig	Checks the integrity of Oracle E-Business Suite patches downloaded from My Oracle Support.	None.
AD Configuration	adutconf.sql	Reports standard information about the installed configuration of Oracle E-Business Suite.	If a patch edition exists, the tool can be invoked from the patch file system. If a patch edition does not exist, the tool can be invoked from the run file system.
AD Controller	adctrl	Manages parallel workers in AD Administration and AutoPatch.	If adop is run in hotpatch mode, you should run adctrl from the run file system. Otherwise, run adctrl from the patch file system.
AD File Identification	adident	Reports the version and translation level of an Oracle E-Business Suite file.	None.

AD Utility Name	Executable or Script	Description	Usage Restrictions
AD File Character Set Converter	adncnv	Converts a file from one character set to another.	You should run adncnv from the same \$APPL_TOP as the source file (to be converted) resides on.
AD Merge Patch*	admrgpch	Merges multiple patches into a single merged patch.	In Release 12.2, adop is the recommended tool for merging patches. If you still wish to run AD Merge Patch, you should do so from the run file system.
AD Relink	adrelink.sh	Relinks Oracle E-Business Suite executable programs with the Oracle server product libraries.	The tool will regenerate the executables on whichever file system it was invoked on.
AD Splicer	adsplce	Adds off-cycle products.	If a patch edition exists, the tool can be invoked from the patch file system. If a patch edition does not exist, the tool can be invoked from the run file system.
AD Job Timing Report		Reports a summary of the timing for jobs run by parallel workers.	If adop is run in hotpatch mode, you should run adtimrpt.sql from the run file system. Otherwise, run adtimrpt.sql from the patch file system.
AD Online Patching	adop	Applies patches and other system updates.	Must always be run from the run file system.

AD Utility Name	Executable or Script	Description	Usage Restrictions
Patch Application Assistant	admsi.pl	Generates customized installation instructions for a patch.	<p>If you are applying a patch in hotpatch mode, you must run Patch Application Assistant from the run file system.</p> <p>If you are not using hotpatch mode, run Patch Application Assistant in the patch file system after the prepare phase.</p>
Rapid Install	rapidwiz	Provides a wizard for entering parameters that are specific to a new installation or an upgrade of an Oracle E-Business Suite system.	None.

Online Patching and AD Utilities

The use of online patching in Oracle E-Business Suite Release 12.2 has implications for the operation of AD Admin and AD Splicer. Specifically, if these utilities are run during an online patching cycle and perform tasks that change the file system, this will trigger an fs_clone during the prepare phase of the next online patching cycle. That is to say, the run file system will be cloned (copied) to the patch file system.

Such tasks include:

- GEN_MESSAGES
- GEN_FORMS
- GEN_REPORTS
- GEN_JARS
- RELINK
- COPY_FILES
- CONVERT_CHARSET

- CMP_INVALID
- CMP_MENU
- CMP_FLEXFIELDS

Important: AD Admin and AD Splicer can be invoked from the run edition only if there is no patch edition. An error will be displayed if you try to run any of these utilities from the run edition and a patch edition exists. In other words, before running them you will either set your environment to the patch file system, or, if no patch edition exists, set your environment to the run file system.

For more information about the effect of online patching on Oracle E-Business Suite maintenance activities, refer to Choosing the Correct File System For Maintenance Tasks, page 7-1 in this book.

Web-Based Utilities

Oracle Applications Manager (OAM) provides a Web-based interface where system administrators can monitor system status, administer services, examine system configuration, manage Oracle Workflow, view applied patches, and measure system usage. It provides a concise overview of the state of your Oracle E-Business Suite system, and serves as a gateway to utilities for tasks such as managing system configuration, reviewing patch history, determining which patches will bring your system up to date, registering additional products and languages, and other maintenance activities.

The Web-based maintenance utilities are listed in the following table. Their operation is described further in Part 2 of this book and in *Oracle E-Business Suite Setup Guide*.

Oracle Applications Manager Utilities

OAM Utility Name	Description
Applied Patches	Uses key patch information in the patch history database. You can search the database to create reports in several formats.
AutoConfig	Use to view current context files, edit parameters contained in the context files, view previous context files, and compare current context files against previous ones.

OAM Utility Name	Description
File History	Enables the viewing of files that have been updated by a patch.
License Manager	Registers additional Oracle E-Business Suite products, country-specific functionalities, or languages. You can also use License Manager to change the base language for your system.
Patch Wizard	Determines patches that have not been applied, but that should be applied to keep the system current. Downloads and merges patches from My Oracle Support.
Register Flagged Files	Used to record any files in which you have made customizations. Replaces the need to use applcust.txt, which contained the record for all customized files in previous releases.
Software Updates	Provides an overview of all patching-related information for your system.
Timing Reports	Helps you monitor jobs that are running or view statistics of completed AutoPatch and AD Administration maintenance sessions.

Online Help

Both the AD utilities and the OAM utilities provide a help function.

Command Line Help

For the AD command line utilities, you can request a list of arguments by entering the utility name with `help=y` appended. For example, to access help for AD Administration, enter the command:

```
adadmin help=y
```

The arguments and options that you can use to refine the operation of a utility are listed, along with a brief description of how they work. Below is an example of the command line help for AD Administration:

```
usage: adadmin [help=y]

adadmin
  [printdebug=y|n] [localworkers=<localworkers>]
  [flags=hidepw|trace]

adadmin Non-Interactive mode
  [defaultsfile=<${APPL_TOP}/admin/SID/defaultsfile>]
  [logfile=<logfile>] [interactive=y|n]
  [workers=workers>] [menu_option=ASK_NAME>] [restart=y|n]
```

where

Key to options:localworkers = The number of workers to run on the local machine. Used in Distributed AD.

flags = Generic flags passed to AD utilities. The available values for AD Admin are hidepw and trace.

defaultsfile = The defaults file filename, located under \$APPL_TOP/admin/SID/ directory.

menu_option = Skips the AD Admin menu and executes the task supplied on the command line. Valid values are listed below.

```
RELINK          Relink Applications programs
GEN_MESSAGES    Generate message files
GEN_FORMS       Generate form files
GEN_REPORTS     Generate reports files
GEN_JARS        Generate product JAR files
VALIDATE_APPS   Validate APPS schema
CMP_INVALID     Compile APPS schema
CMP_MENU        Compile menu information
CREATE_GRANTS   Recreate grants and synonyms for APPS schema
CMP_FLEXFIELDS  Compile flexfield data in AOL tables
MAINTAIN_MLS    Maintain multi-lingual tables
CHECK_DUAL      Check DUAL table
RELOAD_JARS     Reload JAR files to database
COPY_FILES      Copy files to destinations
CHECK_FILES     Check for missing files
LIST_SNAPSHOTS  List snapshots
UPDATE_CURRENT_VIEW Update current view snapshot
CREATE_SNAPSHOT Create named snapshot
EXPORT_SNAPSHOT Export snapshot to file
IMPORT_SNAPSHOT Import snapshot from file
DELETE_SNAPSHOT Delete named snapshot(s)
CONVERT_CHARSET Convert character set
SCAN_APPLTOP    Scan the APPLTOP for exceptions
SCAN_CUSTOM_DIR Scan a CUSTOM directory for exceptions
```

Obtaining Help in OAM

OAM Help is available by clicking the *Help* link in the top right-hand section of any Oracle Applications Manager screen.

OAM Site Map and Help Link

ORACLE Applications Manager

Support Cart Setup Home Logout Help

Applications Dashboard | Site Map

Site Map: VIS

Administration Monitoring Maintenance Diagnostics And Repair

System Configuration	Application Services	Workflow
Hosts	Generic Services	Home
AutoConfig	Request Processing Managers	Work Item Metrics
License Manager	Transaction Managers	Agent Activity
	Parallel Concurrent Programming Setup	Background Engines
		Notification Mailer
		Service Components
		Purge

Concurrent Requests	Service Fulfillment Manager	Others
Submit New	Service Fulfillment Manager	Applications Manager Log
Pending		Knowledge Base
Running		
Completed (Last Hour)		

TIP Only the items to which you have access are clickable.

Support Cart Setup Home Logout Help

Copyright 2001, 2006 Oracle Corporation. All Rights Reserved.
About Oracle Applications Manager Version 2.3.1

For example, from the OAM Site Map, OAM displays page-specific help describing the features of the Site Map page.



Individual help topics may include topical essays, procedures, and page descriptions. The help associated with the utilities and features discussed in this book provides navigation paths, field definitions, and general information about using the page.

Command Line Utilities

The AD maintenance utilities were developed to perform specific Oracle E-Business Suite maintenance and reporting tasks from the command line. For example, you use AD Online Patching (`adop`) to apply all types of patches to your system, and you use AD Administration to perform routine maintenance tasks.

However, even though the utilities each have a specialized function, they are designed to complement each other, so many employ similar operations. This section summarizes

the operations that AD utilities have in common. Subsequent chapters describe each utility's features in detail.

Common AD Operations

Many AD utilities employ similar features and operations as they perform processing tasks. For example, most rely on prompts to gather values for system-specific processes, and all automatically create log files to record processing actions. This section describes some of these common operations.

Prompts

Many AD utilities prompt for information necessary for completing a task. Prompts typically include a description of the information needed, and may include a default answer (in square brackets). You can just press the [Return] key to accept the default.

For example:

```
The ORACLE username specified below for Application Object Library
uniquely identifies your existing product group: APPLSYS
```

```
Enter the ORACLE password of Application Object Library [APPS] :
```

```
Press [Return] to accept the default value, or type a new value
after the colon and press [Return]. Read the prompts carefully
to make sure you supply the correct information.
```

Interactive and Non-Interactive Processing

The AD utilities perform processing tasks interactively by default. That means the utility prompts for system-specific information at the point where it needs it, making it necessary for you to be present during the entire operation in order to respond to the prompts.

AD Administration, `adop`, `ad`, and AD Controller can run some file system and database tasks non-interactively: you store the required information in a defaults file, and the utility reads the information from this file instead of prompting you for the input. Non-interactive processing is useful for scheduling routine tasks that require little or no user intervention.

Note: For more information, see *Performing Maintenance Tasks Non-Interactively*, page 7-19 in this book. See also *Monitoring and Controlling Parallel Processes*, page 7-58 in this chapter.

Special Parameter for Using `adop` and AD Administration Non-Interactively

When running AD Online Patching, AD Administration, or AD Splicer in non-interactive mode, the "stdin=y" option can optionally be used to prompt for passwords in the standard input. The default is for passwords to be supplied without prompting.

Log Files

All AD utilities record their processing actions and any errors that they encounter in log files. Many utilities prompt you for the name of the log file that will record the processing session, with a display such as this:

```
<utility name> records your <utility name> session in a text file you
specify.
Enter your <<utility name> log file name or press [Return] to accept the
default name shown in brackets.
```

```
Filename [<utility name>.log] :
```

The default file name is <utility name>.log. For example, for AD Administration, the default log file is adadmin.log.

AD Administration log files are stored in the following locations:

UNIX:

```
$APPL_TOP/admin/<SID>/log
```

Windows:

```
%APPL_TOP%\admin\<SID>\log
```

Restart Files

Restart files contain information about what processing has already been completed. They are located in \$APPL_TOP/admin/<SID>/restart (UNIX) or in %APPL_TOP%\admin\<SID>\restart (Windows).

If a utility stops during processing due to an error, or you use AD Controller (in the case of parallel processing) to shut down workers while they are performing processing tasks, you can restart the utility. If you do, it looks for restart files to determine if there was a previous session. If the files exist, the utility prompts you to continue where the processing left off, or to start a new process. If you choose to continue, it reads the restart files to see where the process left off, and continues the process from that point.

Caution: Do not modify or delete any manager or worker restart files unless specifically directed to do so by Oracle Support Services.

By default, AD utilities delete their restart files when processing is complete, but leave backup versions with the extensions .bak, .bk2, or .bk3.

Warning: Restart files record passwords for your Oracle E-Business Suite products. You should restrict access to all restart files (located in \$APPL_TOP/admin/<SID>/restart). If you are running a utility with options=nohidepw, the log files may also contain passwords on lines prefixed with HIDEPW.

Configuration and Environment Files

Most AD utilities require access to system parameters stored in various configuration and environment files when processing maintenance tasks. For example, it may be necessary to know the location of an Oracle Application Server ORACLE_HOME or the Database (RDBMS) ORACLE_HOME.

Configuration and environment files are generated by AutoConfig during an installation or upgrade. You typically do not have to manually update or maintain the information in these files. They are updated when you run AutoConfig.

Note: For more information, see AutoConfig in *Oracle E-Business Suite Concepts*. See also My Oracle Support Knowledge Document 1380535.1, *Using AutoConfig to Manage System Configurations in Release 12.2*.

The following table lists configuration and environment files commonly used by the AD command line utilities, and in some cases, by the OAM Web-based utilities.

Note: <CONTEXT_NAME> defaults to <SID>_<hostname>.

AutoConfig-Managed AD Utility Files

File name	Location	Description
adconfig.txt	\$APPL_TOP/admin	Contains environment information used by all AD utilities. Warning: Do not update this file manually.
<CONTEXT_NAME>.env (UNIX)	RDBMS ORACLE_HOME	Used to configure the environment when performing maintenance operations on the database.
<CONTEXT_NAME>.cmd (Windows)		
APPS<CONTEXT_NAME>.en v (UNIX)	APPL_TOP	This file calls the environment files needed to set up the APPL_TOP and the Applications ORACLE_HOME.
APPS<CONTEXT_NAME>.c md (Windows)		

File name	Location	Description
<CONTEXT_NAME>.env (UNIX)	APPL_TOP	Called by APPS<CONTEXT_NAME>.env (UNIX) or APPS<CONTEXT_NAME>.cmd (Windows) file to set up the APPL_TOP. This file calls either adovars.env (UNIX) or adovars.cmd (Windows).
<CONTEXT_NAME>.cmd (Windows)		
<CONTEXT_NAME>.env (UNIX)	\$INST_TOP/appl/admin	Called by APPS<CONTEXT_NAME>.env (UNIX) or APPS<CONTEXT_NAME>.cmd (Windows).
<CONTEXT_NAME>.cmd (Windows)		
adovars.env (UNIX) adovars.cmd (Windows)	APPL_TOP/admin	Called by the <CONTEXT_NAME>.env (UNIX) or <CONTEXT_NAME>.cmd (Windows) file located in the APPL_TOP. Used to set environment variables for Java and HTML.

The following configuration and environment files are also used by most AD utilities, but are not created by AutoConfig.

Warning: Do not update any of these AutoConfig-managed files manually.

Non-AutoConfig AD Utility Files

File name	Location	Description
applora.txt	APPL_TOP/admin	Contains information about required init.ora parameters for runtime.

File name	Location	Description
applorau.txt	APPL_TOP/admin	Contains information about required command lineinit.ora parameters for install and upgrade.
applprod.txt	APPL_TOP/admin	The AD utilities product description file, used to identify all products and product dependencies.
applterr.txt	APPL_TOP/admin	The AD utilities territory description file. It contains information on all supported territories and localizations.
fnenv.env	FND_TOP	Sets additional environment variables used by Oracle Application Object Library. The default values should be applicable for all customers.

Feature Version Numbers

In order to use some AD Administration and adop features, the version number of the feature must be the same in both the file system and the database. There may be times when these feature versions do not match. For example, if a patch did not run successfully to completion, it may have updated the file system, but not the database. In this case, the file system version and the database version could be different.

When you start AD Administration or adop, an information matrix scrolls on the screen. It indicates the status (Active=<Yes or No>) and version numbers of the following features: CHECKFILE, PREREQ, CONCURRENT_SESSIONS, PATCH_HIST_IN_DB, PATCH_TIMING, and SCHEMA_SWAP.


The matrix is for information only. No action is required unless the feature versions do not match. If they do not, you can use the OAM Applied Patches utility to determine which patches were applied successfully and verify the version level.

The AD Interface

Some AD utilities are designed to perform a single function. For example, you run AD Relink only to relink executables programs with the server product libraries. These utilities do not use menus or input screens. All user interaction is from the command line in the form of prompts.

However, other utilities have multiple functions, which are presented on menus or input screens. For example, when you run AD Administration, the first screen you see is the main menu.

AD Administration Main Menu



From this screen, choose one of the submenus, and then from there, choose the process you want to run.

Note: You can run the AD Utilities while an online patching cycle is in progress. The only additional step needed is to set the environment to the patch edition. An error will be raised if you try to run a utility from the run edition while a patch cycle is in progress (that is, a patch edition exists).

Command Line Arguments

You can direct the way the AD utilities operate by adding modifiers to the utility's start command. These modifiers may be in the form of arguments, flags, or options. They all refine the actions performed by a utility.

Command line arguments, flags, and options are in the "token=value" format, where token is the name of the modifier. You should enter both the argument and the value in lowercase type (the utility automatically converts the "token" portion to lowercase, but it cannot convert the "value").

For example:

```
$ adadmin LOGFILE=TEST.LOG
```

The token ("LOGFILE") will be converted to lowercase, but the value (TEST.LOG) is not recognized by the utility. The correct way to enter this command is:

```
$ adadmin logfile=test.log
```

You can enter more than one token=value argument on a single command line by separating them with one blank space as in the following adadmin command.

```
$ adadmin printdebug=y flags=hidepw
```

In some cases, you can include more than one value for a token. In this case, separate the values with commas. For example:

```
$ adadmin flags=nohidepw,trace
```

Comma-separated lists must not contain blank spaces. For example, the following command is not valid and will give an error:

```
$ adadmin flags=nohidepw, trace
```

Some arguments are used only by a specific utility. For example, adop makes extensive use of command line arguments and options that are unique to that utility: these are listed and discussed in the Patching section of this book.

Other command line arguments are used by several utilities. These are listed in the following table.

AD Utility Command Line Arguments

abandon	Description
Used by	AD Administration, adop.
Purpose	Tells AD utilities to abandon an existing non-interactive session. Can be used only when interactive=n is also specified.
Values	y or n
Default	n, meaning that the last utility run non-interactively did not successfully complete the processing.
Example	adadmin interactive=n abandon=y

AD Utility Command Line Arguments

defaultsfile	Description
Used by	AD Administration, adop, AD Controller.
Purpose	Specifies the defaults file which stores answers to interactive AD utility questions. Normally used non-interactively.
Values	A fully-qualified filename. Must be under the \$APPL_TOP/admin/<SID> directory.
Default	None, meaning that no defaults file is used.
Example	adctrl defaultsfile=/d1/apps/prodappl/admin/prod1/ prod_def.txt

AD Utility Command Line Arguments

help	Description
Used by	All AD utilities.
Purpose	Summarizes available command line options.
Values	y or n
Default	n
Example	adadmin help=y

AD Utility Command Line Arguments

interactive	Description
Used by	AD Administration, adop, AD Controller.

interactive	Description
Purpose	Tells AD utilities whether to run either interactively or non-interactively.
Values	y or n
Default	y, meaning that the utility runs interactively.
Example	adadmin interactive=n

AD Utility Command Line Arguments

localworkers	Description
Used by	AD Administration, adop.
Purpose	Specifies the number of workers to run on the primary node in a Distributed AD environment.
Values	1 to the maximum supported by your database, but not more than 999, inclusive
Default	Defaults to the value of the workers argument, which means all workers run on the primary node.
Example	adadmin workers=8 localworkers=3

AD Utility Command Line Arguments

logfile	Description
Used by	All AD Utilities.
Purpose	Tells AD utilities what log file to use. Normally used when running a utility non-interactively.

logfile	Description
Values	A file name (<i>not</i> a fully-qualified path name)
Default	None, meaning that the utility will prompt for the log file name.
Example	adctrl logfile=test.log

AD Utility Command Line Arguments

menu_option	Description
Used by	AD Administration, AD Controller.
Purpose	When running one of these utilities non-interactively, used to connect the actions in a defaults file with a specific menu item.
Values	See list of menu options in the description of these utilities. Must be used with interactive=n and defaultsfile=<name of defaults file>.
Default	N/A
Example	adctrl interactive=n defaultsfile=\$APPL_TOP/admin/prod/ctrldefs.txt menu_option=SHOW_STATUS

AD Utility Command Line Arguments

parallel_index_threshold	Description
Used by	AD Administration, adop.

parallel_index_threshold	Description
Purpose	Specifies the number blocks in a table. If a table contains fewer blocks than the threshold setting, indexes are created with parallel workers and serial DML. If the table contains more blocks than the threshold setting, indexes are created with one worker and parallel DML.
Values	0 to 2147483647; if set to 0, indexes are created with parallel workers and serial DML
Default	20000; meaning a threshold of 20,000 blocks
Example	adadmin parallel_index_threshold=15000

AD Utility Command Line Arguments

printdebug	Description
Used by	All AD Utilities.
Purpose	Tells AD programs to display extra debugging information. In some cases, the amount of extra debugging information is substantial.
Values	y or n
Default	n
Example	adadmin printdebug=y

AD Utility Command Line Arguments

restart	Description
Used by	AD Administration, adop, AD Controller.

restart	Description
Purpose	Tells AD utilities running non-interactively to restart an existing session. Only valid when interactive=n is also specified.
Values	y or n
Default	n, meaning that the utility running non-interactively will expect to run a completely new session.
Example	adadmin interactive=n restart=y

AD Utility Command Line Arguments

wait_on_failed_job	Description
Used by	AD Administration, adop.
Purpose	Directs the utilities to wait for user input in a non-interactive session when a job fails.
Values	y or n
Default	n
Example	adadmin wait_on_failed_job=yes

AD Utility Command Line Arguments

workers	Description
Used by	AD Administration, adop.
Purpose	Specifies the number of workers to run. Normally used when running the utility non-interactively.

workers	Description
Values	1 to the maximum supported by your database, but not more than 999
Default	No, meaning that the program prompts for the number of workers to run
Example	adadmin workers=8

AD Flags Argument

The flags= argument is used by all AD utilities. It passes one of several generic flags to the utility. Enter one flag or a comma-separated list of flags. The default is None.

flags= Argument Options

hidepw	Description
Default	hidepw
Purpose	Directs the utilities to either hide or show passwords in AD Utility log files.
Comments	By default, lines in an AD utility log file containing passwords are modified to hide the passwords. When nohidepw is specified, each line containing hidden passwords is followed by a corresponding line prefixed with HIDEPW;, showing the original line with passwords.
Example	adadmin flags=nohidepw

flags= Argument Options

logging	Description
Default	logging

logging	Description
Purpose	Tells the AD utility whether to create indexes using logging or nologging.
Comments	<p>Using flags=nologging when creating indexes may increase performance. However, flags=nologging makes database media recovery incomplete and does not work with standby databases.</p> <p>Logging is the default in adop to support database media recovery and standby databases. We do not recommend using flags=nologging for production systems unless you make a complete backup both before and after running adop.</p> <p>flags=nologging affects indexes created through ODF only, not SQL scripts. The XDF utility always creates indexes with logging.</p>
Example	adop flags=logging

flags= Argument Options

trace	Description
Default	notrace
Purpose	Tells the AD utility whether to log all database operations to a trace file.
Comments	<p>Database trace files created while running an AD utility may aid debugging. The flags=trace option creates multiple trace files for the AD utility and the AD workers. A new trace file is created each time the AD utility or a worker reconnects to the database.</p> <p>Note that flags=trace only traces database operations internal to the AD utility itself. Database operations in SQL scripts or external programs run by the AD utility are not recorded by flags=trace.</p>

trace	Description
Example	adadmin flags=trace

Note: Many AD utilities accept additional arguments to those listed. However, these should be used only under the explicit direction of Oracle Support.

Running AD Utilities

Important: In an online patching environment such as Oracle E-Business Suite Release 12.2, AD utilities such as AD Admin and AD Splice can be invoked from the run edition if (and only if) there is no patch edition. An error will be displayed if a user tries to run any of these utilities from the run edition and a patch edition exists.

To run AD utilities, set the environment to define the system configuration parameters. For example, a utility may require the directory path to the Applications ORACLE_HOME. This parameter, and others, make up your system environment.

Important: Before setting the environment, Windows users must also configure Windows services.

Once you have pointed the utility to the correct environment, you start it by entering the utility name.

Note: See Configuration and Environment Files, page 7-40 in this chapter.

Setting the Environment:

To set the Oracle E-Business Suite environment, complete the following steps. See the applicable Installation and Upgrade Notes for any additional platform-specific steps.

Important: Remember that there are two file systems in Release 12.2, run and patch. The environment must be set correctly on both.

1. Log in as applmgr (Applications file system owner).
2. Run the environment (UNIX) or command (Windows) file for the current APPL_

TOP and database.

UNIX:

The environment file is typically APPS<CONTEXT_NAME>.env, and is located under APPL_TOP. From a Bourne, Korn, or Bash shell, enter the following command:

```
$ . APPS<CONTEXT_NAME>.env
```

Windows:

Using either Windows Explorer or the Run option from the Start menu, enter the command:

```
%APPL_TOP%\envshell.cmd
```

This creates a command window with the required environment settings for Oracle E-Business Suite. All subsequent commands should be run in this window.

3. If you have made any changes to the environment, check that it is correctly set by entering the following commands:

UNIX:

```
$ echo $TWO_TASK  
$ echo $ORACLE_HOME  
$ echo $PATH
```

Windows:

```
C:\> echo %LOCAL%  
C:\> echo %ORACLE_HOME%  
C:\> echo %PATH%  
C:\> echo %APPL_CONFIG%
```

For UNIX, the ORACLE_HOME must be set to the proper database directory, and TWO_TASK or LOCAL must identify the correct database. For Windows, APPL_CONFIG must be set to <CONTEXT_NAME>.

4. Ensure that there is sufficient temporary disk space.
You should have at least 50 MB in the temporary directories denoted by \$APPLTMP and \$APPLPTMP (UNIX), or %APPLTMP% and %APPLPTMP% (Windows). You should also have space in the operating system's default temporary directory, which is usually /tmp or /usr/tmp (UNIX) or C:\temp (Windows).
5. If you are running an AD utility to relink or update Oracle E-Business Suite product files or modify Oracle E-Business Suite database objects, shut down the concurrent manager, Web server listeners, forms server listeners if the files are on a node that contains the associated servers. For example, if the files are on the node that contains the concurrent processing server, shut down the concurrent managers.

Note: For more information, see Administer Concurrent Managers

in *Oracle E-Business Suite System Setup Guide*.

Configuring Windows Services:

If you are running AD utilities on a Windows platform, you must first shut down all forms services, Web listener services, and concurrent manager services. In addition, you must verify that the database and database listeners are running.

To view and change the status of a service, follow these steps:

1. Select Start > Settings > Control Panel, and double-click on Services.
2. Highlight the appropriate service name and click Stop or Start as appropriate. The following table lists the services and status required when running an AD utility:

Windows Services and AD Utility Status Requirements

Service Type	Service Name	Status
Concurrent Manager Services	OracleConcMgr<CONTEXT NAME>	Stopped
Database Services	OracleService<SID>	Started
Database Listener	Oracle<SID>_<DB_VERS>R DBMSTNSListener<SID>	Started

Starting a Utility:

To start an AD utility, enter the utility's executable name on the command line. For example, to start AD Administration, you would enter the command:

```
$ adadmin
```

Note: For more information, see *Command Line Utilities*, page 7-37 in this chapter for a list of AD executables.

Exiting or Stopping a Utility:

When menu-driven utilities complete a processing task, they return you to the main menu, where you either choose another process or Exit. AD Administration is an example. Other utilities do not use a menu format. In this case, the utility exits automatically when processing is complete. *adop*, *AD Merge Patch*, and *File Character Set Converter* *File Character Set Converter* are examples.

Before it begins processing tasks, you can stop a utility by entering `abort` at any prompt. You can use this command only for utilities that display prompts, and only when a prompt is displayed on the screen.

In some cases, a utility may begin the processing actions, but quits before the actions are complete (because of an error). Or, during a parallel processing session, you may decide to stop the processing actions by shutting down the workers.

Note: For more information, see the Troubleshooting chapter in this section for additional details about shutting down and restarting workers.

Restarting a Utility:

You can restart a utility by entering the executable name on the command line. When you restart, the utility prompts you to enter a new log file, or to specify the log file from the interrupted session. When you reuse the log file from a previous session, the utility adds the message "Start of <utility name> session" to the end of the file and appends messages from the continued session as it generates them.

The utility prompts you to do one of the following:

- **Continue Session (the default)**

The utility checks the progress of the previous session in the restart files, and begins processing at the point where your last session stopped.

- **Start New Session**

The utility asks you to confirm your choice if you choose not to continue the previous session. It starts the process from the beginning.

If the process that stopped was running in parallel, a `FND_INSTALL_PROCESSES` table may exist. If it does, the utility asks if you want to drop the table. This message serves as a warning to make you aware of the existing AD session. Determine if any other utility is running in another session or on another node. If you are sure that the AD utility that is currently running is not needed, you can drop the `FND_INSTALL_PROCESSES` table and continue with the newer AD session that you started.

Note: For more information, see Restart Files, page 7-39 in this chapter.

Using Parallel Processing

Processing Tasks in Parallel

Parallel processing is typically used by AD Administration and AutoPatch to:

- Compile invalid objects.
- Run database driver tasks, such as SQL scripts.
- Generate various kinds of files, such as forms, report, and message files.

Workers complete processing tasks assigned to them by the manager. The utilities themselves determine the list of tasks to be performed and prioritize them for execution. They also prompt for the number of workers to perform the tasks. For example, when AutoPatch is applying a database driver, it creates a list of database tasks and prompts you to specify the number of workers that should run concurrently to execute these tasks.

The worker processes are instances of the adworker program. This program can only be called by the manager processes, and cannot be run stand-alone.

Managers

The manager assigns each worker a unique ID and inserts a row for each worker in the FND_INSTALL_PROCESSES table. It creates this table to serve as a staging area for job information, and as a way to communicate with the worker. Communication is accomplished using two columns: CONTROL_CODE and STATUS.

The manager updates the table with a subset of the list of jobs, one job per worker. For example, if there are five workers, then the table holds five jobs (even though there may be 100 or more jobs involved in the complete action). The manager starts the workers and uses the CONTROL_CODE and STATUS columns to assign tasks. It polls these two columns continuously, looking for updates from the workers. As a worker finishes its assignment, the manager updates each row with the next task in the list, and leaves another message for the worker.

Once all jobs are complete, the manager tells the workers to shut down, and then drops the FND_INSTALL_PROCESSES table (after it is sure all workers have actually shut down).

Workers

Each worker updates the STATUS column, giving the manager a report on its progress. As the jobs are completed, the manager updates the table with the next job in the queue, and updates the CONTROL_CODE and STATUS columns telling the worker to start processing. If there is a failure, the worker reports a failed status.

For certain tasks, some worker processes spawn other child processes that do the actual work. The spawned child process returns a status code to the worker that spawned it. The worker interprets the code to determine if the job has been completed successfully. Examples of child processes are SQL*Plus and FNDLOAD.

Deferred Jobs

The first time a job fails, the manager automatically defers the job and assigns a new

one to the worker. If the deferred job fails the second time it is run, the manager defers it again only if the total runtime of the job is less than ten minutes. If the deferred job fails a third time (or if the job's total runtime is not less than ten minutes the second time it is run) the job stays at failed status and the worker waits. At this point, you must address the cause of the failure, and then restart the job.

Note: For more information, see *Running AD Controller Interactively*, page 7-58 in this chapter.

The deferred job feature uses the AD_DEFERRED_JOBS table. This table is created when the FND_INSTALL_PROCESSES table is created, and is dropped when the FND_INSTALL_PROCESSES table is dropped.

Determining Number of Workers

The AD utilities provide a default number of workers of twice the number of CPUs on the database server. Oracle recommends you choose a number of workers between 2-4 times the number of CPUs. For example, if there are four CPUs on the database server, you should choose something in the range of 8-16 workers.

The AD utilities calculate a maximum number of workers that your database can support (up to 999). You cannot enter a number of workers greater than the database can support.

Note: In Release 12.2, AD utilities execute during runtime. Therefore, the number of available DB processes is also taken into account when calculating the number of workers. During periods of presumed high activity, estimating the requirement for DB processes can limit the number of AD workers more than the actual number of DB processes.

Worker Log Files

In addition to the information recorded in the <utility name>.log file, utilities that process jobs in parallel write details about errors to worker log files. The adwork<number>.log files (adwork001.log, adwork002.log...) reside in the \$APPL_TOP/admin/<SID>/log directory, where <SID> is the value of the ORACLE_SID or TWO_TASK variable (UNIX), or in %APPL_TOP%\admin\<SID>\log, where <SID> is the value of ORACLE_SID or LOCAL (Windows).

Concurrent requests run by AutoPatch and AD Administration create their own log files.

Note: For more information, see *Log and Output Filenames in Oracle E-Business Suite Setup Guide*.

Worker Restart Files

Restart files are used to continue processing at the point where it stopped. Each worker may also have a restart file called `adworkxxx.rf9`. These files are stored in `$APPL_TOP/admin/<SID>/restart` (UNIX) or in `%APPL_TOP%\admin\<SID> \restart` (Windows). The worker creates the restart file when the manager assigns it a job, and deletes the restart file when it finishes the job.

Caution: Do not modify or delete any manager or worker restart files unless explicitly told to do so by Oracle Support.

The Troubleshooting chapter in this section discusses various error situations when running a utility and how to resolve them.

Parallel Support for Data Manipulation Language (DML)

To reduce downtime when creating indexes, the `parallel_index_threshold` argument for AD utilities is set to a default value of 20,000. This means that if a table contains less than 20,000 blocks, the AD utilities create indexes with parallel workers and serial DML (just as in earlier releases). If a table contains 20,000 blocks or more, indexes are now created with only one worker and parallel DML. You can adjust this threshold value by specifying the `parallel_index_threshold` argument on the AD utility command line.

Monitoring and Controlling Parallel Processes

AD sessions that use parallel processing may run to completion without user intervention. However, it is often useful to determine how many jobs have been completed or whether processing has stopped for some reason. AD Controller is a utility that you can use to determine the status of AD Administration or AutoPatch workers and to control their actions. You can run AD Controller interactively or non-interactively. It must be run in its own window, not in the same window as AD Administration or AutoPatch.

Note: For more information, see Interactive and Non-Interactive Processing, page 7-38 in this chapter.

You choose options that display worker status, restart workers, or issue commands to the manager from the AD Controller main menu.

Running AD Controller Interactively

Follow these steps to access AD Controller.

1. Log in as `applmgr` and set the environment as described in Setting the Environment, page 7-52 in this chapter.

Important: AD Controller (adctrl) can only be run from the **patch edition** of the file system. Attempting to run it from the run edition will give an error. You can identify which edition you are in by checking the value of the FILE_EDITION environment variable.

2. Start AD Controller with the adctrl command. This will prompt you to:
 - Confirm the value of APPL_TOP.
 - Specify an AD Controller log file (the default is adctrl.log). The AD Controller log file is written in the current working directory.
 - Supply the Oracle Application Object Library user name and password.

3. Choose an option from the main menu.

Once you respond to the prompts, the main menu appears.

AD Controller Menu

```
AD Controller Menu
-----
1. Show worker status
2. Tell worker to restart a failed job
3. Tell worker to quit
4. Tell manager that a worker failed its job
5. Tell manager that a worker acknowledges quit
6. Restart a worker on the current machine
7. Exit

Enter your choice [1] : 1
```

Type a number to select an option. Press [Return] at any time to return to the AD Controller main menu.

Note: See the Troubleshooting Applications DBA Operations chapter in this book for instructions on using each menu option.

Running AD Controller Non-Interactively

You can run AD Controller without user intervention by creating a defaults file, which captures information you supply at the interactive prompts in a file that you can later use to run AD Controller without user intervention. Creating a defaults file and running AD Controller non-interactively works in much the same way as it does for AD Administration.

Note: For more information, see Scheduling Non-Interactive Maintenance in this book.

Like AD Administration, the same defaults file can be used to run different AD Controller commands: a single file can contain all your choices for the different menu options. In order to choose which task the defaults file will run, you add `menu_option=<menu choice>` to the utility start command. This overrides any menu-specific key stroke information stored in the defaults file initially, and allows you to use the defaults file for any of the AD Controller menu items. It also ensures that the menu option you intended for the defaults file is always valid, even if the menu items are renumbered or relocated in subsequent releases.

The available options are listed in the following table.

AD Controller Menu Options

Menu Option	Effect
ACKNOWLEDGE_QUIT	Tell manager that a worker acknowledges quit
INFORM_FAILURE	Tell manager that a worker failed its job
RESTART_JOB	Tell worker to restart a failed job
SHOW_STATUS	Show worker status
SHUTDOWN_WORKER	Tell worker to quit
START_WORKER	Restart a worker on the current machine

Note: The menu options for running AD Administration are listed in the Preparing for Non-Interactive Processing, page 7-62 section of this book.

The following is an example of running AD Controller non-interactively to show worker status:

```
$ adctrl interactive=n defaults_file=$APPL_TOP/admin/prod/ctrldefs.txt \
logfile=adctr.log menu_option=SHOW_STATUS
```

Using any menu option on the command line, except for SHOW_STATUS, requires that you also use the worker_range=<range> option. See the AD Controller command line help for details.

Distributing Processing Tasks Across Nodes

AD uses its existing manager-worker job system employed in parallel processing to include Distributed AD. This parallel processing feature allows workers in the same AD session to be started on multiple application tier servers to utilize all available resources. Because the AD workers create and update file system objects, as well as database objects, Distributed AD must be used only on systems that are using a shared application tier file system to ensure the files are created in a single, centralized location.

While running either AD Administration or AutoPatch on the primary node, you start an AD Controller session from any of the nodes in the shared application tier file system environment to perform any standard AD Controller operation, using both local and non-local workers.

Note: For more information, see the Distribute Processing with Distributed AD, page 7-21 of this book.

About System Maintenance

After your system is installed, it will be necessary to perform certain maintenance tasks to keep it running smoothly. For example, you will generate form files, maintain snapshot information, relink executables, compile or validate the APPS schema, and so on. Some tasks are routine and should be performed on a regular basis. Other tasks are non-routine and generally performed infrequently.

You run maintenance tasks from the command line using AD Administration. Once you start this utility, it presents the tasks in menu form, grouped generally by type of activity you will perform. For example, the tasks associated with compiling and reloading Applications database entities are grouped on the same menu.

In addition to the AD Administration maintenance tasks, this chapter describes AD Relink, a command line utility used to relink AD executables.

Important: You *cannot* relink AD utilities executables using AD Administration.

AD Administration Overview

AD Administration manages most of the maintenance tasks required for your Oracle E-Business Suite system. Currently, these maintenance tasks are grouped by types on the AD Administration main menu.

When you start AD Administration from the command line, it prompts you for the basic system-specific information it needs. For example, you need to supply a name for the log file where processing actions and error messages will be recorded.

Note: For more information, see Prompts, page 7-62 in Chapter 1.

Once you respond to these prompts, AD Administration displays the main menu, which serves as the gateway to various submenus where you select the individual maintenance tasks. For example, on the Generate Applications Files menu, you can run tasks that generate message files, forms files, report files, message files, or product JAR files. These submenu tasks may also require you to respond to prompts to collect task-specific information. For example, some tasks require you to enter the number of workers you want to employ to process the jobs associated with the task.

Note: For more information, see Processing Tasks in Parallel, page 7-55 in Chapter 1.

When you respond to AD Administration prompts, you are running the utility interactively. However, like AutoPatch and AD Controller, you can also run AD Administration non-interactively, specifying a previously created defaults file that contains the information necessary to run a specific maintenance task without user intervention.

Note: For more information, see Interactive and Non-Interactive Processing, page 7-38 in Chapter 1.

Prompts

In addition to the basic prompts described in Chapter 1, AD Administration may require additional information that is specific to one of the submenu tasks. If so, it displays additional prompts. For example, when running the Generate Product JAR files task from the Generate Applications Files menu, AD Administration prompts you as follows:

```
Do you wish to force generation of all jar files? [No]:
```

The task-specific prompts are described more fully in the discussion of each task.

Preparing for Non-Interactive Processing

The discussion of command line prompts assumes you are running AD Administration interactively. You respond to the standard prompts and those required for specific tasks you choose from the AD main menu and submenus. AD Administration can also run some tasks non-interactively by using the information you store in a defaults file, instead of requiring you to respond to prompts.

Note: For more information, see *Interactive and Non-Interactive Processing*, page 7-38 in this book.

Specifying a Menu Option in the AD Administration Defaults File

The same defaults file can be used to run different AD Administration tasks a single file can contain all your choices for the different menu options. In order to choose which task the defaults file will run, you add `menu_option= <menu choice>` to the utility start command. This overrides any menu-specific key stroke information stored in the defaults file initially, and allows you to use the defaults file for any of the AD Administration menu items. It also ensures that the menu option you intended for the defaults file is always valid, even if the menu items are renumbered or relocated in subsequent releases.

Defaults File menu_option Values

menu_option Value	Corresponding AD Administration Menu Choice
GEN_MESSAGES	Generate message files
GEN_FORMS	Generate form files
GEN_REPORTS	Generate reports files
GEN_JARS	Generate product JAR files
RELINK	Relink Applications programs
COPY_FILES	Copy files to destinations
CONVERT_CHARSET	Convert character set
SCAN_APPLTOP	Scan the APPL_TOP for exceptions

menu_option Value	Corresponding AD Administration Menu Choice
SCAN_CUSTOM_DIR	Scan a CUSTOM directory for exceptions
LIST_SNAPSHOT	List snapshots
UPDATE_CURRENT_VIEW	Update current view snapshot
CREATE_SNAPSHOT	Create named snapshot
EXPORT_SNAPSHOT	Export snapshot to file
IMPORT_SNAPSHOT	Import snapshot from file
DELETE_SNAPSHOT	Delete named snapshot
CHECK_FILES	Check for missing files
CMP_INVALID	Compile APPS schema
CMP_MENU	Compile menu information
CMP_FLEXFIELDS	Compile flexfield data in AOL tables
RELOAD_JARS	Reload JAR files to database
VALIDATE_APPS	Validate APPS schema
CREATE_GRANTS	Recreate grants and synonyms for APPS schema
MAINTAIN_MLS	Maintain multi-lingual tables
CHECK_DUAL	Check DUAL table

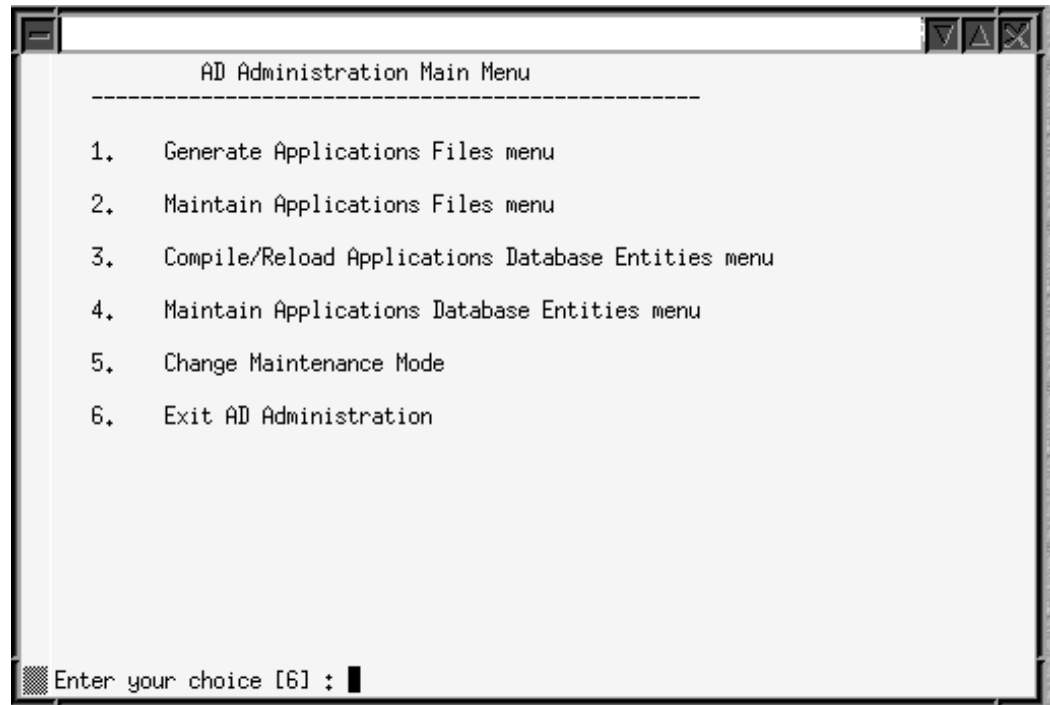
The AD Administration Interface

You start AD Administration from the command line. However, all maintenance tasks are initiated from the AD Administration Main Menu. This section describes some of the common features used to run this utility.

Main Menu

After you start AD Administration and respond to the prompts, the AD Administration Main Menu appears.

AD Administration Main Menu



This menu displays the submenus where the individual maintenance tasks are grouped. To choose a submenu, type the number of the menu at the prompt. To exit AD Administration, press [Return].

Available Options

Depending on your system configuration, the submenus for AD Administration may show slightly different option names and numbers from the ones displayed here.

Running AD Administration Interactively

Complete the steps in this section to display the AD Administration Main Menu and access the submenus and the maintenance tasks.

1. Set the environment.

You must set the environment in order to apply the environment variables that define your system. This task is common to many AD utilities. See *Setting the Environment*, page 7-52 in this book for the preparatory steps.

2. From any directory, start AD Administration with this command:

```
$ adadmin
```

The utility starts and displays the first prompt.

3. Respond to prompts.

Supply the information requested by the AD Administration prompts. Prompts unique to an option are described with the option.

When you complete the prompts, the Main Menu appears.

4. Choose maintenance tasks.

On the Main Menu, choose a submenu. The submenus and the options they display are described fully beginning with Generate Applications Files in the next section.

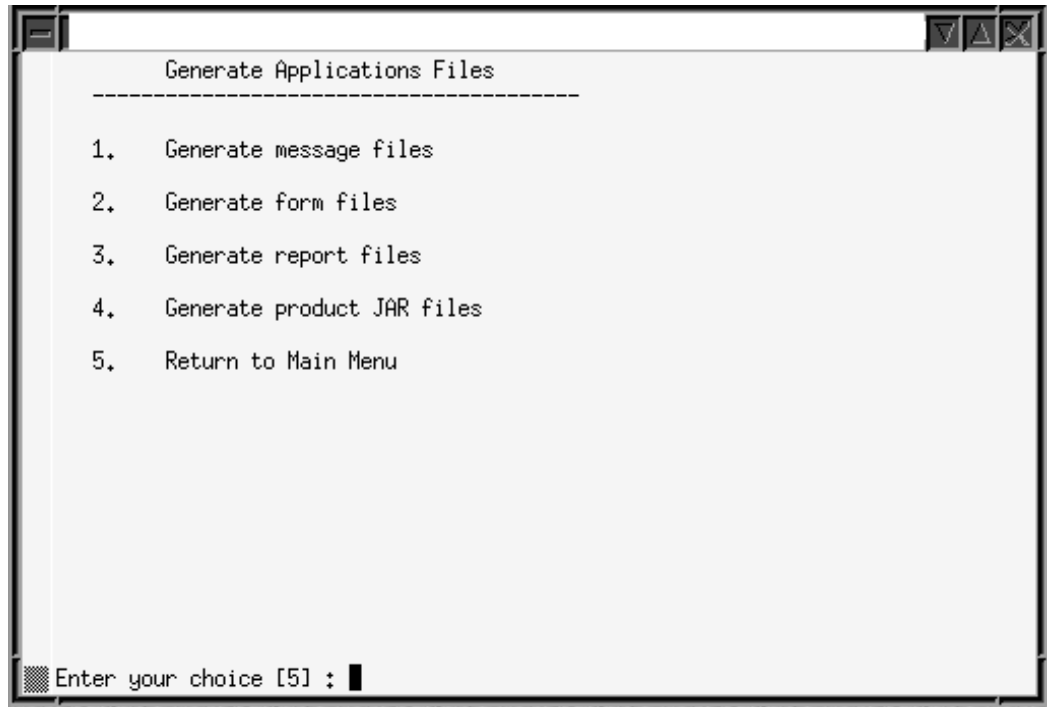
5. Exit AD Administration.

You can exit AD Administration from the Main Menu by choosing option 6 (Exit AD Administration) at the screen prompt. You can also choose to exit the utility at any prompt by typing `abort` on the command line. See *Restart Files*, page 7-39 in this book for information about restarting AD utilities after using the `abort` command.

Generating Applications Files

You may need to generate Applications files from time to time during your Applications life cycle. You access the associated tasks from the Generate Applications Files menu.

Generate Applications Files Menu



If system users are having difficulty accessing messages, forms, or reports, you may be able to resolve the issue by generating the associated files. Or, when you apply a patch that adds or changes product functionality, you may want to generate the associated files after you apply the patch, instead of running the generate driver during the patching downtime. The Generate Files tasks may be performed on any server, as required.

You do not have to shut down your system to generate files. However, users that access the files being generated (for example, for Human Resources forms) must log off.

Generate Message Files

Oracle E-Business Suite uses files to display messages. This task generates binary message files (extension .msb) from Oracle Application Object Library tables.

Caution: Run this task only when instructed to do so in a patch readme file, or by Oracle Support Services.

Generate Form and Report Files

These activities are carried out in much the same way.

- **Generate forms files**

Generates executable Oracle forms files (extension .fmx) from the binary forms definition files (extension .fmb). The definition files are located under AU_TOP, and the executable files are stored under each product's directory.

- **Generate report files**

Generates the binary Oracle Reports report files (extension .rdf).

The prompts and behavior work in similar fashion, except as noted:

- Ask for the number of workers and generate selected objects for selected products in parallel
- Display the current character set (from NLS_LANG) and ask if you want to generate form or report objects in this character set
- Ask if you want to regenerate Oracle Forms PL/SQL library files, menu files, and executable files (forms files only)
- Ask for the products associated with the form or report objects
- Ask if you want to generate specific form or report objects for each selected product
- Display the current set of installed languages and ask if you want to generate form or report files in these languages
- Create a list of all objects to generate
- Display the list of objects to be generated (specific objects or all objects)

Generate Product JAR Files

Generate Java archive (JAR) files whenever you upgrade the Oracle Developer technology stack, or when advised by Oracle Support Services. This task signs JAR files (if on a Web server) and also does the following:

- Generates product JAR files in JAVA_TOP and copies them to APPL_TOP
- Generates other Java-related files under APPL_TOP and JAVA_TOP
- Recreates Java libraries (appsborg.zip and appsborg2.zip) under APPL_TOP and JAVA_TOP

When you run the task, it prompts:

```
Do you wish to force generation of all jar files? [No]
```

If you choose No, it generates only JAR files that are missing or out-of-date. If you choose Yes, all JAR files are generated (more time-consuming).

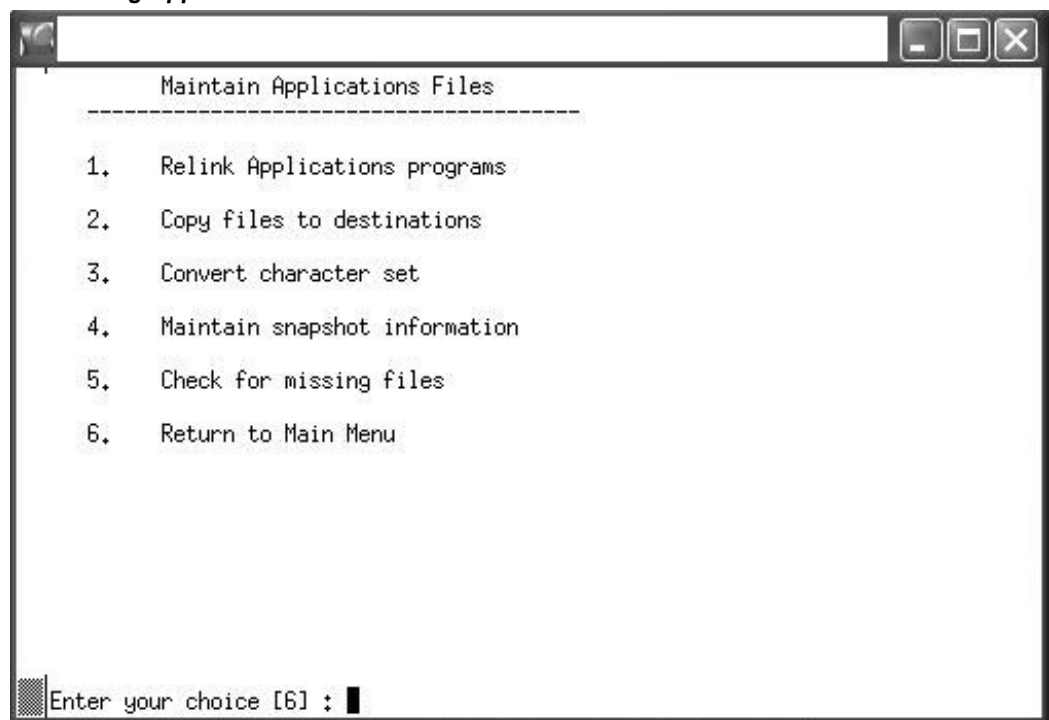
If AD Administration displays a list of warnings or errors and objects that did not

generate successfully and asks if you want to continue as if successful, review the log file to determine if the problems require attention. If you choose not to continue and restart your session at a later time, AD Administration attempts to regenerate only the files that did not generate successfully.

Maintaining Applications Files

Certain maintenance tasks are required to keep your Applications files up to date. For example, you may need to copy product files to a central location or convert files in the APPL_TOP to another character set. These tasks are grouped on the Maintain Applications Files menu.

Maintaining Applications Files Menu



You can run any of these tasks by choosing it from this menu.

Relink Applications Programs

Relinks Oracle E-Business Suite executable programs with the Oracle server libraries so that they function with the Oracle database. For each product, you can choose whether to link all executables or specific ones only.

The default is to relink without debug information. You should use the debug option only when requested to do so by Oracle Support Services.

Important: AD Administration cannot be used to link executables for the AD products themselves. You must use AD Relink for this. See Relinking AD Executables in this chapter.

Copy Files to Destinations

Copies files from each product area to central locations where they can be easily referenced by non-Applications programs. This option uses revision-based copy logic to ensure that the destination file versions are the same as, or higher than, the source file versions.

Oracle recommends that you do not use the force option to overwrite existing files, unless so instructed by Oracle Support Services. Copying files with this option updates all JAR files, resulting in them all being downloaded to each client again and causing runtime performance degradation.

The file types and their respective destinations are shown in the following table:

Copy Files to Destinations Summary

These files are copied to: (UNIX)	... are copied to: (Windows)
Java files	\$JAVA_TOP	%JAVA_TOP%
HTML files	\$OAH_TOP	%OAH_TOP%
Media files	\$OAM_TOP	%OAM_TOP%

The directories for the variables are specified in the adovars.env file (UNIX) or the adovars.cmd file (Windows).

When this option is used to copy reports files, the default destination is under AU_TOP.

Convert Character Set

Prepares the files in the APPL_TOP for conversion to another character set, and then performs the conversion.

Note: For more information, see Globalization Support in *Oracle E-Business Suite Concepts*.

When you choose this option, AD Administration presents another submenu, which contains options for scanning your files in preparation for the conversion. The scan searches for exceptions - files that will have incomplete (lossy) conversions - so that you

can fix potential problems before you actually convert the character set. Choose one of the following scan options.

Tip: Always verify the compatibility of the database character set before converting the APPL_TOP character set.

The options are:

1. Scan the APPL_TOP for exceptions.

Scans the APPL_TOP and creates three files in the admin\<SID>\out directory.

Scan APPL_TOP for Exceptions Output Files

File	Contents
admanifest_excp.lst	Lists files that will not be converted because of lossy conversion.
admanifest.lst	Lists files that can be converted.
admanifest_lossy.lst	Lists files with lossy conversions, including line by line detail.

Review the files listed in admanifest_excp.lst. Fix files that report lossy conversion before you convert the character set. Repeat this task until there are no entries in admanifest_excp.lst. If you need to see more detail, review admanifest_lossy.lst.

2. Scan a CUSTOM directory for exceptions.

Collects the same information as the first task, but scans custom Applications directories rather than the APPL_TOP directory.

Note: With this option, adadmin may list additional files (such as .rdf, .doc, and .zip) as exceptions in admanifest_excp.lst. This is because the CUSTOM directory can be modified by users, so the file extension is not enough for adadmin to determine whether a file can be successfully converted. In contrast, users cannot modify the files under \$APPL_TOP, so the file extensions there are a reliable guide to whether a file can be successfully converted.

3. Convert character set.

Run this task only if admanifest_excp.lst has no entries. It prompts you for the manifest file (admanifest.lst) created when you ran the scan option(s).

The utility backs up the product source files and the APPL_TOP/admin source files. It saves product files in the <PROD>_TOP directories in the format <prod>_s_<char_set>.zip. It saves admin source files in the APPL_TOP/admin directory in the format admin_s_<char_set>.zip

Maintain Snapshot Information

There are two types of snapshots: *APPL_TOP snapshots* and *global snapshots*. An APPL_TOP snapshot lists patches and versions of files in the APPL_TOP. A global snapshot lists patches and latest versions of files in the entire Applications system (that is, across all APPL_TOPs).

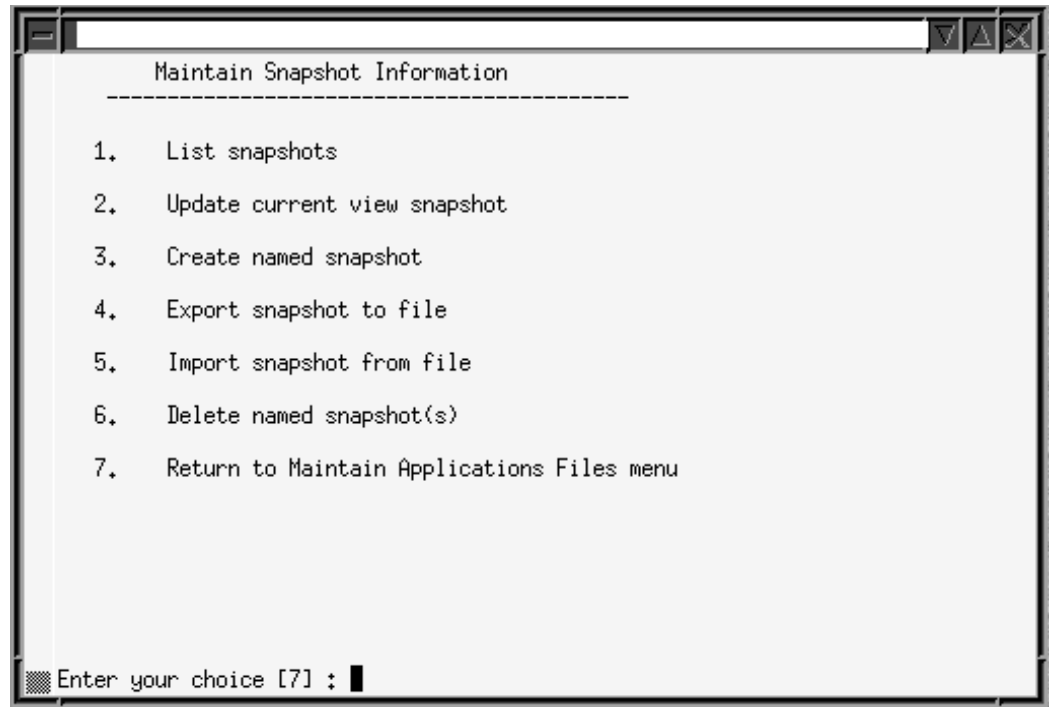
Both APPL_TOP snapshots and global snapshots may be either *current view snapshots* or *named view snapshots*. A current view snapshot is created once and updated when appropriate to maintain a consistent view. A partial view snapshot allows you to synchronize only selected files from a current view. A named view snapshot is a copy of the current view snapshot at a particular time (not necessarily the latest current view snapshot), and is not updated.

Patch Wizard uses the information contained in the global current view snapshot to determine which patches have already been applied. AutoPatch uses the APPL_TOP current view snapshot to determine if all prerequisite patches have been applied to that APPL_TOP. Snapshot information is stored in the AD_SNAPSHOTS, AD_SNAPSHOT_FILES, and AD_SNAPSHOT_BUGFIXES tables.

During a new installation, Rapid Install creates a current snapshot as a baseline. Each time you run AutoPatch, it automatically creates a new (updated) snapshot so that the information is current as of the application of the patch.

Snapshot information maintenance is performed by choosing Maintain Snapshot Information from the Maintain Applications Files menu, and then selecting the required option.

Maintain Snapshot Information Main Menu



These options allow you to:

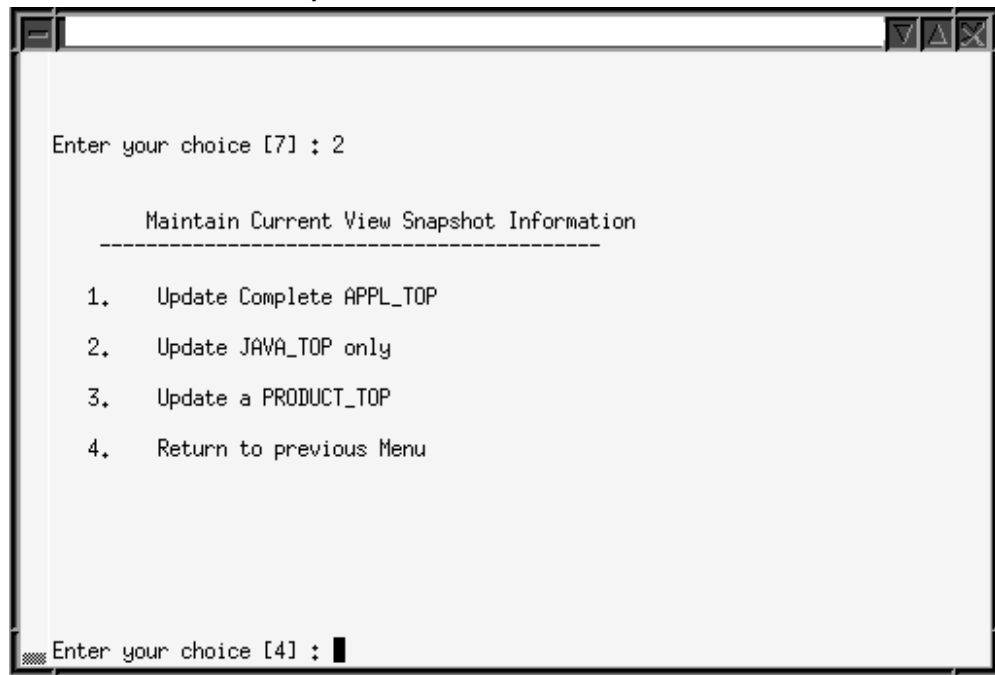
- List snapshots stored in the system
- Update a current view snapshot (full or partial APPL_TOP and global)
- Create a named snapshot (you select a current view snapshot to copy and name)
- Export a snapshot to file (you select a snapshot to export to a text file)
- Import a snapshot from a text file (you select a snapshot to import from a text file)
- Delete a named snapshot

Maintain Current View Snapshot Information

When you maintain a current view snapshot, you can choose to synchronize selected files (to maintain a partial snapshot), instead of synchronizing all files for the entire APPL_TOP. Use this option when you have copied only a few files to the APPL_TOP.

1. Select the Update Current View Snapshot option from the Maintain Snapshot Information menu.

Maintain Current View Snapshot Information Menu



2. From the Maintain Current View Snapshot Information menu, you can select one of the following options:

- **Update Complete APPL_TOP**

This is the original functionality of the Update Current View Snapshot option. It synchronizes all the files in your APPL_TOP.

- **Update JAVA_TOP only**

Synchronizes only the files in the JAVA_TOP. At the prompt, enter the path to the JAVA_TOP subdirectory where the files were copied. If the files were copied to more than one directory, press *Enter*. AD Administration scans the entire JAVA_TOP and updates the information in both the current view and the global view snapshots.

- **Update a <PRODUCT>_TOP**

Synchronizes only the files in a specific <PRODUCT>_TOP. Enter the product abbreviation, then provide the subdirectory information at the prompt.

Enter the path to a single subdirectory in the <PRODUCT>_TOP. If the files were copied to more than one directory in the <PRODUCT>_TOP, press *Enter*. AD Administration scans the entire <PRODUCT>_TOP and updates the information in both the current and the global view snapshots.

Check for Missing Files

Verifies that all files needed to run Oracle E-Business Suite for the current configuration are in the current APPL_TOP. Choose this task if you suspect there are files missing in your APPL_TOP.


Managing Database Entities

Database entities are database objects or data in the database related to Oracle E-Business Suite. Tasks for managing entities are grouped into two options on the AD Administration Main Menu: one for compiling or reloading entities and one for verifying their integrity.

Compiling or Reloading Database Entities

To compile or reload database entities, choose the Compile/Reload Applications Database Entities Menu option from the AD Administration Main Menu.

Compile/Reload Applications Database Entities Menu



You run the tasks on this menu any time you need to compile or reload database objects; for example, after you upload new menu entries, or apply a patch that changes the setup of flexfields. Run these tasks only on the node where the core AD technology directories are located.

Compile APPS schema

Spawns parallel workers to compile invalid database objects in the APPS schema.

Note: For more information, see *Compiling Invalid Objects* in this chapter.

Compile Menu Information

Compiles menu data structures. Choose this task after you have uploaded menu entries to the FND_MENU_ENTRIES table, or if Compile Security concurrent requests submitted from the Menus form (after changing menu entries) fail for any reason.

AD Administration asks if you want to force compilation of all menus. If you choose the default (No), only menus with changes are compiled. If you enter Yes, all menus are compiled. Compiling all menus is generally not required.

Compile Flexfields

Compiles flexfield data structures in Oracle Application Object Library (FND) tables. Choose this task after you apply a patch that changes the setup of flexfields. Patches usually indicate when you should perform this step

Flexfields automatically compile data when you use them for the first time, so running this task is generally not required. However, compiling flexfields at a specific time can alleviate potential runtime performance issues. For example, you may choose to compile them when system usage is known to be low, rather than automatically on first use.

Reload JAR files to Database

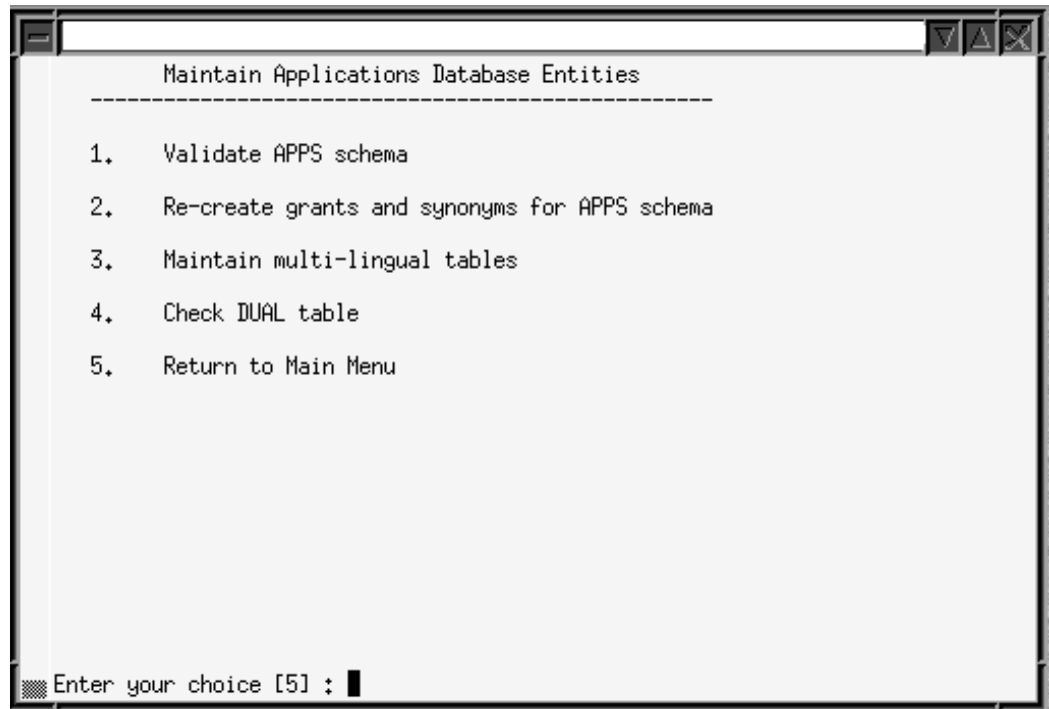
Reloads all appropriate Oracle E-Business Suite JAR files into the database. Choose this task if all Oracle E-Business Suite Java classes are removed from your database; for example, if the database Java Virtual Machine (JVM) is reloaded because of database corruption.

Maintaining Applications Database Entities

During normal system use, the integrity of your database can be compromised, for example through user error or after you apply an inappropriate patch. It is advisable to verify the integrity of database entities as a regular maintenance procedure, or whenever the behavior of your system indicates that database entities may have been corrupted.

To perform these maintenance tasks, select the Maintain Applications Database Entities Menu option from the AD Administration Main Menu.

Maintain Applications Database Entities Menu



Some tasks on this menu report on issues, or potential issues, with database entities, and others actually remedy the issues.

Validate APPS schema

Verifies the integrity of the APPS schema. It produces a report named <APPS schema name>.lst. that lists issues and potential issues, grouped by the action required:

- Issues you *must* fix (not specific to the APPS schema)
- Issues you *must* fix (specific to the APPS schema)
- Issues you may want to address (specific to the APPS schema)

The report is located in \$APPL_TOP/admin/<SID>/out (UNIX), where <SID> is the value of the ORACLE_SID or TWO_TASK variable, or in %APPL_TOP%\admin\<SID>\out (Windows), where <SID> is the value of the LOCAL variable. Each section of the file contains instructions for resolving the issues that are listed. Most issues can be fixed by either compiling invalid database objects or recreating grants and synonyms.

Recreate Grants and Synonyms for APPS Schema

This task recreates grants and synonyms for the Oracle E-Business Suite public schema (APPLSYS PUB), recreates grants on some packages from SYSTEM to APPS, and spawns

parallel workers to recreate grants and synonyms linking sequences and tables in the base schemas to the APPS schema.

Typically, you run this task after the Validate APPS schema task has reported issues with missing grants and synonyms.

Maintain Multi-Lingual Tables

Run this task after you add a language. It prompts you for the number of workers, then updates all multilingual tables.

Check DUAL Table

Some Oracle E-Business Suite products must access the DUAL table. It must exist in the SYS schema and contain exactly one row. This task verifies the existence of this table and the single row.

Important: If the DUAL table does not exist, or if it does not contain only one row, the Oracle E-Business Suite products that access it will fail to operate correctly.

Using AD Relink

You use AD Relink to relink AD executables with the Oracle server product libraries when required, to ensure they will keep functioning properly with the Oracle database.

Note: AD executables can *only* be relinked using AD Relink. This is in contrast with other product executables, which are relinked using the Relink Applications Executables task on the AD Administration Maintain Applications Files menu.

Relinking AD Executables

A number of options are available to provide extra control over the relinking process. For example, you can relink multiple AD executables in one operation.

Log Files:

As you run AD Relink, it creates a log file (adrelink.log) where it records errors and messages. AD Relink appends information about the latest relink action to the end of the file. This file is located in APPL_TOP/admin/log. If an error occurs while you are using AD Relink, or if you are not sure that the relinking was successful, review this file to see what issues should be fixed.

Relinking errors encountered during an AD Administration or an AutoPatch session are

recorded in the main log files for those utilities. See Log Files, page 7-39.

A new log file is created each time AD Relink runs. To recover disk space, or just as good housekeeping practice, you can delete `adrelink.log` files that are no longer needed (for example, after a relink operation has been found to have completed successfully).

Command Line Arguments:

You can modify or refine the operation of AD Relink with the command line arguments in the following table.

AD Relink Command Line Arguments

force	Description
Purpose	Indicates which executable programs to relink
Values	<ul style="list-style-type: none">• <code>n</code>, (relink only if the libraries or object files are more recent than the current executable program)• <code>y</code> (relink regardless of the status of the libraries or object files)
Default	none (must enter either <code>y</code> or <code>n</code>)
Example	<code>adrelink force=n</code>

AD Relink Command Line Arguments

backup_mode	Description
Purpose	Indicates whether you want to back up executables
Values	<ul style="list-style-type: none">• <code>none</code> (do not back up any executables)• <code>all</code> (back up all executables)• <code>file</code> (back up files according to instructions in <code>adlinkbk.txt</code>)

backup_mode	Description
Default	backup_mode=file
Example	adrelink force=n backup_mode=all

Note: These command line arguments are intended for use with the AD Relink utility only.

Files that are critical to running Oracle E-Business Suite are listed in the adlinkbk.txt file, which is located in APPL_TOP/admin. Using the backup_mode=file argument directs AD Relink to back up only these files.

The AD Relink Interface:

You run AD Relink from the command line. It does not use menus or input screens.

Running AD Relink:

Run AD Relink as follows.

1. Set the environment.

You must set the environment to indicate the location of the configuration parameters that define your system. This task is common to many AD utilities. See *Setting the Environment*, page 7-52 in this book for the basic steps.

2. Relink files.

Run AD Relink with the appropriate command for your operating system:

UNIX

```
S adrelink.sh force=n "ad <executable name>"
```

Windows

```
C:\> sh adrelink.sh force=n "ad <executable name>"
```

Applications DBA Reporting and Tracking Tasks

Timing Information

Use the reports in this section to gather job timing statistics.

Requirement:

How can I monitor the time it takes to complete individual system patching or maintenance sessions?

Discussion:

When you run AD Administration or AutoPatch, they automatically capture timing information about processing sessions that run parallel tasks.

Stored in database tables, this information can be accessed as the *AD Job Timing Report*. You can access the content of this report from either the OAM Timing Reports feature or the command line.

The information captured includes:

- Jobs run successfully on the first try
- Failed jobs that were restarted and then run successfully
- Failed jobs that were skipped
- Time-consuming jobs
- Job timing information
- Summary information about each parallel phase

- Overall timings for each session
- Status of an in-progress patching session

Actions:

Information about timing sessions for both AD Administration and adop is collected in a single action and can be viewed in Oracle Applications Manager.

Note: For more information about this interface and its contents, see adop Timing Details, page 5-18 and AD Administration Timing Details, page 5-21 in the Patching section of this book.

General System Reporting

Use the reports in this section to gather general system statistics and status information.

Installed Configuration Information:**Requirement**

How can I view information such as undo information, list of operating units, or NLS init.ora settings?

Discussion

The AD Configuration script (adutconf.sql) is an SQL script that reports on the configuration of an Oracle E-Business Suite system. You can use the report (adutconf.lst) in troubleshooting, or simply to document the status of your installation. For example, it contains information about undo segments, registered products and schemas, Reporting Currencies settings, and NLS database initialization parameters.

Actions

1. Log in as applmgr and set the environment as described in Setting the Environment in Running AD Utilities, page 7-52.
2. Use the command for your platform to run the script. The output file is written to adutconf.lst in the current working directory.

UNIX

```
$ cd $APPL_TOP/admin/<SID>/out
$ sqlplus <APPS schema username>/<APPS schema password> \
@$AD_TOP/sql/adutconf.sql
```

Identifying File Versions and Translation Levels:

Requirement

I want to obtain information about file versions.

Discussion

When collecting information about your system, perhaps for sending to Oracle Support, you may need to determine the version and translation levels of your files.

Actions

You can obtain version and translation levels of your files by running AD File Identification (`adident`).

Oracle E-Business Suite Reporting Tools

As you use your Oracle E-Business Suite system, you perform maintenance tasks that modify and enhance your system. Oracle E-Business Suite includes tools that enable you to create numerous reports about system status. For example, you can generate a report about the version and translation level of your files. You can also generate reports that contain statistics about how many maintenance sessions are complete, number of jobs in each session, and the time it took to complete the session and individual jobs.

Note: Many of the Oracle E-Business Suite reporting capabilities are related to patching. For details of those reports, see the Patching section of this book.

Categories of Reporting Tool

The AD reporting utilities introduced here are described in more detail later in this section. They are all run from the command line.

AD Job Timing Report

Produced automatically by AutoPatch and AD Administration to report on long-running processes, this report can be run manually from the command line to provide summary information about AD utility sessions.

AD Configuration Report

This report contains information about the installed configuration of Oracle E-Business Suite, including product group information, whether Multi-Org or MRC functionality is installed, base language and other installed languages, and so on.

AD File Identification Report

This report identifies the version and translation level of Oracle E-Business Suite files.

AD Job Timing Report

When you run AutoPatch or AD Administration, they automatically generate an AD Job Timing report (adt<session_id>.lst) that shows how long it takes to complete a parallel processing session, and provides information about the actions of workers as they process jobs during the session. These reports include timing statistics for the entire session, the phases in the session (AD Administration does not group jobs by phases), and individual jobs.

At any time during an AutoPatch or an AD Administration session, you can run a script to create an AD Job Timing report that shows the progress of the current session. Or you can go to the APPL_TOP/admin/<SID>/out directory to view an adt<session_id>.lst report from a previous session.

For AutoPatch and AD Administration sessions, the adt<session_id>.lst report is very similar to the web-based Timing Report you can access via Oracle Applications Manager.

AD Job Timing Report Interface:

You can view job timing statistics from the Timing Reports page in Oracle Applications Manager. You can also run the AD Job Timing Report for AD Administration jobs from the command line. There are no menus or input screens.

Running AD Job Timing Report:

1. Set the environment.

Set the environment in order to apply the environment variables that define your system. This task is common to many AD utilities. See *Setting the Environment*, page 7-52 in this book for the basic steps.

2. Run AD Job Timing report.

Run the report with this command, where <session_id> is the session of the timing statistics you want to see, and <output file> is the name of the file where the statistics will be written.

UNIX

```
$ cd $APPL_TOP/admin/<SID>/out
$ sqlplus <APPS username>/<APPS password>
@$AD_TOP/admin/sql/adtimrpt.sql \
<session id> <output file>
```

AD Configuration Report

The AD Configuration utility is a SQL script that reports standard information about the installed configuration of Oracle E-Business Suite. Run this task in order to debug or document the status of your installation. Running AD Configuration generates a report file (adutconf.lst) that contains the following:

- SQL*Plus PAUSE and NEWPAGE settings
- Undo information
- Information about the product group
- Whether Multi-Org is installed, and list of operating units
- Whether Multiple Reporting Currencies (MRC) functionality is installed
- List of registered products
- Information on all registered schemas
- Information about all installed products, including shared and dependent products
- Status of localization modules
- NLS init.ora settings

AD Job Configuration Report Interface:

You run AD Configuration and supply the information it needs from the command line. There are no menus or input screens.

Running AD Configuration Report:

1. Set the environment.

Set the environment in order to apply the environment variables that define your system. This task is common to many AD utilities. See *Setting the Environment*, page 7-52 in this book for the basic steps.

2. Run AD Configuration report.

Use the following commands. The report output file is written to adutconf.lst in the current working directory.

UNIX

```
$ cd $APPL_TOP/admin/<SID>/out
$ sqlplus <APPS schema username>/<APPS schema password> \
@$AD_TOP/sql/adutconf.sql
```

AD File Identification Report

The AD File Identification utility creates a report that identifies the version and translation level of Oracle E-Business Suite files. It is useful when collecting information about your site for Oracle Support Services.

AD File Identification Report Interface:

You run AD File Identification and supply the information it needs from the command line. There are no menus or input screens.

Running AD File Identification:

Run this utility as follows.

1. Set the environment.

Set the environment in order to apply the environment variables that define your system. This task is common to many AD utilities. See *Setting the Environment*, page 7-52 for the basic steps.

2. Run AD File Identification.

Use the following commands. The output is displayed on the screen.

UNIX

```
$ addident Header <file 1> [ <file 2> <file 3> ... ]
```

AD Check Digest

The AD Check Digest utility checks the integrity of downloaded patches. Oracle provides MD5 and SHA-1 digests for each Oracle E-Business Suite patch. The MD5 digest is a 128-bit string output that uniquely identifies the patch and the SHA-1 is a 160-bit string output. The patch digests are viewable from the My Oracle Support download page for a particular patch. Use AD Check Digest to verify whether the computed digests for the downloaded patch match the digests published on My Oracle Support.

AD Check Digest Interface:

You run AD Check Digest and supply the information it needs from the command line. There are no menus or input screens.

Parameters:

The following parameters are used for running AD Check Digest.

AD Check Digest Parameters

Parameter	Meaning
-file	<p>Patch file name and path. This parameter is required.</p> <p>When the <code>-file</code> parameter is specified without the <code>-md5</code> and <code>-sha1</code> parameters, AD Check Digest computes the MD5 and SHA-1 digests for the patch.</p>
-md5	<p>The MD5 output from the My Oracle Support patch download page.</p> <p>When you specify the <code>-md5</code> parameter, AD Check Digest compares the MD5 value you provide with the MD5 digest computed for the patch file.</p>
-sha1	<p>The SHA-1 output from the My Oracle Support patch download page.</p> <p>When you specify the <code>-sha1</code> parameter, AD Check Digest compares the SHA-1 value you provide with the SHA-1 digest computed for the patch file.</p>

Running AD Check Digest:

Run this utility as follows.

1. Set the environment.

You must set the environment in order to apply the environment variables that define your system. This task is common to many AD utilities. See *Setting the Environment*, page 7-52 in this book for the basic steps.

2. Run AD Check Digest.

Use the following commands. The output is displayed on the screen.

UNIX

```
$ adchkdig -file <File> [ -md5 <MD5_digest> -sha1 <SHA-1_digest> ]
```

Troubleshooting Applications DBA Operations

Managing Worker Processes

AD Administration and AutoPatch can perform processing jobs in parallel to speed the time it takes to complete them. This section describes the procedures for reviewing these processes and handling situations where processing has been interrupted.

Note: For more information, see Using Parallel Processing, page 7-55.

Reviewing Worker Status:

Requirement

How can I monitor the progress of parallel processing jobs?

Discussion

When AD Administration and AutoPatch process jobs in parallel, they assign jobs to workers for completion. There may be situations that cause a worker to stop processing. AD Controller is a utility you can use to determine the status of workers and manage worker tasks. You use it to monitor the actions of workers and the status of the processing jobs they have been assigned.

Actions

To review worker status, perform these steps:

1. Set the environment by executing (sourcing) the patch file system environment file:

```
$ source <patch APPL_TOP path>/APPS<CONTEXT_NAME>.env
```

Note: For more information, see Setting the Environment in

2. Start AD Controller by entering `adctrl` on the command line.

3. Review worker status.

Select "Show worker status" from the AD Controller main menu. AD Controller displays a summary of current worker activity. The summary columns are:

- Control Worker is the worker number
- Code is the last instruction from the manager to this worker
- Context is the general action the manager is executing
- Filename is the file the worker is running (if any)

The following table describes the types of status that may be assigned to a worker and reported in the Status column.

Worker Status Values

Status	Meaning
Assigned	The manager assigned a job to the worker, and the worker has not started.
Completed	The worker completed the job, and the manager has not yet assigned it a new job.
Failed	The worker encountered a problem.
Fixed, Restart	The worker should retry the failed operation now that the problem has been fixed.
Restarted	The worker is retrying a job or has successfully restarted a job (note that the status does <i>not</i> change to Running).
Running	The worker is running a job
Wait	The worker is idle.

If the worker status shows as Failed, the problem may need to be fixed before the

AD utility can complete its processing. This is described next.

Determining Why a Worker Failed:

Requirement

One of the workers has failed. How do I determine the cause of the failure?

Discussion

When a worker fails its job, you do not have to wait until the other workers and the manager stop. Use the worker log files (`adworknnn.log`) to determine what caused the failure. These log files are written to `APPL_TOP/admin/<SID>/log`. You can find the worker log file and copy it to a temporary area so that you can review it. If the job was deferred after the worker failed, there may be no action required on your part.

The first time a job fails, the manager defers the job and assigns a new worker. If the deferred job fails a second time, the manager defers it a second time only if the runtime of the job is less than ten minutes. If the deferred job fails a third time, or if the job's runtime is greater than ten minutes, the job stays at a failed status and the worker waits for intervention.

Actions

1. Set the environment by executing (sourcing) the patch file system environment file:

```
$ source <patch APPL_TOP path>/APPS<CONTEXT_NAME>.env
```

Note: For more information, see *Setting the Environment in Running AD Utilities*, page 7-52.

2. Start AD Controller by entering `adctrl` on the command line.
3. Identify the worker that encountered a problem.

Workers that have encountered problems stop processing jobs and show a status of Failed. Follow the steps in the *Reviewing Worker Status*, page 9-1 section in this chapter to determine which workers have a status of Failed.

4. Review the log file to find out why the worker failed.

The following is an example of a worker failure message:

```

AD Worker error:
The following ORACLE error:

ORA-01630: max # extents (50) reached in temp segment in tablespace
TSTEMP
occurred while executing the SQL statement:

CREATE INDEX AP.AP_INVOICES_N11 ON AP.AP_INVOICES_ALL (PROJECT_ID,
TASK_ID)
NOLOGGING STORAGE (INITIAL 4K NEXT 512K MINEXTENTS 1 MAXEXTENTS 50
PCTINCREASE 0 FREELISTS 4) PCTFREE 10 MAXTRANS 255 TABLESPACE APX

AD Worker error:
Unable to compare or correct tables or indexes or keys because of
the error
above

```

In this example, the worker could not create the index AP_INVOICES_N11 because the maximum number of extents in the temporary tablespace was reached.

5. Determine how best to resolve the problem that caused the failure. For example, search My Oracle Support for potential causes. If you cannot identify a fix, you may wish to open a service request with Oracle Support.

Handling a Failed Job:

Requirement

I have reviewed the log file for the failed worker and determined the problem. What do I do next?

Discussion

A worker usually runs continuously in the background and when it fails to complete the job it was assigned, it reports a status of Failed. When the manager displays an error message, confirm the failed status of a worker by using AD Controller to review worker status. If the job was deferred after the worker failed, no action may be required.

Note: For more information, see Using Parallel Processing, page 7-55.

Actions

Perform the following steps:

1. Set the environment by executing (sourcing) the patch file system environment file:

```
$ source <patch APPL_TOP path>/APPS<CONTEXT_NAME>.env
```
2. Start AD Controller by entering `adctrl` on the command line.

Note: For more information, see Setting the Environment in Running AD Utilities, page 7-52.

3. Identify the failed file.

The Worker and Filename columns in the AD Controller worker status screen show the numbers of the workers that failed and list the name of the files that failed to run.

4. Review the worker log file.

Each worker logs the status of tasks assigned to it in a log file called adworkxxx.log, where nnn is the worker number. For example, adwork001.log for worker 1 and adwork007.log for worker 7. These files are in the \$APPL_TOP/admin/<SID>/log directory on the patch file system. Review adworkxxx.log for the failed worker to determine the source of the error.

5. Resolve the error.

Resolve the error using the information provided in the log files. Contact Oracle Support Services if you do not understand how to resolve the issue.

6. Restart the failed job.

Choose Option 2 from the AD Controller main menu to tell the worker to restart a failed job.

7. Verify worker status.

Choose Option 1 again. The Status column for the worker that failed should now say Restarted or Fixed, Restart.

Note: When all workers are in either Failed or Wait status, the manager becomes idle. At this point, you must take action to get the failed workers running again.

Terminating a Hanging Worker Process:

Requirement

A worker process has been running for a long time. What should I do?

Discussion

When running AD utilities, there may be situations when a worker process appears to hang, or stop processing. If this occurs, it may be necessary to terminate the process manually. Once you do, you must also restart that process manually.

Caution: A process that appears to be hanging could actually just be a long-running job.

To terminate a process, start AD Controller, obtain the ID of the worker, and then stop

any hanging processes. Once you make the necessary changes, you can restart the job or worker.

Note: For more information, see *Restarting a Failed Worker*, page 9-7.

Actions

1. Set the environment by executing (sourcing) the patch file system environment file:

```
$ source <patch APPL_TOP path>/APPS<CONTEXT_NAME>.env
```

Note: For more information, see *Setting the Environment in Running AD Utilities*, page 7-52

2. Start AD Controller by entering `adctrl` on the command line.

3. Determine what the worker process is doing.

Use the AD Controller worker status screen to determine the file being processed and check the worker log file to see what it is doing:

- Check whether the process is consuming CPU.
- Review the file to see what actions are being taken.
- Check for correct indexes on the tables (if the problem appears to be performance-related).
- Check for an entry for this process in the `V$SESSION` table. This may provide clues to what the process is doing in the database.

4. Get the worker's process ID.

If the job is identified as "hanging," determine the worker's process ID.

UNIX:

```
$ ps -a | grep adworker
```

Windows:

Invoke the Windows Task Manager (with Ctrl-Alt-Delete or Ctrl-Shift-Esc) to view processes.

5. Determine what processes the worker has started, if any.

If there are child processes, get their process IDs. Examples of child processes include `SQL*Plus` and `FNDLOAD`.

6. Stop the hanging process, using the command that is appropriate for your operating system.

7. Fix the issue that caused the worker to hang. Contact Oracle Support Services if you require assistance doing this.
8. Restart the job or the worker.
See Restarting a Failed Worker, page 9-7 in this chapter for more information.

Restarting Processes

This section describes some situations where you may need to choose the restart option in AD Controller.

Restarting a Failed Worker:

Requirement

I need to restart a failed worker.

Discussion

If a worker has failed, or if you have terminated a hanging worker process, you need to restart the worker manually.

Some worker processes spawn other processes called child processes. If you terminate a child process (that is hanging), the worker that spawned the process shows Failed as the status. After you fix the problem, choose to restart the failed job. Once the worker is restarted, the associated child processes are started as well.

Actions

Perform these steps:

1. Set the environment by executing (sourcing) the patch file system environment file:

```
$ source <patch APPL_TOP path>/APPS<CONTEXT_NAME>.env
```

Note: For more information, see Setting the Environment in Running AD Utilities, page 7-52.

2. Start AD Controller by entering `adctrl` on the command line.
3. Choose Option 1 to review worker status.
4. Take the appropriate action for each worker status.

If the worker shows Failed, choose Option 2 to restart the failed job. When prompted, enter the number of the worker that failed.

If the worker shows Running or Restarted status, but the process is not really running, select the following menu options:

- Option 4: Tell manager that a worker has failed its job. When prompted, enter the number of the hanging worker.
- Option 6: Tell manager to start a worker that has shut down on the current machine. When prompted, enter the number of the worker that failed.

Caution: Do not choose Option 6 if the worker process is running. Doing so will create duplicate worker processes with the same worker ID.

The worker will restart its assigned jobs and spawn the necessary child processes.

Restarting an AD Utility After Machine Failure:

Requirement

While I was running an AD utility, the machine crashed. What is the best way to the restart the utility?

Discussion

Because the manager cannot automatically detect a machine crash, you must manually notify it that all jobs have failed and manually restart the workers. If you restart the utility without doing this, the utility status and the system status will not be synchronized.

Actions

Perform these steps:

1. Set the environment by executing (sourcing) the patch file system environment file:

```
$ source <patch APPL_TOP path>/APPS<CONTEXT_NAME>.env
```

Note: For more information, see *Setting the Environment in Running AD Utilities*, page 7-52.

2. Start AD Controller by entering `adctrl` on the command line.
3. Select the following options:
 - Option 4: Tell manager that a worker has failed its job (specify 'all' for workers)
 - Option 2: Tell worker to restart a failed job (specify 'all' for workers)
4. Restart the AD utility that was running when the machine crashed.

Shutting Down and Restarting Managers

This section discusses some reasons for shutting down and reactivating managers.

Shutting Down a Manager:

Requirement

How do I stop an AD utility while it is running?

Discussion

There may be situations when you need to shut down an AD utility while it is running. For example, you may need to shut down the database during an AutoPatch or AD Administration session.

You should perform this shutdown in an orderly fashion so that it does not affect your data. The best way to do this is to shut down the workers manually so that the AD utility quits in an orderly fashion.

Actions

Perform these steps:

1. Start AD Controller

Set the environment by executing (sourcing) the patch file system environment file:

```
$ source <patch APPL_TOP path>/APPS<CONTEXT_NAME>.env
```

Note: For more information, see *Setting the Environment in Running AD Utilities*, page 7-52.

2. Start AD Controller by entering `adctrl` on the command line.
3. In `adctrl`, select Option 3 and enter 'all' for the worker number. Each worker stops when it either completes or fails its current job.
4. Verify that no worker processes are running. Use the appropriate command for your platform.

UNIX:

```
$ ps -a | grep adworker
```

Windows:

Invoke Windows Task Manager (with Ctrl-Alt-Delete or Ctrl-Shift-Esc) to view the relevant processes.

5. When all workers have shut down, the manager and the AD utility quit.

Restarting a Manager:

Requirement

No workers are running jobs, when they should be doing so. What is the problem?

Discussion

A restarted worker resumes the failed job immediately as long as the worker process is running. The other workers change to a Waiting status if they cannot run any jobs because of dependencies on the failed job, or because there are no jobs left in the phase. When no workers are able to run, the manager becomes idle and messages like the following will appear on the screen:

```
ATTENTION: All workers either have failed or are waiting:
```

```
FAILED: file cedropcb.sql on worker 1.  
FAILED: file adgrnctx.sql on worker 2.  
FAILED: file aftwf01.sql on worker 3.
```

```
ATTENTION: Please fix the above failed worker(s) so the manager can  
continue.
```

Actions

Complete the following steps for each failed worker:

1. Start AD Controller.

Note: For more information, see Setting the Environment in Running AD Utilities, page 7-52.

2. Determine the cause of the error.

Choose Option 1 to view the status. Review the worker log file for the failed worker to determine the source of the error.

3. Resolve the error.

Use the information provided in the log files. Contact Oracle Support Services if you do not understand how to resolve the issue.

4. Restart the failed job.

Choose Option 2 on the AD Controller menu to tell the worker to restart a failed job. The worker process restarts, causing the AD utility to become active again.

Technology Inventory Utility

Technology Inventory Utility

This chapter describes the *Technology Inventory Utility* that was introduced in Oracle E-Business Suite Release 12. This command-line utility generates reports that list the installed technology stack components and versions on the various nodes of a Release 12.x Oracle E-Business Suite system. The reports can be generated in either HTML (the default) or text format. Separate reports are generated for the database and application tiers.

Since there are major differences in technology components between Release 11*i* and Release 12.x, this utility will also be useful for those who wish to become familiar with the components and versions employed by Release 12.x.

Running the Technology Inventory Utility

The Technology Inventory Utility generates a consolidated report that summarizes the version levels of all installed technology stack components.

Set your Oracle E-Business Suite environment, then run one of the following commands:

On UNIX:

Application tier:

```
perl $FND_TOP/patch/115/bin/TXKScript.pl
-script=$FND_TOP/patch/115/bin/txkInventory.pl
-txktop=$APPLTMP
-contextfile=$CONTEXT_FILE
-appspass=apps
-outfile=$APPLTMP/Report_Inventory.html
```

Database tier:

```
perl $ORACLE_HOME/appsutil/bin/TXKScript.pl
-script=$ORACLE_HOME/appsutil/bin/txkInventory.pl
-txktop=$ORACLE_HOME/appsutil/temp
-contextfile=$CONTEXT_FILE
-appspass=apps
-outfile=$ORACLE_HOME/appsutil/temp/Report_Inventory.html
```

Note: To generate the report in text format, append `-reporttype=text` to the relevant command, and change the outfile name to have a .txt suffix instead of a .html suffix.

On Windows:

Application tier:

```
perl %FND_TOP%\patch\115\bin\TXKScript.pl
-script=%FND_TOP%\patch\115\bin\txkInventory.pl
-txktop=%APPLTMP%
-contextfile=%CONTEXT_FILE%
-appspass=apps
-outfile=%APPLTMP%\Report_Inventory.html
```

Database tier:

```
perl %ORACLE_HOME%\appsutil\bin\TXKScript.pl
-script=%ORACLE_HOME%\appsutil\bin\txkInventory.pl
-txktop=%ORACLE_HOME%\appsutil\temp
-contextfile=%CONTEXT_FILE%
-appspass=apps
-outfile=%ORACLE_HOME%\appsutil\temp\Report_Inventory.html
```

Note: To generate the report in text format, append `-reporttype=text` to the relevant command, and change the outfile name to have a .txt suffix instead of a .html suffix.

Parameters

The following table describes the parameters for the utility:

Parameter	Usage
txktop	Temporary working directory used by perl modules. Required parameter.
contextfile	Location of the Applications context file. If not specified, default is picked from environment.
appspass	APPS schema password. If not specified, default password is used.

Parameter	Usage
outputfile	Location of the report being generated. If not specified, the default location is \$APPLTMP/TXK.

Output from the Technology Inventory Utility

The report generated on both the application and database tiers has the following common header:

Parameter	Usage
Date	Date on which report was generated.
Hostname	Details of host on which report was generated.
Enabled Services	Services enabled on the host where report was generated (application tier only).
Instance	Name of the instance.
Platform	OS name of the host where report was generated.
OS	OS release version of the host where report was generated.
DB Host	Details of the host where database is located.
Context File	Location of the context file specified when the report was generated.
Report File	Location of the report that was generated.
XML Definition File	Lists actions executed to obtain the contents of the report.

The contents of the main report reflect the role of the node on which the utility is run: Database, Web, Forms, or Concurrent Processing.

Future Directions

A future release of Oracle E-Business Suite will build on the Inventory Utility to validate the components in use.

Managing Oracle Fusion Middleware Log Files

Collecting and Managing Log Files

This chapter describes debugging techniques and options for the Oracle Fusion Middleware 11g components of E-Business Suite Release 12.2. These include:

- Oracle Process Manager and Notifications Server, which manages the HTTP Server.
- Oracle HTTP Server (Apache), which acts as a single Entry Point Proxy Server.
- Oracle WebLogic Server, which contains Oracle E-Business Suite Application Server Deployments.

This chapter provides steps to enable extra debug and automate collection of log files required for debugging the various issues that you may encounter.

Important: Any increase to logging levels should be reversed once the necessary debug information has been collected. Oracle does not recommend that detailed logging is left activated for extended periods, especially on production systems.

Collecting Log Files at Lower Debug Levels

This section should be followed if you are planning to create a service request for an issue that has already occurred or which is not easily reproducible. The procedures described do not enable additional debug information to be collected: rather, they simply collect log files at whatever debug levels are currently configured.

For some problems, the debug information collected here may not suffice: in such cases, you should follow the steps in the section Collecting Log Files at Higher Debug Levels.

Oracle Process Manager and Oracle HTTP Server Log Files

To gather the log files for OPMN and OHS services, run the following command as the owner of the application tier file system:

```
$ zip -r /tmp/\uname -n`_`date +%m%d%y.%H%M`_OPMN_OHS.zip \  
$IAS_ORACLE_HOME/instances/*/diagnostics/logs/OHS/EBS_web_component/*log  
* \  
$IAS_ORACLE_HOME/instances/*/diagnostics/logs/OPMN/opmn/*
```

This will create zip files in the /tmp directory with names including the server name, date, and time. For example:

server.customer.com_080712.1405_OPMN_OHS.zip. (Here, the date is specified in MMDDYY format.) To support a service request, you may collect these files from the /tmp directory and upload them to the request.

Oracle WebLogic Server Log Files

To gather the log files for Oracle Fusion Middleware components such as Node Manager, Admin Server, forms, oacore and oafm services, run the following command as the owner of the application tier file system:

```
$ zip -r /tmp/\uname -n`_`date +%m%d%y.%H%M`_FMW.zip \  
$IAS_ORACLE_HOME/./wlserver_10.3/common/nodemanager \  
$EBS_DOMAIN_HOME/servers/oa*/logs/* \  
$EBS_DOMAIN_HOME/servers/forms*/logs/* \  
$EBS_DOMAIN_HOME/servers/AdminServer/logs/* \  
$EBS_DOMAIN_HOME/sysman/log/*
```

This will create zip files in the /tmp directory with names including the server name, date, and time. For example:

server.customer.com_080712.1405_OPMN_OHS.zip. (Here, the date is specified in MMDDYY format.) To support a service request, you may collect these files from the /tmp directory and upload them to the request.

Collecting Log Files at Higher Debug Levels

This section describes how to enable the higher logging levels required to capture the additional debug information that is sometimes needed for resolving problems. You may be directed to this section by Oracle Support. Follow the steps below to enable debug for the appropriate component, in readiness for reproducing the issue and (typically) sending the collected log files to Oracle.

Note: Generally speaking, increasing the log levels in this way should only have a minimal impact on performance (maximum 5% degradation).

Oracle Process Manager

To enable debug for opmn, open the file

`$IAS_ORACLE_HOME/instances/<instance>/config/OPMN/opmn/opmn.xml`
with a text editor (such as `vi`) and change the line:

```
<debug comp="" rotation-size="1500000"/>
```

to:

```
<debug comp="ons[all];pm[all]" rotation-size="1500000"/>
```

Save your changes and exit the editor.

Alternatively, you can increase the opmn logging level for the current session by issuing the command:

```
$IAS_ORACLE_HOME/instances/<yourinstance>/bin/opmnctl set target=debug  
comp="ons[all];pm[all]"
```

This will enable debug for *all* components running under the opmn process. However, this may not be practical in a large number of situations.

Oracle HTTP Server Access Logging

To adjust OHS logging levels, access the Enterprise Manager console at `http://server.domain:wls_admin_port/em` and log in as `weblogic` with password `welcome1`.

Expand the Web Tier menu, right-click 'EBS_web_component', and choose 'Administration - Log Configuration'. Here you have the option to set the log level. Set it to 'Trace:32'.

Node Manager

To adjust the log level for Node Manager, open the file `$IAS_ORACLE_HOME/..wls_server_10.3/common/nodemanager/nodemanager.properties` with a text editor (such as `vi`) and add the line:

```
adjust LogLevel=INFO to LogLevel=FINEST
```

Start and Stop Logging

To adjust start/stop log levels, access the Enterprise Manager console at `http://server.domain:wls_admin_port/em` and log in as `weblogic` with password `welcome1`.

Expand 'WebLogic Domain', then expand 'EBS_domain_X'. Right-click 'AdminServer' and select 'Logs - Log Configuration'. Here you have the option to set the logging level. Set it to 'Trace:32'.

Oracle WebLogic Server Logging

Part 1: Enable Logging Level

1. To adjust Oracle WebLogic Server logging level, access the Enterprise Manager console at `http://server.domain:wls_admin_port/em` and log in as `weblogic` with password `welcome1`.

2. Expand 'Environment - Servers'.
3. Select the servers you wish to enable debug for. Typically, this will be the oacore servers.

Note: If you have more than one server defined in a cluster, you will need to enable debug for each server individually. Only enable logging for managed servers that are experiencing issues.

4. Choose the 'Logging' tab, select the 'Advanced' link, and change severity level to 'Trace'.
5. Change Severity levels for standard out and domain log broadcaster to 'Debug'.
6. On this same configuration page, you will see the setting 'Limit number of retained files'. You may need to increase this value, as extra logging might otherwise result in log files being overwritten because of increased log file rotation.

Part 2: Enable Debug Areas

Continuing from Part 1 above:

1. Select the 'Debug' tab and expand the 'WebLogic' section.
2. Select the Servlet group for which you wish to enable debug:

Component	When to Enable Debug
Servlet	You are experiencing general issues.
JDBC	You are experiencing JDBC connection problems.
ClassLoader	You are experiencing issues relating to classloading, such as missing classes on startup of services.

3. Choose 'Enable' to activate debug for these components.

Restart Application Tier Services and Clear Log Files

You are now ready to restart application tier services using the scripts located in \$ADMIN_SCRIPTS_HOME.

At this stage, it is desirable to clear the existing log files. Using a new set of log files will enable the problem to be more easily pinpointed. This is especially true of production

systems where logs can be large and consequently hard to debug.

Log files should be cleared while the application tier services are down. You can do this manually, or use the following script.

Script For Automatically Clearing Contents of Log Files

Source the environment, ensure the application tier services are shut down (as noted above), then run the following script as the applmgr user:

```
for files in $IAS_ORACLE_HOME/instances/*/diagnostics/logs/OPMN/opmn
$EBS_DOMAIN_HOME/sysman/log
do
do
for file in `ls $files/*`
do
echo "Clearing file: " $file
cat /dev/null > $file
done
done
for files in
$IAS_ORACLE_HOME/instances/*/diagnostics/logs/OHS/EBS_web_component
$EBS_DOMAIN_HOME/servers/*/logs
$IAS_ORACLE_HOME/./wlserver_10.3/common/nodemanager
do
do
for file in `ls $files/*log* $files/*out*| grep -v lck`
do
echo "Clearing file: " $file
cat /dev/null > $file
done
done
```

Tip: To identify the log files without making any changes, comment out the line `cat /dev/null > $file` by placing a # character at the beginning. When you are happy with the proposed changes, remove the # and re-run the script.

For additional information on logging definition, rotation, location, and control, refer to My Oracle Support Knowledge Document 1366187.1, *Oracle Applications E-Business Suite 12.2 Fusion Middleware Log Files: Locate, View, and Control*.

Reproduce Issue and Upload Log Files

You are now ready to reproduce the issue and gather the log files, typically for uploading to your service request.

Logging Features in Oracle E-Business Suite

Overview

Oracle E-Business Suite provides several types of logs in its products. This chapter focuses on the Logging Framework, which enables you to set up and view log messages in Oracle Applications Manager.

Other types of log files are discussed in the documentation for their respective product area.

The Logging Framework

The Oracle E-Business Suite Logging Framework provides the ability to store and retrieve log messages for debugging, error reporting, and alerting purposes.

You can set up, view, and purge log messages through HTML-based user interface pages that are located in Oracle Applications Manager.

These messages must be written into the code by developers. For more information, refer to the *Oracle E-Business Suite Developer's Guide*.

Configuring the Logging Framework

The following sections cover configuration of the logging framework.

Using Middle-tier Properties to Configure Logging

All middle-tier property settings take precedence over profile option settings. Configuring logging using Java system properties (usually by setting the Apache JServ system properties in the `jserv.properties` file) is a quick way to turn on logging for all sites or users, regardless of current profile option settings. Middle-tier properties only affect the middle-tier code, and do not affect the PL/SQL layer logging.

Using Java

Java system properties can be defined for controlling logging for each JVM.

The following examples show how to turn on logging for all modules and levels using Java system properties.

For this example, we assume that the JVM has write permission for the file "/path/to/apps.log". This file can be changed to any other file for which the JVM has write permission.

If you plan to log to a file, it is highly recommended that you explicitly override the default file "aferror.log" by setting AFLOG_FILENAME. The default does not specify a full file path, and may not be writable by the middle-tier process in some cases.

Command Line JVM System Properties

To enable logging for an application (for example, MyClass) that is run from the command line, add the parameter values to the command line:

```
/local/java/jdk1.2.2/bin/java \  
-DAFLOG_ENABLED=TRUE -DAFLOG_LEVEL=STATEMENT \  
-DAFLOG_MODULE=% -DAFLOG_FILENAME=/path/to/apps.log MyClass
```

Apache JServ Java System Properties

To enable logging using Apache JServ JVM system properties, add the following to the jserv.properties file (typically located in \$IAS_ORACLE_HOME/Apache/Jserv/etc/):

```
wrapper.bin.parameters=-DAFLOG_ENABLED=TRUE  
wrapper.bin.parameters=-DAFLOG_LEVEL=STATEMENT  
wrapper.bin.parameters=-DAFLOG_MODULE=%  
wrapper.bin.parameters=-DAFLOG_FILENAME=/path/to/apps.log
```

A convenient location for the log file in this case is the log directory used by Jserv (\$IAS_ORACLE_HOME/Apache/Jserv/logs/).

Using C

Environment variables can be defined for controlling logging for each C process.

The following examples show how to turn on logging for all modules and levels using C Environment variables.

For this example, we assume that the C process has write permission for the file "/path/to/apps.log". This file can be changed to any other file for which the C process has write permission.

If you plan to log to a file, it is highly recommended that you explicitly override the default file "aferror.log" by setting AFLOG_FILENAME. The default does not specify a full file path, and may not be writable by the middle-tier process in some cases.

```
#!/bin/csh  
setenv AFLOG_ENABLED Y  
setenv AFLOG_LEVEL STATEMENT  
setenv AFLOG_MODULE %  
setenv AFLOG_FILENAME /path/to/apps.log  
./C-Executable
```


Using Oracle Application Object Library Profile Options to Configure Logging

You can configure logging by setting Oracle Application Object Library (FND) profile options. The following table lists profile option names and sample values:

Profile Options

Profile Option Name	User Specified Name	Sample Value
AFLOG_ENABLED	FND: Debug Log Enabled	"Y"
AFLOG_MODULE	FND: Debug Log Module	"%"
AFLOG_LEVEL	FND: Debug Log Level	"ERROR"
AFLOG_FILENAME	FND: Debug Log Filename	"/path/to/apps.log"

The available levels are Site, Application, Responsibility, and User. User settings override Responsibility settings, Responsibility settings override Application settings, and Application settings override Site settings.

To emphasize this point, the following is a summary of the impacts of the different profile option levels:

- User: Affects only the given user.
- Application: Affects all users for the specific application.
- Responsibility: Affects all users in any application for that responsibility.
- Site: Affects all users, applications, and responsibilities.

Note: When setting up logging at the Site level, we strongly recommend that you set the logging level to UNEXPECTED. ERROR or EXCEPTION are also possibilities. We strongly discourage setting the logging level for a site to anything other than UNEXPECTED, ERROR, or EXCEPTION.

Using Logging to Screen

In addition to the above methods where log messages are written to a file or the database, Logging to Screen provides:

- The ability to enable logging on a per HTTP request or per HTTP session basis.

- Dynamic configuration which does not require restarting any servers or changing any log profiles.
- A convenient lightweight mechanism to diagnose performance issues. Each message is timestamped to the millisecond.

If Logging to Screen is enabled, then the Java log messages generated for a particular HTTP Request-Response are buffered in memory and appended to the end of the generated HTML page.

This feature does not affect any existing configurations of file or database logging. File or database logging continues to behave per the configured middle tier log properties and/or log profile values.

Note that this mechanism currently provides only Java layer messages. Regular file or database logging should be used if messages from other layers (e.g., PL/SQL) are needed.

Enabling Logging to Screen in Oracle Application Framework Pages

For security reasons, this feature is only accessible if the "FND: Diagnostics" Profile is set to "Yes".

Use the following procedure to enable Logging to Screen in pages based on the Oracle Application Framework:

1. Click the **Diagnostics** button.
2. Select **Show Log to Screen** from the drop-down list.
3. Choose an appropriate log level.
4. Optionally, enter a module filter criteria such as **jtf***. [In URLs, use the asterisk symbol (*) as a wildcard character, not the percent sign (%).]

Enabling Logging to Screen in CRM Technology Foundation Pages

For security reasons, this feature is only accessible if the "FND: Diagnostics" Profile is set to "Yes".

To enable logging to screen in pages based on the CRM Technology Foundation, append the following to the page's URL:

jtfdebug

Specify the logging level that should be displayed on the current screen.

jtfdebugfilter

(Optional) If desired, this parameter can be used as a filter to display messages based on a Java package name.

For example: `<current_url>&jtfdebug=STATEMENT&jtfdebugfilter=jtf*`

[In URLs, use the asterisk symbol (*) as a wildcard character, not the percent sign (%).]

Startup Behavior

At startup, applications do not have access to profile values. If middle-tier properties are not set, then at startup, the system defaults to logging as follows:

- Logs are stored in the file `aferror.log` (in the current directory).
- Logs are stored at the level `UNEXPECTED`.
- Logs are stored for all modules.

After a connection to the database has been established, the site-level log profiles are read. When the user, responsibility, and application have been established, the Oracle Application Object Library (FND) profiles are read for that user.

For Java and PL/SQL applications, the logging system is initialized by `FND_GLOBAL.INITIALIZE` (which is called from `APPS_INITIALIZE`), which is called normally as part of the startup of every Java application session, form, report, or concurrent program. At that point, it has user information and will log with the proper user profiles. Before the `FND_GLOBAL.INITIALIZE`, if the logging system is called it will self-initialize and log with the site-level profile values.

For Java applications, this is the sequence of startup steps:

1. If any of the log parameters are set as Java system properties, then use them.
2. Logging is not disabled using the Java system property `AFLOG_ENABLED=FALSE`, and if any of the remaining log parameters are not set as system properties, then retrieve the corresponding Oracle Application Object Library (FND) profile option values from the database. User-level profile values override responsibility-level profile values, which override application-level profile values, which override site-level profile values.
3. If any of the log parameters are not set either as system properties or as profile values (or they are not accessible due to an error), then use the default values.

Guidelines for the Logging Framework

Set up your system for logging according to the following guidelines. We recommend that you use Oracle Applications Manager as the user interface for any log management tasks.

Recommended Default Site-Level Settings

For normal operations, we recommend that you configure your system as follows:

- Enabled: On

- Logging Level: UNEXPECTED
- Log Repository: Database
- Module Filter: %

Caution: If you set the default site-level logging level to STATEMENT or PROCEDURE, a decrease in system performance could result. Under that configuration, the large amount of generated log messages might significantly slow down the system. Furthermore, if the site-level logging level is set to a low severity for a long time, then the FND_LOG_MESSAGES table could potentially run out of space.

Recommended Settings for Debugging

If you need to lower the logging level in order to gather information about a system error, use the following recommended configurations. (As stated above, the default logging level should be UNEXPECTED. This maintains optimum system performance.)

Using Logging to Screen

For Java-based pages that are based on the Oracle CRM Technology Foundation or the Oracle Application Framework, if you have access to the browser that is displaying the generated HTML, you can use the Logging to Screen feature to view further details if an error is reported. See: Using Logging to Screen, page 12-1.

This lightweight mechanism works best in cases where:

- You are interested in Java layer messages only.
- Debugging of is required for a particular request-response. For example, a JSP request from a browser.
- Debugging is required for all request-responses within a specific session.

Pinpointing an Error to a Specific User

You can use Oracle Application Object Library profiles to enable logging for the specific user, responsibility, and application that were active when the error occurred. Ask the user to log in again for the profile changes to take effect. Remember to return the profiles to their usual values after debugging has been completed.

If you suspect that certain code is causing the problem, then use hierarchical module filters to restrict which messages are logged. For example: fnd.common.%

Set the logging level according to the appropriate level of detail. Recall that EVENT messages report key progress events, while EXCEPTION, ERROR, and UNEXPECTED

messages report failures.

For High Volumes

For high load, high volume scenarios, you can log middle-tier messages to a local file, which is faster than logging to a remote database. To do so, define the `AFLOG_FILENAME` property to write all middle tier logging to a local file. Be sure to limit the number of generated messages:

- Use Oracle Application Object Library FND Profiles to restrict logging according to:
 - Specific users
 - Specific responsibilities
 - Specific applications
- If you suspect that certain code is causing the problem, then use hierarchical module filters to restrict which messages are logged. For example: `fnd.common%`
- Set the logging level according to the appropriate level of detail. Recall that `EVENT` messages report key progress events, while `EXCEPTION`, `ERROR`, and `UNEXPECTED` messages report failures.

For maintenance purposes, you should periodically rotate log files and purge old messages from the database table.

Updating Configuration Properties

If you have configured logging using Middle-tier properties, you need to restart the affected processes for any changes to be picked up.

If you have configured logging using FND Log Profiles, you need to request the user to log in again (no restart is needed)

If you have configured logging using Logging to Screen, the update is immediate. No re-login or restart is needed.

How to Completely Disable Logging

Use the following procedure to completely disable logging:

- If logging is configured using middle-tier properties, then set the `AFLOG_ENABLED` middle-tier properties to `FALSE` in all appropriate middle-tier configuration files (for example, `jserv.properties`) and/or startup scripts.
- If logging is configured using Oracle Application Object Library profiles in the database, use the logging setup screen in Oracle Applications Manager to turn off logging for all applications, responsibilities, and users. For details, see Oracle

Applications Manager online help.

See the "Updating Configuration Properties" section above for details on how and when the modified values come into effect.

Purging Log Messages

You should periodically delete old log messages to account for the space limitations of the database table. In addition, you should periodically rotate log files.

There are several ways to purge log messages. They are described below:

Using a Concurrent Program

The concurrent program "Purge Debug Log and System Alerts" (Short name: FNDLGPRG) is the recommended way to purge messages. This program purges all messages up to the specified date, except messages for active transactions (new or open alerts, active ICX sessions, concurrent requests, and so on). This program is by default scheduled to run daily and purge messages older than 7 days. Internally this concurrent program invokes the FND_LOG_ADMIN APIs, which are described later in this document.

Using Oracle Applications Manager

Navigate to **System Alerts and Metrics** from the **Navigate to** drop-down list on the Applications Dashboard. Then click **Logs**.

Using the Oracle CRM System Administrator Console

Navigate to **Settings > System > Debug Logging**.

Using PL/SQL

You can use the FND_LOG_ADMIN PL/SQL package to delete log messages.

For example:

```
SET SERVEROUTPUT ON
declare
    del_rows NUMBER;
BEGIN
    del_rows := fnd_log_admin.delete_all;
    DBMS_OUTPUT.PUT_LINE(del_rows || ' rows deleted');
END;
```

Viewing Log Messages

This section summarizes the different user interfaces that can be used to view and work with log messages, and how to access log messages from each UI.

CRM System Administrator Console

Navigate to **Settings > System > Debug Logging**.

Oracle Application Framework Pages

When working in Oracle Application Framework pages, you can use the following procedure to view log messages.

1. Pages based on the Oracle Application Framework have a global button labeled **Diagnostics**. Click this button to open a window where you can choose **Show Log**. (Note that this "Diagnostics" global button does not refer to the Diagnostics feature in Oracle Applications Manager that enables management and execution of diagnostic tests.)
2. Select **Show Log** to open the Logs page within Oracle Applications Manager. The Logs page is part of the System Alerts and Metrics feature.

Note: For the Diagnostics global button to be visible, the profile option FND_DIAGNOSTICS must be set to YES.

Oracle Applications Manager

The Logging features in Oracle Applications Manager can be accessed in the following ways:

- From the global Diagnostics button in Oracle Application Framework, select **Show Log**
- Using the System Administration responsibility, navigate to Oracle Applications Manager > Logs
- From the Oracle Applications Manager Site Map, navigate to Monitoring > Logs (Under Current Activity)
- In a Forms-based application, navigate to Help (menu) > Diagnostics > Logging

OAM gives you the capability to perform the following:

- Search using various criteria (Concurrent Program Request ID, Session ID, User ID, and so on).
- Drill down to view related logs
- View log attachments
- Add logs and attachments to the Support Cart

- Select valid values using Interdependent Lists of Values
- Export/Download Logs in CSV format
- Delete logs
- Configure logs based on logging profiles
- View a log summary
- Drill down from the OAM Concurrent Processing Request page to the request log

Access to the OAM Logging functionality is controlled using Oracle Application Object Library function security. The seeded responsibilities System Administration and System Administrator have the logging functions assigned. For other responsibilities, use the following:

- Full access (Search/Configure/Delete): OAM_BF_SYSLOG_ALL_MENU (Log Search: All Functions) (Menu)

This menu contains:

- Search: OAM_BF_SYSLOG_READ_ONLY_MENU (Log Search: Read Only) (Menu)
- Delete: OAM_BF_SYSLOG_DELETE (Log Delete) (function)
- Configure: OAM_BF_SYSLOG_CONFIG (Log Setup) (function)

Oracle Forms

Navigate to **Help > Diagnostics > Logging**.

Log Files in Applied Technology Products

The following table describes some of the log files used by Applied Technology products. These log files do not necessarily use the Logging Framework described above.

Product	User-friendly log filename or description	Default log file name(s), if any (generally for those log files not available in the UI)	Log file location and/or navigation path	For more information, see:
Oracle Alert	Action Log. This log shows the text of the action, if the action is a message action, operating system script, or SQL statement script action that is defined in Oracle Alert.		(Review Alert History form > Find Actions button > Action Log button)	Reviewing Action History, <i>Oracle Alert User's Guide</i>
Oracle Alert	Reply Log		Choose the Reply Log button to open a window and display response log. (Review Alert History form > Find Actions button > Response > Reply Log)	Reviewing Action History, <i>Oracle Alert User's Guide</i>

Product	User-friendly log filename or description	Default log file name(s), if any (generally for those log files not available in the UI)	Log file location and/or navigation path	For more information, see:
Oracle Alert	Request Log. This allows you to view the complete log file of the check request.		(Review Alert History form > Find Check button > Choose Request Log from the Tools menu)	Reviewing Alert Check Information, <i>Oracle Alert User's Guide</i>
Oracle Application Framework	error.log / jserv.log (An Apache/Servlet log file that identifies if any problems occur when starting the JVM)	error.log or jserv.log	Accessible through the Quick Apache configuration page.	<i>Oracle Application Framework Developer's Guide</i> , available from My Oracle Support Document 1315485.1.
Oracle Application Framework	fnd.log, for FND logging at the middle tier	fnd.log	/tmp/fnd.log	<i>Oracle Application Framework Developer's Guide</i> available from My Oracle Support Document 1315485.1.
Oracle Application Object Library - Concurrent Processing	Concurrent Request log file	l<request ID>.req	Default is \$APPLCSF/\$APPLLOG. BUT new variable: \$APPLDLM - if set to product, log files will go under \$APPLCSF/<PROD>/\$APPLLOG	Log and Output Filenames, <i>Oracle E-Business Suite Setup Guide</i>

Product	User-friendly log filename or description	Default log file name(s), if any (generally for those log files not available in the UI)	Log file location and/or navigation path	For more information, see:
Oracle Application Object Library - Concurrent Processing	Host Language Concurrent Program	FCP_LOG	same as above (nothing special)	Host Language Concurrent Programs, <i>Oracle E-Business Suite Setup Guide</i>
Oracle Application Object Library - Concurrent Processing	Internal Concurrent Manager Log File	<mgrname.mgr>	Default is \$APPLCSF/\$APPLLOG. BUT new variable: \$APPLLDMM - if set to product, log files will go under \$APPLCSF/system/\$APPLLOG. UI: Concurrent Processes window > Internal Manager Log, OAM > Site Map > Administration > Request Processing Managers > [Service] > View Processes > ICM Log (B)	Controlling the Internal Concurrent Manager from the Operating System, <i>Oracle E-Business Suite Setup Guide</i>

Product	User-friendly log filename or description	Default log file name(s), if any (generally for those log files not available in the UI)	Log file location and/or navigation path	For more information, see:
Oracle Application Object Library - Concurrent Processing	Concurrent Manager log file	w<number>.mgr	same as ICM , or Concurrent Requests window > Manager Log	Log and Output Filenames, <i>Oracle E-Business Suite Setup Guide</i>
Oracle Application Object Library - Concurrent Processing	Internal Monitor process log	i<number>.mgr	same as ICM log	Concurrent Processes Window, <i>Oracle E-Business Suite Setup Guide</i>
Oracle Application Object Library - Concurrent Processing	ICM Activation/Deactivation Logs	CM_<SID>.log and CS_<SID>.log	\$FND_TOP/\$APPLLOG	Setting Up, Starting, and Shutting Down Concurrent Managers, <i>Oracle E-Business Suite Setup Guide</i>
Oracle Application Object Library - Concurrent Processing	FRD log file	NA	GUI: OAM UI - "Forms Sessions for Process ID" > View Diagnostics	OAM Generic Collection Service, <i>Oracle E-Business Suite Setup Guide</i>
Oracle Application Object Library - Printing (Pasta)	Pasta error log file	user-defined	user-defined	Configuration File Options, <i>Oracle E-Business Suite Setup Guide</i> and Command Line Parameters, <i>Oracle E-Business Suite Setup Guide</i>

Product	User-friendly log filename or description	Default log file name(s), if any (generally for those log files not available in the UI)	Log file location and/or navigation path	For more information, see:
Oracle Application Object Library - Tablespace Migration Utility (OATM)	Log file for the generation of migration commands	fndgmcmd <timestamp>.log	Created in the working directory from which the user runs the PERL program.	Phase 1: Preparatory Steps, <i>Oracle E-Business Suite Setup Guide</i>
Oracle Application Object Library - Tablespace Migration Utility (OATM)	Log file for the migration of tables with LONG columns.	fndmlong<timestamp>.log	Created in the working directory from which the user runs the PERL program.	Phase 2: Migration Steps, <i>Oracle E-Business Suite Setup Guide</i>
Oracle Application Object Library - Tablespace Migration Utility (OATM)	Log file for the process to execute the script fndemseq.sql .	fndemseq<timestamp>.log	Created in the working directory from which the user runs the PERL program.	Phase 2: Migration Steps, <i>Oracle E-Business Suite Setup Guide</i>
Oracle Application Object Library - Tablespace Migration Utility (OATM)	Log file for the execution of execute the SQL script fndemcmd.sql	fndemcmd<timestamp>.log	Created in the working directory from which the user runs the PERL program.	Phase 2: Migration Steps, <i>Oracle E-Business Suite Setup Guide</i>

Product	User-friendly log filename or description	Default log file name(s), if any (generally for those log files not available in the UI)	Log file location and/or navigation path	For more information, see:
Oracle Application Object Library - Tablespace Migration Utility (OATM)	Log file for the process for enabling all the constraints, triggers, policies and start queues	fndenabl<timestam p>.log	Created in the working directory from which the user runs the PERL program.	Phase 3: Post-Migration Steps, <i>Oracle E-Business Suite Setup Guide</i>
Oracle Applications Manager	Forms Runaway Processes	NA	Site Map > Monitoring (subtab) > Current Activity (heading) > Forms Runaway Processes (link) Overview	Additional Monitoring Features and Options, page 14-12
Oracle Applications Manager	OAM log file		Navigation: Site Map > Administration > Applications Manager Log (under Others)	Oracle Applications Manager Log, page 14-47

Product	User-friendly log filename or description	Default log file name(s), if any (generally for those log files not available in the UI)	Log file location and/or navigation path	For more information, see:
Oracle E-Business Suite Integrated SOA Gateway	Design-time Logs, recorded during service generation and deployment life cycle for an interface that has the design-time log enabled		Accessible through the Integration Repository user interface.	Logging for Web Services, <i>Oracle E-Business Suite Integrated SOA Gateway Implementation Guide</i>
Oracle E-Business Suite Integrated SOA Gateway	Runtime Logs, recorded during the invocation of Oracle E-Business Suite services by Web service clients		These logs are captured and viewed through the Service Monitor user interface.	Monitoring and Managing SOAP Messages Using Service Monitor, <i>Oracle E-Business Suite Integrated SOA Gateway Implementation Guide</i>
Oracle E-Business Suite CRM Technology Foundation (JTT)	Page Flow Logging reports	NA	Settings tab > System > Properties > Page Flow Logging > Reports	Viewing Page Flow Logging Reports, <i>Oracle E-Business Suite CRM System Administrator's Guide</i>
Oracle iSetup	Extract Log File	NA	Migrations > Extract > View Log	Viewing an Extract, <i>Oracle iSetup User's Guide</i>
Oracle iSetup	Transform Log File	NA	Transforms > Extract > View Log	Viewing a Transform, <i>Oracle iSetup User's Guide</i>

Product	User-friendly log filename or description	Default log file name(s), if any (generally for those log files not available in the UI)	Log file location and/or navigation path	For more information, see:
Oracle Report Manager	Logs for Oracle Report Manager concurrent requests		Access through UI. Use standard concurrent processing UI or Oracle Report Manager: Home > Completed Requests > Details icon > (B) View Log	Report Manager Home Page, <i>Oracle Report Manager User's Guide</i>
Oracle Web Applications Desktop Integrator	Oracle Web Applications Desktop Integrator log	bne.log or value specified in BNE Server Log Filename profile option	FND_TOP/log or value specified in BNE Server Log Path profile option	Setting Profile Options, <i>Oracle Web Applications Desktop Integrator Implementation and Administration Guide</i>

Product	User-friendly log filename or description	Default log file name(s), if any (generally for those log files not available in the UI)	Log file location and/or navigation path	For more information, see:
Oracle Workflow	Oracle Workflow Builder log	wfbldr.log, if View menu > Log > To File is selected	If View menu > Log > To File is selected, then the log file is stored in the Oracle home on the client PC where Oracle Workflow Builder is installed, ORACLE_HOME\WF\DATA\us\wfbldr.log; or access through the Oracle Workflow Builder UI: View menu > Log > Show	Oracle Workflow Builder Menus, <i>Oracle Workflow Developer's Guide</i>
Oracle Workflow	Log of command-line diagnostic tests for notification mailers when you run the tests using oracle.apps.fnd.wf.mailer.Mailer	test.log or name specified in - Dlogfile parameter	Directory from which the diagnostic tests are run	Running Command-Line Notification Mailer Diagnostics, <i>Oracle Workflow Administrator's Guide</i>

Product	User-friendly log filename or description	Default log file name(s), if any (generally for those log files not available in the UI)	Log file location and/or navigation path	For more information, see:
Oracle Workflow	SQL trace file for an Oracle Workflow Business Event System agent listener, if you enable SQL tracing for the agent listener using the SQL_TRACE_LEVEL parameter	<INSTANCE>_ora_<PID>_WFAL_<componentId>_<time stamp>.trc. For example: WF11G_ora_254_WFAL_10002_20100302.trc	The location specified in the USER_DUMP_DEST parameter as listed in the V\$PARAMETER view	Scheduling Listeners for Local Inbound Agents, <i>Oracle Workflow Administrator's Guide</i>

Product	User-friendly log filename or description	Default log file name(s), if any (generally for those log files not available in the UI)	Log file location and/or navigation path	For more information, see:
Oracle Workflow	Logs for Oracle Workflow service component containers.	The log file names are determined by Oracle E-Business Suite Logging Framework properties or profile options, if written to a file.	Either written to the Generic Service Management log file or to the log file location specified in your Oracle E-Business Suite Logging Framework properties or profile options. View through Oracle Applications Manager UI, from the Service Components page or Component Details page. Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Service Components status icon > (B) View Log or Navigation:	Service Components, page 15-6

Product	User-friendly log filename or description	Default log file name(s), if any (generally for those log files not available in the UI)	Log file location and/or navigation path	For more information, see:
Oracle Workflow	Logs for Workflow Background Process concurrent requests		<p>Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Service Components status icon > (B) View Details icon > (B) View Log</p>	Engines, page 15-86
			<p>Location: Access through UI. Use standard concurrent processing UI or Oracle Applications Manager, from the Background Engines page. Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Background Engines status icon > (B) View Log</p>	

Product	User-friendly log filename or description	Default log file name(s), if any (generally for those log files not available in the UI)	Log file location and/or navigation path	For more information, see:
Oracle Workflow	Logs for Purge Obsolete Workflow Runtime Data concurrent requests		Access through UI. Use standard concurrent processing UI or Oracle Applications Manager, from the Workflow Purge pages. Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Purge status icon > Request Log link or Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Purge status icon > (B) View Purge Requests > (B) View Log	Purging Workflow Data, page 15-90

Product	User-friendly log filename or description	Default log file name(s), if any (generally for those log files not available in the UI)	Log file location and/or navigation path	For more information, see:
Oracle Workflow	Logs for Workflow Control Queue Cleanup concurrent requests		Access through UI. Use standard concurrent processing UI or Oracle Applications Manager, from the Control Queue Cleanup page. Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Control Queue Cleanup status icon > (B) View Log	Workflow Control Queue Cleanup, page 15-96

Running Diagnostics

Diagnostics in Oracle E-Business Suite

Oracle E-Business Suite delivers various diagnostic tools, including, for example:

- Diagnostics in Oracle Applications Manager, page 14-37
- Oracle Diagnostics Framework, page 13-1

In addition, diagnostic features are available from the Help menu in Forms-based applications in Oracle E-Business Suite.

Oracle Diagnostics Framework

Oracle Diagnostics Framework provides the infrastructure to execute diagnostic tests either for troubleshooting or for simply sanity-checking the Oracle E-Business Suite instance periodically or after applying any patch. Oracle E-Business Suite Diagnostics provides application specific troubleshooting tools that can help shorten problem-resolution time.

For more information, see Introduction to Oracle E-Business Suite Diagnostics, *Oracle Diagnostics Framework User's Guide*.

Controlling Access to the Oracle Forms-based Applications Diagnostics Menu

The Diagnostics menu is one of several menus available in Oracle Forms-based applications in Oracle E-Business Suite. These menus allow the user to invoke standard Oracle Forms functions, such as "Clear Record" as well as Oracle E-Business Suite-specific functions. For more information on the Oracle E-Business Suite pulldown menus for Forms-based applications, see: Pulldown Menus and the Toolbar, *Oracle E-Business Suite Developer's Guide*.

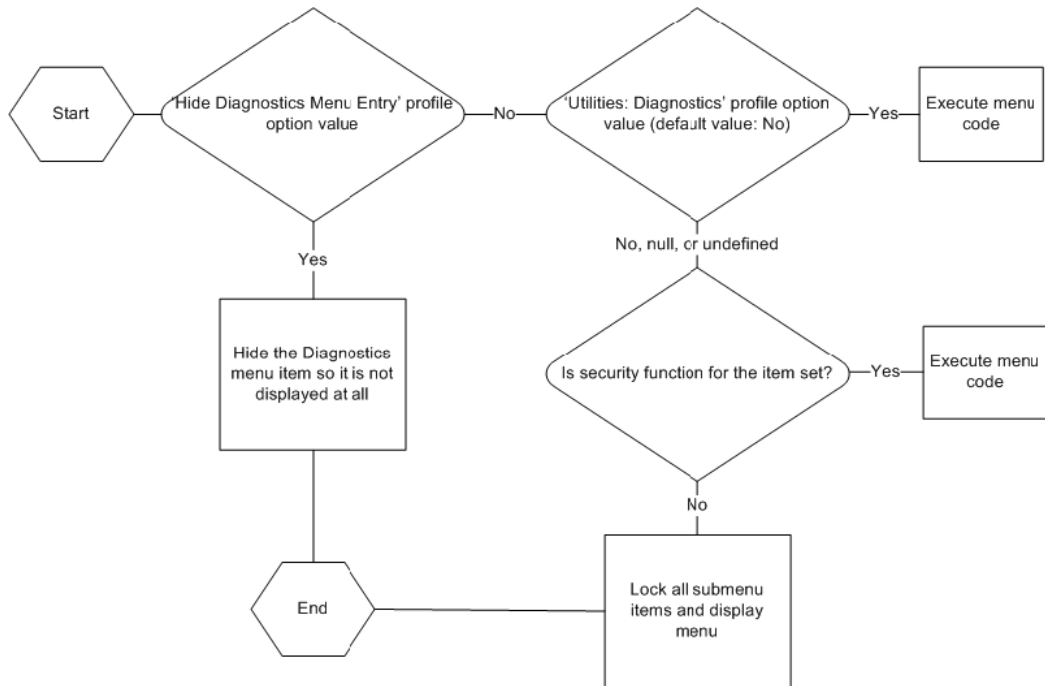
The Diagnostics menu can be found under the Help menu > Diagnostics.

In previous releases, access to the Diagnostics menu and submenu items was controlled by two profile options, Hide Diagnostics Menu Entry and Utilities: Diagnostics:

- Hide Diagnostics Menu Entry - If this profile is set to Yes, the Diagnostics menu is hidden from the user.
- Utilities: Diagnostics - Utilities: Diagnostics determines whether a user can automatically use the following Diagnostics submenu items: Examine, Trace, Debug, Properties, and Custom Code. If Utilities:Diagnostics is set to Yes, then users can automatically use these features. If Utilities:Diagnostics is set to No, then users must enter the password for the APPS schema to use these Diagnostics features. When the user no longer needs to use these features, he or she can select the menu entry Disable Secured Diagnostics.

Beginning with Release 12.1.3, access to the above Diagnostics submenu items can be controlled by the profile Utilities:Diagnostics or by security functions using Role-Based Access Control (RBAC). Whether or not a submenu item is available is checked on an as-needed basis by the system when the user selects the submenu item. If the menu item is not available to the user, the message "Function not available to this responsibility. Change responsibilities or contact your System Administrator."

The following flowchart illustrates how the menu items are secured:



Seeded Security Functions and Permission Sets

The following table lists the seeded securing functions and their corresponding Diagnostic menu items.

Note that there are read-only functions for some menu items.

Securing Function Name	Securing Function User-Friendly Name	Internal Menu Name	Runtime Menu Name
FND_DIAGNOSTICS_EXAMINE	FND Diagnostics Menu Examine	DIAGNOSTICS • EXAMINE	Diagnostics • Examine
FND_DIAGNOSTICS_EXAMINE_RO	FND Diagnostics Menu Examine	DIAGNOSTICS • EXAMINE	Diagnostics • Examine

Securing Function Name	Securing Function User-Friendly Name	Internal Menu Name	Runtime Menu Name
FND_DIAGNOSTICS_TRACE	FND Diagnostics Trace	TRACE <ul style="list-style-type: none"> • NO_TRACE • REGULAR • BINDS • WAITS • BINDS_AND_WAITS • PLSQL_PROFILING 	Diagnostics <ul style="list-style-type: none"> • No Trace • Regular Trace • Trace with Binds • Trace with Waits • Trace with Binds and Waits • PL/SQL Profiling
FND_DIAGNOSTICS_VALUES	FND Diagnostics Values	PROPERTIES_MENU <ul style="list-style-type: none"> • ITEM • FOLDER 	Diagnostics - Properties <ul style="list-style-type: none"> • Item • Folder
FND_DIAGNOSTICS_VALUES_RO	FND Diagnostics Values Read Only	PROPERTIES_MENU <ul style="list-style-type: none"> • ITEM • FOLDER 	Diagnostics - Properties <ul style="list-style-type: none"> • Item • Folder
FND_DIAGNOSTICS_CUSTOM	FND Diagnostics Custom	CUSTOM_CODE_MENU <ul style="list-style-type: none"> • NORMAL • OFF • CORE • SHOW_EVENTS 	Diagnostics - Custom Code <ul style="list-style-type: none"> • Normal • Off • Core Code Only • Show Custom Events

Securing Function Name	Securing Function User-Friendly Name	Internal Menu Name	Runtime Menu Name
FND_DIAGNOSTICS_PERSONALIZE	FND Diagnostics Personalize	CUSTOM_CODE_MEN NU <ul style="list-style-type: none"> CUSTOMIZE 	Diagnostics - Custom Code <ul style="list-style-type: none"> Personalize
FND_DIAGNOSTICS_PERSONALIZE_RO	FND Diagnostics Personalize Read Only	CUSTOM_CODE_MEN NU <ul style="list-style-type: none"> CUSTOMIZE 	Diagnostics - Custom Code <ul style="list-style-type: none"> Personalize

The following table lists seeded permission sets.

Permission Set Name	Permission Set Code	Permissions Assigned
FND Diagnostics Examine Menu	FND_DIAGNOSTICS_EXAMINE_PS	FND Diagnostics Menu Examine
FND Diagnostics Examine Read Only	FND_DIAGNOSTICS_EXAMINE_RO_PS	FND Diagnostics Menu Examine Read Only
FND Diagnostics Custom Menu	FND_DIAGNOSTICS_CUSTOM_PS	FND Diagnostics Custom
FND Diagnostics Personalizations Menu	FND_DIAGNOSTICS_PERSONALIZATION_PS	FND Diagnostics Personalize
FND Diagnostics Personalizations Menu Read Only	FND_DIAGNOSTICS_PERSONALIZATION_RO_PS	FND Diagnostics Personalize Read Only
FND Diagnostics Properties Menu	FND_DIAGNOSTICS_PROPERTIES_PS	FND Diagnostics Values
FND Diagnostics Properties Menu Read Only	FND_DIAGNOSTICS_PROPERTIES_RO_PS	FND Diagnostics Values Read Only

Permission Set Name	Permission Set Code	Permissions Assigned
FND Diagnostics Trace Menu	FND_DIAGNOSTICS_TRACE_PS	FND Diagnostics Trace
FND Diagnostics Menu Developer	FND_DIAGNOSTICS_DEVELOPER_PS	<ul style="list-style-type: none"> • FND Diagnostics Examine • FND Diagnostics Personalize • FND Diagnostics Trace • FND Diagnostics Values • FND Diagnostics Custom
FND Diagnostics Menu Support	FND_DIAGNOSTICS_SUPPORT_PS	<ul style="list-style-type: none"> • FND Diagnostics Examine Read Only • FND Diagnostics Personalize Read Only • FND Diagnostics Trace • FND Diagnostics Values Read Only • FND Diagnostics Custom

To give a user access to a secured menu item using RBAC, grant a permission set containing a permission for that item to a role. Then assign that role to the user.

Example - Granting access to the Examine menu item only

For example, say you want to give a certain user full access to the "Examine..." menu item only. You want to make the other menu items that are controlled by the profile Utilities:Diagnositics inaccessible to that user.

1. Ensure that the profile option Hide Menu Entry is set to No, because that profile controls whether the Diagnostics menu is visible at all.
2. Ensure that the Utilities: Diagnostics profile is set to No, so that the other menu items will be unavailable.
3. Grant the seeded permission set **FND Diagnostics Examine Menu** to a role.

4. Assign the role to the user.

The user should then be able to access the "Examine..." menu item, but not the other items controlled by the Utilites:Diagnostics profile.

Related Topics

Overview of Access Control with Oracle User Management, *Oracle E-Business Suite Security Guide*

Monitoring Oracle E-Business Suite with Oracle Applications Manager

Overview of Monitoring

Using Oracle Applications Manager (OAM), you can monitor components of your Oracle E-Business Suite instance.

For more information on setting up and using Oracle Applications Manager, see *Oracle E-Business Suite Setup Guide*.

The Applications Dashboard

The Applications Dashboard provides a "snapshot" of your Oracle E-Business Suite system. Information is grouped under the following tabs: Overview, Performance, Critical Activities, Diagnostics, Business Flows, Security, and Software Updates.

From the Dashboard you can navigate to the Site Map, or use the drop-down menu to navigate to any of the following pages:

- Application Services
- Configuration - Overview
- Forms Sessions
- Database Status
- Applied Patches
- Patch Wizard
- Workflow Manager

Oracle Applications Manager uses the collection program OAM Applications Dashboard Collection (short name: FNDOAMCOL) to gather the information displayed. The default repeat interval for this program is 10 minutes. To immediately regather the data and update the display for a particular region, click the corresponding Refresh icon. If the OAM Applications Dashboard Collection request is not running when you log in to the Oracle Applications Manager, a request will be submitted automatically under your username.

Note: The status of Web Components is collected manually from the Dashboard.

Overview

This page provides an overview of the general status of your system. It includes the following regions:

Applications System Status

Use this region to view the status of each host machine in your system. The display shows which services are installed on which host machine and the statuses of these services.

Services displayed that represent more than one service component (such as Forms) indicate the status of the worst-case component. For example, if the Forms Listener is down, but the other Forms components are running, the down status will be indicated on this page.

The Database, Concurrent Processing, Forms, and Web status indicators drill down to the Applications System Status page where you can view the status of each individual service.

- Host - the host name.
- Platform - the host's operating system.
- Admin - indicates whether the Admin server has been installed on the host machine.
- Database - indicates the status of the database instance installed on the host machine.
- Concurrent Processing- indicates the status of the Internal Concurrent Manager and the services managed by the ICM.
- Forms - indicates the status of the Forms Server components: Forms Listener, Metrics Server, Metrics Client, and OAM Generic Collection Service.

- Web - indicates the status of the Apache Web Listener.

Configuration Changes

The purpose of this region is to alert you to system-level changes that have occurred in the last 24 hours. Use this data to help diagnose sudden changes in the functioning of your applications system.

To see the list of Patches Applied, click on the number to drill down to the Patch Summary page.

To see the list of Site Level Profile Options, click on the number to drill down to the Site Level Profile Settings page.

To see the list of Applications Context Files Edited, click on the number to drill down to the Applications Configuration Parameters page. Changes made to context files can impact your overall processing configuration and the functioning of business processes.

System Alerts

This region lists the number of system alerts in the categories listed below. If your system is functioning well, there should be no new alerts reported. When an alert of a particular type first occurs, it is counted as a new alert. It remains new until the status is manually changed by the administrator. If an alert of the same type occurs again while the original alert is still in open or new status, it is counted as a new occurrence.

- New Alerts - alerts that have not yet been acknowledged by the administrator. An alert is acknowledged when it is manually moved from a status of "New" to a status of "Open" or "Closed."
- New Occurrences - additional occurrences of alerts that are in new status.
- Open Alerts - all alerts that are in an open status. An alert must be manually moved from the new to open status.
- Open Occurrences - all occurrences of alerts that are currently open. Click on the number for any of these to drill down to the System Alerts and Metrics page.

Web Components Status

This region lists the status of the web components. Status values may be "Up," "Down," or "Warning".

The status of each Web component is determined by testing the corresponding URL as defined in the component's Web agent profile option. The Warning status will be displayed if the profile option is not set. Otherwise, a status of Up or Down will be returned based on the success of the URL test. The profile options are listed with their corresponding components below.

- Servlet Agent - Uses profile option APPS_SERVLET_AGENT (Apps Servlet Agent). If down, the Self-Service Framework-based Applications will not function, as well as all other servlet-based features. Look for errors in the Apache error and access logs (see above for location). Also, execute the Servlet Ping from the System Administration Diagnostics menu.
- JSP Agent - Uses profile option APPS_SERVLET_AGENT (Apps Servlet Agent). If down, execute the JSP Ping from the System Administration Diagnostics menu.
- TCF - Uses profile option APPS_SERVLET_AGENT (Apps Servlet Agent). If down, try running the AOL/J Diagnostic or the Servlet Ping utilities from the System Administration Diagnostics menu.

Applications System Status

This page lists each Applications Server and its status. Each server type expands to display the host name, which expands to display the status of each server component.

Navigation: Applications Dashboard (Overview page) > (drill down on) Database, Concurrent Processing, Forms, or Web column (under Applications System Status)

- Administration
- Database- expands to display the instance name and status. Drill down on the instance name to display the Database Status Details page. Concurrent Processing - expands to display concurrent managers and services controlled by the Internal Concurrent Manager. These expand to display the instances of the managers and services and their statuses. Drill down on the instance names to display the Service Instances page.
- Forms - expands to display the Forms server components: the Forms Listener, the Metrics Server, the Metrics Client, and the OAM Generic Collection Service. The component names expand to display the service instances. Drill down on the instance name to display the Service Instances page.
- Web - expands to display the web component: the Apache Web Listener. The component name expands to display the service instance name. Drill down on the instance name to display the Service Instances page.

Click on the Focus icon for an item to display only its status and the status of its children.

Applications Dashboard - Performance

The Performance region lists Activity and System Throughput indicators. Each of the values listed for Activity and System Throughput links to the related detail page.

Activity

- Forms Sessions - the number of running Forms sessions. Drills down to the Forms Sessions page.
- Database Sessions - the number of active database sessions. Clicking the value runs the Show Active Database Sessions request and returns the results page.
- Running Concurrent Requests - drills down to the Search for Requests Results page showing all currently running requests.
- Service Processes - drills down to the System Activity page. Service processes include all concurrent manager processes and all processes managed by the ICM. If you have set up your system to have other services managed by the GSM, those services are included as well.
- Services Up - the number of service instances whose target services match the actual services. Services Down - the number of service instances whose target services do not match the actual services.
- Invalid Database Objects - drills down to the Invalid Database Objects page displaying the search results for invalid objects owned by the APPS schema.
- Unsent Workflow E-Mail

System Throughput (last 24 hours)

- Completed Concurrent Requests - the percentage of concurrent requests submitted in the last 24 hours that have completed.
- Sent Workflow E-Mail - the percentage of Workflow e-mail sent successfully.

Applications Dashboard - Critical Activities

The Critical Activities region lists concurrent programs that perform maintenance activities. The programs are grouped by activity type and by application. To display only a particular group, click the group's **Focus** icon.

To add or delete a program to the critical activities list, click the **Modify Monitored Program List** button to access the Modify Monitored List page.

To change the frequency that a monitored program is run, click the **Update Frequency** button.

For each critical activity, the following are displayed:

- Program Name - Drills down to the Activity Summary page showing work metrics for those programs that have been instrumented to compute them.

- Request ID - The last run request ID. Drills down to display the request in the concurrent request Search Results screen.
- Last Run Date
- Outcome - indicates the completion status of the request.
- Oracle Recommended Frequency - The frequency that Oracle recommends a critical program be run (if applicable).
- On Schedule (Oracle Recommended) - indicates whether the Oracle recommended schedule has been met (if applicable).
- Onsite Frequency - the frequency that the program is currently scheduled to run. To change the frequency, use the Update Frequency button to access the Update Frequency for Monitored Critical Activities page.
- On Schedule (Onsite Frequency) - indicates whether the onsite schedule has been met.
- Success Rate - the percentage of completed requests that completed with a status of normal. Drill down on the value to display a success rate chart showing completion status percentage rates of Normal, Warning, and Error. Mouse over the chart to display the numeric values.

Modify Monitored List

Navigation: Applications Dashboard > Critical Activities (B) Modify Monitored Program List

Use this page to add or remove programs to the critical activities list.

To add a program to the Monitored list, select the program from the Not Monitored list and click the Move shuttle button.

To add all programs from the Monitored list, click the **Move All** shuttle button.

To remove a program from the Monitored list, select the program from the Monitored list and click the **Remove** shuttle button.

To remove all programs from the Monitored list, click the **Remove All** shuttle button.

Click **OK** to apply your changes.

Update Frequency for Monitored Critical Activities

Navigation: Applications Dashboard > Critical Activities (B) Update Frequency

Use this screen to update the frequency that your critical activity programs are run.

The following are listed for each critical activity program:

- Program Name

- Application
- Program Type
- Oracle Recommended Frequency - the run frequency recommended by Oracle (if applicable).
- Onsite Frequency - the frequency that the program is currently scheduled to run. To change the run schedule for a program, update the **Onsite Frequency** field and click **OK**. Note that this is the target frequency and may not be the frequency that the program actually runs. Monitor the success of the target frequency with the On Schedule (Onsite Frequency) field on the Applications Dashboard - Critical Activities page.

Critical Activities - Activity Summary

Navigation: Applications Dashboard > Critical Activities > [Program Name]

This page displays work metrics for those maintenance programs that have been instrumented to compute them. The display can be filtered by the table name or value.

- Name - the name of the table that will be purged by the program.
- Value - the number of rows in the table that will be purged if the program is run.

Applications Dashboard - Business Flows

Oracle Applications Manager allows you to monitor and support business flows within Oracle E-Business Suite. User-defined key business flows are correlated with the system components responsible for the execution of those flows.

Navigation: Applications Dashboard > Business Flows tab

From the OAM console you can:

- View the hierarchical representation of the business flows.
- Monitor system alerts, errored requests, and errored work items for a business flow.
- View the setup status for the business flows and associated subflows.

The Key Business Flows region displays the current listing of business flows, with these columns:

- Status - Indicates the setup status of the business flow. Business flows that are not fully set up are listed as unavailable
- Edit

To create a new business flow, click **Create**. Click View Details for a selected business

flow to view additional information for that business flow. Click the **Edit** icon for a selected business flow to update it.

Create or Edit a Business Flow

Use these pages to create or edit a business flow.

Navigation: Applications Dashboard > Business Flows tab > Create (B) or Edit icon for a selected business flow

Enter a name and description for the business flow.

Enter in a child flow or component for the business flow. Choose from the following:

- **New Business Flow** - If you select New Business Flow you are prompted for a name and description of the new business flow. You can later update the new subflow with children of its own.
- **Existing Business Flow** - You are prompted to choose a business flow from a list of values.
- **Work Item Type** - You are prompted to choose a workflow item type from a list of values.
- **Component** - Select from Concurrent Program, Service, Form or Function. You are prompted for a component name from a list of values.

View Business Flow Details

This page displays details for a selected business flow.

Navigation: Applications Dashboard > Business Flows tab > View Details (B) for selected business flow

Subflows and components of the business flow are shown in hierarchical format. You can expand or collapse nodes on the hierarchical tree.

Business Flow Monitoring and Setup

Maintain your business flow monitoring from this page.

Navigation: Setup (global icon) > Business Flows (side navigation)

Schedule Requests

OAM provides the following concurrent program to help you maintain your business flow setup. Schedule requests for the concurrent program from the link provided.

- **Metrics Refresh** - schedule requests for the OAM: KBF Metrics Rollup Program to update the setup status of your business flows.

Setup Monitoring

For each of the business flows listed, you can view whether monitoring is enabled and enable or disable monitoring.

Select a business flow and click Update to enable or disable monitoring. Click **View Details** to view if monitoring is enabled.

Applications Dashboard - Security

Information on this page helps you detect and diagnose security issues on your Oracle E-Business Suite System.

Navigation: Applications Dashboard > Security (tab)

Click the **Manage Security Options** button to manage SQL*Net access for your middle-tier hosts.

Security Alerts

Security Alerts can be raised either at runtime by the application code, or at the failure of security-related diagnostic tests. The table is organized by severity, which can be Critical, Error, or Warning. It provides numerical counts of new and open alerts. Where enabled, you can drill down on the numerical links to view and manage the details of an alert and any associated diagnostic test reports. Alert details and test reports can be added to the Support Cart.

Security Test Failures

This table shows security-related diagnostic tests that failed when they were executed. The table specifies the most recent time that the test failed, and provides links that open detailed test reports. For a specific test, clicking the **Diagnose** icon will re-execute the test - this is useful to verify that the error still exists. For a specific application, clicking the Diagnose icon allows you to re-execute all failed tests in that application for the chosen security level.

Resources

Links to security-related documents on My Oracle Support are located here. Documents include:

- Best Practices for Securing Oracle E-Business Suite
- Oracle Support Services Security Alert - Frequently Asked Questions
- Security Announcements and Notes

Security-Related Tests

You can manage Oracle E-Business Suite Diagnostics tests from the Dashboard.

The two key tests accessible from the OAM Security tab are:

- Best Practices: Database Security Tests
- Best Practices: Oracle E-Business Suite Security Tests

These tests are described further in My Oracle Support Document 1337420.1, *Secure Configuration Guide for Oracle E-Business Suite Release 12.2*.

For more information on Oracle E-Business Suite Diagnostics, see: *Oracle Diagnostics Framework User's Guide*.

Manage Security Options

Use this button to access Security Options.

Managing SQL*Net Access from Application Tier Hosts

These pages allow you to restrict SQL*Net access to the database from your application tier hosts. If you enable the SQL*Net Access security option, you can select which hosts have SQL*Net access to the database. If you disable the SQL*Net Access security option, then all application tier hosts have SQL*Net access to the database.

View SQL*Net Access

Use the View SQL*Net Access page to see how SQL*Net Access is currently configured for your middle-tier hosts.

Navigation: Applications Dashboard > Security (tab) > Manage Security Options (B)

If the Manage SQL*Net Access security option is disabled, a message here indicates that it is disabled. All hosts have SQL*Net access to the database in this case.

If this feature is enabled, the table of hosts indicates which hosts have SQL*Net access and which do not.

Note: In order for the information on this page to be accurate, the following steps must be run in addition to enabling or disabling the Manage SQL*Net security option:

- Run AutoConfig on the database tier
- Bounce the TNS Listener

The table shows the hosts that have SQL*Net access and includes the following columns:

- Name

- Platform
- Oracle Applications Host - Indicates whether the host is an Oracle E-Business Suite host or not. Application services (Concurrent Processing, Oracle Forms, Web, Admin, and Database services) can run on Oracle E-Business Suite hosts.

Enable SQL*Net Access

Use the Manage SQL*Net Access wizard to enable or disable SQL*Net access to the middle-tier hosts. You can register a new host and grant it access as well from this wizard.

Disable SQL*Net Access

When you disable the SQL*Net Access security option, you allow SQL*Net access to the database from your middle-tier hosts.

Applications Dashboard Collection

Oracle Applications Manager uses the program OAM Applications Dashboard Collection (short name: FNDOAMCOL) to gather the information displayed on the Dashboard under the Overview and the Performance tabs.

The Dashboard Collection Program can selectively enable and disable monitoring of various metrics, and to raise alerts for services when the service has a specified status. The Dashboard Collection Program can collect data for a metric and then raise an alert when a metric reaches a specified threshold. Note that for most components, you can collect data for monitoring purposes in two different ways: (1) through the Dashboard Collection Program, or (2) manually refreshing the data from a Dashboard page.

Metrics for the following data can be monitored for the following using the Dashboard Collection Program. In addition, data for web components can be collected manually in the dashboard.

Activity

- Forms Sessions
- Database Sessions
- Running concurrent requests
- Service processes
- Services up
- Services down
- Invalid database objects

- Unsent Oracle Workflow e-mail

Configuration changes (made in the last 24 hours)

- Patches applied
- Site level profile options
- Applications context files edited

System Alerts

- New alerts
- New occurrences of an alert
- Open alerts

Alerts can be raised for the following services. When a service attains a specified status, an alert is raised.

- Service instances listed under Applications System Status
- Web Components

System Throughput (in the last 24 hours)

- Completed concurrent requests
- Sent Oracle Workflow e-mail

Additional Monitoring Features and Options

From the Monitoring tab on the OAM Site Map, you can access these utilities.

Service Instances for the Forms Listener

Navigation: Site Map > Monitoring > Forms (under Availability)

This page lists the service instances for the Forms Listeners. From this page you can edit information for a selected service instance. You can also view its status, view processes, and view information on its Forms Runtime Processes. Also, you can start, stop, abort, or restart the instance.

SQL Activity

Navigation: Site Map > Monitoring > SQL Activity (under Performance)

This page provides data regarding SQL Activity:

- SQL_HASH
- Physical Reads
- Logical Reads
- Total Sorts
- Execs
- Total Loads
- Load

For more information on these columns, see the Oracle database documentation.

Concurrent Request Runaways

Main Navigation Path: Site Map > Monitoring (subtab) > Performance (heading) > Concurrent Request Runaways (link)

System performance can potentially be affected by database sessions that should have ended when their corresponding concurrent requests were canceled, but for some reason did not.

If any such database sessions are currently active, they will be reported on this page. The table supplies context information for each session: request ID, AUDSID, program, user name, start time, phase, status, Oracle SPID, and PID. You can delete a session by selecting it in the table and clicking Terminate. You can drill down on the links in the request ID, AUDSID, program, and user name columns to view the respective details.

Forms

The following information is shown:

Forms Sessions

Navigation: Site Map - Monitoring > Forms Sessions (under Current Activity)

This page shows information on the current forms sessions. Every open form has its own database session, or "form session."

The profile option "Sign-On: Audit Level" should be set to 'Form' to use this feature. If this profile option is not set to 'Form', the Forms Sessions table will show an empty table even when there are active forms sessions.

To filter the display by Form Name, Username, Responsibility, or Application, make the appropriate selection from the drop-down menu, enter the search string in the field

provided, and click Go.

The following data is shown for each session:

- Form Name
- AUDSID - The auditing session ID. Click on the value to drill down to the Database Session information page.
- RTI_PID - The runtime instance process ID. Click on the value to drill down to the Forms Sessions for Process ID page.
- Username
- Responsibility
- Application
- LRs (Session Logical Reads) - Input/output (I/O) is one of the most expensive operations in a database system. SQL statements that are I/O-intensive can monopolize memory and disk use and cause other database operations to compete for these resources. To prevent single sources of excessive I/O, Oracle lets you limit the logical data block reads per call and per session. Logical data block reads include data block reads from both memory and disk. The limits are set and measured in number of block reads performed by a call or during a session.
- PRs (Physical Reads) - The total number of data blocks read from the disk for the session.
- CPU
- PGA (Session Program Global Area memory) - The PGA is a memory buffer that contains data and control information for a server process. A PGA is created by Oracle when a server process is started. The information in a PGA depends on the configuration of Oracle
- UGA - User Global Area memory used by the session.
- Duration - in HH:MM:SS

Click on the **Session Details** button or the AUDSID to view database information for the selected forms session.

Use the **Diagnostics On/Off** button to turn on or off the Forms Runtime Diagnostics (FRD) for the runtime process. If this button is disabled, make sure your Forms patchset level is 12 or higher (that is, 6.0.8.20 or higher) and then set the environment variable FORMS60_OAM_FRD for the Forms Listener process.

Forms Sessions for Process ID

If you click on the RTI_PID from the Forms Session window, or if you click on the PID from the Forms Runtime Processes window you will see the fields described above as well as the following data for the Process ID:

- Client IP Address
- Server Host Name
- CPU Time
- Memory Usage (KB)
- Diagnostics (On/Off)
- Log File Name

Use the **View Diagnostics** button to view the Forms Runtime Diagnostics (FRD) log file. The log file can be added to the Support Cart.

Forms Runtime Processes

Navigation: Site Map - Monitoring > Forms Runtime Processes (under Current Activity)

This page shows information about Forms runtime processes. You must first register and start a service instance of the OAM Generic Collection Service to collect this information. The Generic Collection Service must be running for the information to be collected.

You can filter your view by Node or Username.

The following columns are shown for each session:

- PID - The ID of the runtime process for the user session. Click this value to drill down to the Forms Sessions for Process ID page.
- Node
- Port - The Apache port of the servlet listener, if any.
- Memory (KB) - The memory used by the runtime process in kilobytes. For HP and AIX platforms, this is the virtual memory size. For all other platforms, this is the resident set size.
- CPU
- Duration
- Client IP Address - The IP address of the client machine used to connect to the

Forms Services.

- Username - The database username used by the Forms application for the user session.
- Diagnostics - On/Off
- Last Update Time

Use the Upload button to refresh the data on this page.

Use the Terminate button to end a selected process.

Click on the Sessions button or click on the PID to view the Forms Sessions for Process ID page.

This page also shows the runtime processes from the Forms Servlet Listener, if any. The Port column for such processes indicates the Apache Listener port.

Forms Listener versus Forms Listener Servlet

The Forms Listener is a process running on a specific port on the server machine. When the connection between the client and the Forms runtime process is established, the client and the runtime process requires that the connection be persistent.

The Forms Listener Servlet is a Java servlet running in a servlet engine. The Web server routes the client requests for the Forms Listener Servlet directly to the servlet instance. Because the web server acts like the end point for the client, the other server machines and ports are no longer exposed to the firewall.

In the Forms Runtime Processes page, the node name and the port are shown for each runtime process. You can distinguish between the Forms Listener process and Forms Listener Servlet process by examining the port numbers. For the Forms Listener process, the port is the Forms server machine port. For the Forms Listener Servlet process, the port is the web server port.

System Activity (Activity Monitors)

Navigation: Site Map > Activity Monitors (under Activity)

This region displays information on the system's activity.

A Database Sessions graph displays the number of database sessions related to the following:

- Login sessions
- Oracle E-Business Suite forms sessions
- Services
- Requests

A Concurrent Requests graph displays the number of requests with the following statuses:

- Pending
- Running
- Waiting on a lock
- Inactive
- Completed in the last hour

Click on the bar for any status to drill down to more information on requests of each status.

Database Session Information

Navigation: Site Map - Monitoring > Forms Sessions (under Current Activity) > (B) Session Details

This page displays detailed information about the selected database session. Click **Terminate** to end the database session.

Summary

- Form or Service Name
- Username
- Responsibility

Instance Attributes

- Logon Time
- Serial Number
- OS PID
- Status
- Session ID
- Oracle SPID
- User
- SQL Hash - If the value shown is a link, you can click on it to view a page showing

the SQL statement that is currently executing, as well as an execution plan for the statement. For more information on execution plans, see the Oracle database documentation.

Client Attributes

- OS User
- Machine
- Process
- Terminal

Application Attributes

- Module
- Module Hash
- Action
- Program

Session Wait Information

- Event
- Wait Time
- Timeouts
- Average Wait
- Total Wait
- Maximum Wait

Tracing Options

Set the trace options to the level desired. Options available are:

- Normal Trace
- Trace with Waits
- Trace Off
- Trace with Binds

- Trace with Binds and Waits

Click **Apply** to apply any changes made to the Tracing Options. Click **View Trace** to view the current trace information.

Current Activity

The following information is shown:

User Monitoring

The feature monitors an Oracle E-Business Suite user's current activity within the system, with respect to the system components. It monitors the user's current activity within forms and concurrent programs. Sign-on Audit should be turned on for Form Activity data to be available.

Invalid Objects

Navigation: Site Map > Monitoring > Invalid Objects (under Current Activity)

This page lists invalid objects in the database. To remove invalid objects, you can compile the APPS schema (for invalid objects in the APPS schema) or run a script provided with the database (for other invalid objects). See *Maintaining the Database*, page 7-14 for more information on compiling objects.

Forms Runaway Processes

Navigation: Site Map > Monitoring (subtab) > Current Activity (heading) > Forms Runaway Processes (link) Overview

You can also access this page by clicking the **View Runaways** button on the Forms Runtime Processes page.

Running Oracle E-Business Suite requires the creation of many system-level processes. On occasion, processes can behave incorrectly and have a negative impact on system performance. In Oracle Applications Manager, you can:

- Configure thresholds (maximum memory size, maximum CPU percent, maximum duration in minutes) for tracking runaway processes. These settings take immediate effect as soon as you click Apply. These settings are used to raise system alerts on the Applications Dashboard.
- See the user name and IP address of runaway processes.
- Terminate processes.
- See the parameters of the OAM Generic Collection Service (the background process which runs on all Forms nodes).

- Open the associated log file.

You can define memory, CPU, and duration thresholds. Memory refers to process memory size, resident set size, or total virtual memory size based on the platform. On a UNIX system, CPU refers to the cumulative execution time of the process. On a Windows NT system, CPU is, CPMemory - Process memory size, Kb, resident set size or total virtual memory size based on the platform. CPU - On UNIX, it is the percentage of CPU use. If the system has both UNIX and Windows NT nodes, then CPU refers to the percentage of CPU use. In all cases Duration refers to the total time elapsed since a connection was established.

The default values of the thresholds are as follows:

- Maximum memory: 1.0 MB
- Maximum CPU: 25%
- Maximum duration: 20.0 minutes

Applications Usage

Navigation: Site Map > Monitoring (tab) > Applications Usage Reports (under Usage)

The Applications Usage page contains links to the following pages:

- Products Installed
- Applications Users Per Module Summary
- Page Access Tracking and Sign-On Audit: Configuration, Reports
- Applications Usage Reports: Purchase Lines Processed, Order Entry Lines Processed, and more

Products Installed

Navigation:

Applications Systems > (B) Configuration > Products Installed

or

Applications Systems > (menu) Applications Usage > (B) Go > Products Installed

This page lists the following information for Oracle E-Business Suite products:

- Application Short Name
- Application Name
- Version

- Status- A product's status can be Installed, Shared, or Inactive. Installed indicates that the product has been licensed and installed. The Shared status is used for products that other products are dependent upon. Products that are neither Installed nor Shared have an Inactive status.

Application Users Per Module Summary

Navigation: All Applications Systems > (pull down menu) Applications Usage > (B) Go > Application Users Per Module Summary

This page lists the following information for Oracle E-Business Suite modules:

- Application Short Name
- Module Name
- Count - number of current users

You can view details for a particular module by selecting its radio button on the left and clicking the **View Details** button. This takes you to a page that lists the following:

- Module Name
- User Name
- Description of User
- Creation Date of User
- Last Log On Date

Click **Show All** to see a format suitable for printing that lists all users. Within the Show All format, click on **Show Set** to see the table format of the list.

Page Access Tracking and Sign-On Audit

Page Access Tracking and Sign-on Audit tracks the accesses of Oracle E-Business Suite JSPs and Oracle Forms for usage pattern analysis and performance statistics. The Reports screen displays the complete flow of accesses across technology stacks within a user session. It also aggregates collected metrics and display summary statistics.

Applications Usage Reports

Use these reports to collect information on specific applications usage. Your License Management Services analyst may ask you to collect such information, or you can use these reports for your own monitoring.

The following reports can generate information on various licensing metrics in a time period you specify. However, for the purposes of License Management, a twelve (12) month period is used.

Purchase Line Items Processed (Internet Supplier Portal, Purchasing Intelligence, and iProcurement)

These reports generate information for the licensing metric Purchase Line. Purchase Line is defined as the total number of purchase line items processed by the application during a 12 month period. Multiple purchase lines may be created on either a requisition or purchase order or may be automatically generated by other Oracle E-Business Suite programs. For iProcurement, Purchase Lines are counted as all line items on an approved requisition created in iProcurement. For Internet Supplier Portal and Purchasing Intelligence, Purchase Lines are counted as the line items on purchase orders processed through each of those applications. This does not include communication on the same Purchase Order. For each application, you may not exceed the licensed number of Purchase Lines during any 12-month period unless you acquire additional Purchase Line licenses from us. You may acquire a different number of Purchase Line licenses for each program (Number of Purchase Lines for iProcurement could be a smaller number than for Internet Supplier Portal).

For Internet Supplier Portal, use the Suppliers script to generate a list of suppliers and their IDs. You can then use this information when running the Purchase Line Items Processed report for Internet Supplier Portal.

Order Entry Lines Processed (Order Management)

This report is used for the licensing metric Order Line, which is defined as the total number of order entry line items processed by the program during a 12 month period. Multiple order entry line items may be entered as part of an individual customer order or quote and may also be automatically generated by the Oracle Configurator. You may not exceed the licensed number of Order Lines during any 12 month period.

Expense Reports Processed (Internet Expenses)

This report is used for the licensing metric Expense Report, which is defined as the total number of expense reports processed by the iExpenses during a 12 month period. You may not exceed the licensed number of Order Lines during any 12 month period.

Invoice Line Items Processed (Accounts Receivables)

This report is used for the licensing metric Invoice Line, which is defined as the total number of invoice line items processed by the program during a 12 month period. You may not exceed the licensed number of Invoice Lines during any 12 month period unless you acquire additional Invoice Line licenses from us.

Custom Reporting Utilities - SQL Extensions

Use this page to run seeded and custom scripts.

Navigation: Site Map > SQL Extensions (under Others)

Click on the icon in the Focus column to display only those reports from the selected group.

Use the **Hide/Show** icon next to the group name to hide or display the reports contained in the group.

The following columns are shown for each report:

- Name - Click on the name of the report to display the report details.
- Description
- Protected - A "locked" icon indicates that a password is required to submit the report.
- Run Report - Click on the icon in this column to run the report. If a password or parameters are required, the SQL File Details page will display. Otherwise, the output of the report will display in the Results page.

Use the **Reload** button to reload the displayed reports from the metadata file.

Adding Custom Scripts to the SQL Extensions Page

You can have your custom scripts automatically discovered by Oracle Applications Manager and available to run from the SQL Extensions page.

1. Create a new SQL script. Multiple SQL statements are allowed within the same file. For example: a report called "Get Sysdate": sysdate.sql
2. Create a directory called /custom/sql for your custom SQL files under <APPL_TOP>/admin. Your directory structure should look like <APPL_TOP>/admin/custom/sql.
3. Copy your SQL files to <APPL_TOP>/admin/custom/**sql** directory.
4. Now log in to Oracle Applications Manager and navigate to Site Map > SQL Extensions.
5. The discovered SQL files will be under the "DefaultC" group.

After the files are discovered, you can customize the grouping, protection, and execution method of these scripts.

Customizing Automatically Discovered Scripts

To customize the grouping, protection, report format, or drill-downs for your automatically discovered scripts, you must edit **oamcustext.amx** located under <APPL_TOP>/admin/custom/xml.

For each discovered script, the oamcustext.amx file will contain an entity similar to the following example that defines the grouping, protection, and report format:

```
<cReport type="SQL" group="DefaultC">
```

```
<title>sysdate.sql</title> <script name="sysdate.sql" protected="yes"
execMode="SQLPLUS" parameters="unknown">
</script>
</cReport>
```

To move your report to a different group

You can change the group that your report displays under.

1. In the oamcustext.amx file, change the value of "group" to the name of the group you want your report to appear in. For example, to change the group to "Custom Reports", the result would be:

```
<cReport type="SQL" group="Custom Reports">
<title>sysdate.sql</title>
<script name="sysdate.sql" protected="yes" execMode="SQLPLUS"
parameters="unknown">
</script>
</cReport>
```

2. Log in to Oracle Applications Manager and navigate to the SQL Extensions page (Site Map > SQL Extensions).
3. Click the **Reload** button to reload the metadata. Your script will appear under the new group.

To change the protection on your report

You can change the password protection that is set on your report.

1. In the oamcustext.amx file set the value of "protected" to "yes", if you want password protection enabled on your script. Set it to "no" to remove password protection. For example, to set the protection to "no", the result would be:

```
<cReport type="SQL" group="Custom Reports">
<title>sysdate.sql</title>
<script name="sysdate.sql" protected="no" execMode="SQLPLUS"
parameters="unknown">
</script>
</cReport>
```

2. Log in to Oracle Applications Manager and navigate to the SQL Extensions page (Site Map > SQL Extensions).

3. Click the **Reload** button to reload the metadata. Your script will appear with the "unlocked" icon.

To change the report format

1. In the oamcustext.amx file set the value of "execMode" to "SQLPLUS" text format, or set it to JDBC for HTML format. For example, to set the report format to HTML, the result would be:

```
<cReport type="SQL" group="Custom Reports">  
<title>sysdate.sql</title>  
  
<script name="sysdate.sql" protected="no" execMode="JDBC"  
parameters="unknown">  
  
</script>  
</cReport>
```

2. Log in to Oracle Applications Manager and navigate to the SQL Extensions page (Sitemap > SQL Extensions).
3. Click the **Reload** button to reload the metadata.

To provide drill-downs from the results of your script

For reports defined in HTML format, you can provide drill-downs from the results of your script to other Oracle Applications Manager pages. Currently drill-downs are supported for requests based on REQUEST_ID and database session information based on AUDSID.

Example:

Suppose your SQL script returns REQUEST_ID as the first column of the report, you can link it to the Request Details page as follows:

1. Ensure that execMode="JDBC"
2. Add the following to the entry for your SQL script:

```
<keyColumns>  
<column position="1" key="REQUEST_ID"/>  
</keyColumns>
```

Here, position="1" indicates that the REQUEST_ID column is the first column reported by your select statement. Currently the possible values for the key attribute are REQUEST_ID and AUDSID.

The new full entry for your SQL script will look like the following:

```
<cReport type="SQL" group="Custom Reports">
```

```

<title>sysdate.sql</title> <script name="sysdate1.sql" protected="no" execMode="JDBC"
parameters="unknown">
</script>
<keyColumns>
<column position="1" key="REQUEST_ID"/>
</keyColumns>
</cReport>

```

Troubleshooting

- If you try to execute a SQL script and encounter the following error message:
An error has occurred!
<filename>(No such file or directory)
The SQL file does not exist under <APPL_TOP>/admin/custom/sql. Make sure you have copied the file into this directory.
- If your SQL script takes input parameters, ensure that you provide the parameters one per line in the **Input Parameters** text field. The result will contain errors if you do not provide the necessary parameters.

Details of Report

Navigation: Site Map > SQL Extensions >(select report name)

This page displays information based on the report definition. Information may include:

- Description
- Report Format - HTML or Text
- Applications Schema Password - If the report is password-restricted, enter the password here.
- Input Parameters - Enter any required or optional parameters.

You can run the report from this window by clicking the **Run Report** button.

Report Results

Navigation: Site Map > SQL Extensions (Run Report)

The contents and format of this page will vary depending on the report run.

Report results returned in HTML allow you to filter the report by a specific Column value.

Use the **Refresh** button to rerun a report from this page.

Click **Add to Support Cart** to add your report results to the Support Cart.

System Alerts, Metrics, and Logs

Overview of System Alerts, Metrics, and Logs

The System Alerts, Metrics, and Logs screens provide information that can help you diagnose potential problems. For example, configuration issues, overdue routine maintenance tasks, and invalid data can cause serious problems requiring either an automated response or manual intervention.

Oracle E-Business Suite applications can report these potential problems as system alerts to Oracle Applications Manager. These alerts can then be tracked in OAM, and administrators can classify alerts as open or closed, as well as keep notes on the steps taken to resolve underlying problems.

In addition, some problems may be more easily detected through external analysis of performance metrics. External analysis allows for easier comparison of current and historical metric values, consideration of metrics from multiple products and components, and end-user defined exception triggers. Such exceptions could include decreasing transaction throughput for a component or excessive completion times for a business process.

System Alerts

Navigation: Site Map >>Monitoring > System Alerts (under Current Activity)

Components in an Applications System such as concurrent programs, forms, service instances, or functions can post exception messages during specific error conditions as defined by the developer of the component. The term "System Alert" denotes a grouping of such exceptions having the same message. The term "Occurrence" is used to denote each member exception of such a group. Each alert is associated with a Severity (Critical, Error or Warning) and a Category (System or Product).

This page shows a summary of the system alerts as well as a list of new alerts.

Alerts are classified by Severity level:

- **Critical** - the alert indicates that an important business flow is impeded, or that a large number of users is affected.
- **Error** - the alert indicates a less severe, more isolated issue.
- **Warning** - the alert indicates that there may be a negative impact on users or business processes.

Alerts are also marked as New or Open. "New" indicates that the alert has just been

posted in the system. "Open" indicates the alert is being resolved.

In the Summary region, Alerts are grouped according to their severity and status of New or Open. The New or Open column indicates how many alerts of the given severity exist. You can click on the number to drill down to details on the alerts.

When a new exception is posted, if an alert already exists with the same message and is in New or Open state, then the new exception is considered an occurrence of the existing alert. If an alert with the same message does not exist then a new one is created (with the state New) and this exception becomes the first occurrence of this alert. A notification is also sent to subscriptions for the newly created alert.

You can change the state of alerts (along with the associated occurrences) in OAM. You can change the state of a new alert to Open to indicate the exception has been acknowledged and the problem is being resolved. Once the problem is resolved you can change the state of the alert to Closed. You can also add notes to alerts; for example, to indicate how the problem was resolved.

You can search for alerts, search for occurrences, and view the notification setup for alerts using the buttons provided.

System Alert Flood Control

Oracle Applications Manager provides the System Alerts feature to inform system administrators of potential problems in Oracle E-Business Suite. For the Oracle Application Object Library messages logged at the level of Unexpected, OAM can raise system alerts. Ideally, system administrators should actively look at these alerts and close them once issue is resolved. However if for some reason, the alerts are not closed, too many new system alerts can flood the system with alerts, occurrences, business events, and notifications. Oracle E-Business Suite provides a mechanism to control the count of new system alerts to avoid a system alert flood.

By default, the system will raise only 500 new alerts. Once this limit is reached for new system alerts, no new alerts or notifications will be raised and a message will be displayed on System Alert and Metric page. To re-enable the alerting, a system administrator should change the status of existing new alerts from OAM. Oracle E-Business Suite also allows system administrators to change the default threshold by using the System Alert Setup button from System Alert and Metrics page can access this page. From the setup page you can also change the number of occurrences per alert. By default only 50 occurrences per alert is logged.

The setup page also provides control to enable the system alert for a particular severity. If critical severity is selected, only critical alerts will be logged. "None" selection will disable the system alert completely and no new alert will be raised.

System Alert Details

This page displays the details associated with a particular system alert. This page includes the summary information for the alert such as severity, category, state, creation date, and the exception message. The occurrences table summarizes the individual

occurrences for this alert. You can select an occurrence and click **View Details** to drill down to the context details for an individual occurrence.

From this page, you can also change the state of the alert as well as navigate to the **Add Notes** page to add notes to the alert.

Search Alerts

This page allows you to search for alerts by Severity, Category, State and Posted Date. The search results are displayed in the same tabular format as in the New Alerts section in the **System Alerts** page. You can also add notes or change the state of the alerts displayed in the results table.

To search for occurrences from this page, click **Search Occurrences**.

Search Occurrences

This page allows the user to search for occurrences of alerts by various criteria. The query criteria are categorized into the following groups:

- **System Alert** - The criteria in this section pertain to the alert to which the occurrence belongs.
- **Component** - The criteria in this section pertain to the component that logged the occurrence.
- **User and Responsibility** - The criteria in this section pertain to the user and responsibility that used the component that generated the alert.
- **Database Session** - The criteria in this section pertain to the database session associated with the transaction during which the exception was logged.
- **Others** - Additional criteria related to the occurrence.

From the results table on this page, users can drill down to view the context details for each occurrence. In addition, the users can also drill down to view the details for the alert to which each occurrence belongs.

To search for alerts from this page, click **Search Alerts**.

System Alert Occurrence Details

This page displays the entire context information associate with an individual alert occurrence. This page is divided into the following three sections:

- **Summary** - This section displays information associated with the alert to which the occurrence belongs.
- **Context** - This section displays all the context information and is further categorized into the following subsections:

- **Component** - Name and application of the component that posted the alert occurrence.
 - **User and Responsibility** - Username, responsibility, and application for the user who ran the Component that posted the alert occurrence.
 - **Database Session** - Database session ID, database instance, session module, and session action associated with the database session for the transaction during which the alert was posted.
 - **Others** - Miscellaneous information such as session ID, node, security group, processes ID, thread ID (if applicable) and JVM ID (if applicable).
- The third section on this page varies based on the type of the transaction during which the alert occurrence was posted. The following types are possible:
 - **Concurrent Request** - Request ID, concurrent program name, a link to the request log, and a link to the output file are available if the transaction is a concurrent request. You can use the Request ID link to drill down to the request details. In addition, you can drill down to view related system logs to view other log messages that were posted during the same transaction.
 - **Concurrent Process** - If the transaction type was a concurrent process (belonging to a service instance), the service instance name, concurrent process ID, and a link to the manager log can be viewed from this section.
 - **Form** - If the transaction was from a Form, the form name is displayed in this section.
 - **ICX** - If the transaction was of type ICX, then the ICX transaction ID is displayed in this section.

In addition, regardless of the transaction type, users can also drill down to view related system logs to view other log messages that were posted during the same transaction.

System Metrics

Navigation: Site Map > Monitoring > System Alerts (under Current Activity) > Metrics (tab)

Not all exception conditions can be immediately detected directly within an Oracle E-Business Suite component, but are best detected through external analysis. Some are detected by measuring certain criteria, such as decreasing transaction throughput for a component or excessive completion times for a business process. External analysis allows for easier comparison of current and historical metric values, consideration of metrics from multiple products and components, and end-user defined exception triggers. These exceptions are analogous to "events" in Oracle Enterprise Manager

where the user specifies the specific conditions that will trigger an alert.

Simple Search Metrics

You can search for metrics based on **Application**, **Component**, **Posted After** date, or **Posted Before** date.

Advanced Search Metrics

Click on the **Advanced Search** button to search for metrics based on detailed criteria.

This page allows the users to search for metrics based on the context information associated with the metrics. The query criteria are categorized into the following groups:

- **Metrics** - The criteria in this section pertain to the metric itself such as metric code, metric value and date on which the metric was posted.
- **Component** - These criteria pertain to the component that logged the metric.
- **User and Responsibility** - These criteria pertain to the user and responsibility that used the component that generated the metric.
- **Database Session** - These criteria pertain to the database session associated with the transaction during which the metric was logged.
- **Others** - This group contains miscellaneous criteria such as node, security group, process ID, Thread ID, and JVM ID.

From the results table, users can drill down to view the context details for each metric.

System Metrics Results Table

The System Metrics results table shows information on:

- **Component** - the application component. A component is a functional unit, such as a concurrent program, form, or Web Application function.
- **Application** - the owning application of the metric.
- **Metric Code** - the internal name of the metric.
- **Value** - the value of the metric.
- **Metric Type** - the data type of the metric.
- **Time** - the time the metric was taken.

System Metric Details

This page shows the following:

Summary

- Metric Code
- Metric Type
- Metric Value
- Time Posted

Context

- Component:
 - Name
 - Application
- Database Session
 - AUDSID
 - DB Instance
 - Session Module
 - Session Action
- User and Responsibility
 - User
 - Responsibility
 - Application
- Others
 - Session ID
 - Node
 - Security Group
 - Process ID

- Thread ID
- JVM ID

Request Summary

- Request ID - Click on the request ID to view details for the request.
- Request Log - Click **View** to view the request log.
- Program Name - the program name.
- Output file - click **View** to view the output file.

System Logs

Navigation: Site Map > Monitoring > Logs (under Current Activity)

System Logs are messages that are logged by Oracle E-Business Suite system components.

Log messages contain a comprehensive set of context information and are useful for pinpointing and diagnosing system problems. They can have the following levels (listed from most serious to least serious):

- 6 - Unexpected: Used for the failure reporting of internal unhandled software failures. Example: Failed to place order due to NullPointerException
- 5 - Error: Used for the failure reporting of external end user errors. Example: Invalid username/password
- 4 - Exception: Used for the failure reporting of internal handled software failures. Example: User Session timed out
- 3 - Event: Used for high-level progress reporting. Example: Order placed successfully
- 2 - Procedure: Used for API-level progress reporting. Example: Entering or exiting an API
- 1 - Statement: Used for low-level progress reporting. Example: Processing records within an API

The system logs screens allow you to work with log messages that have been saved to the database. Please note that if logging has been configured to store messages in a middle tier file, such log messages will not be visible on the UI screens. Also, if a log message would normally raise a system alert but the message is sent to a file instead of the database, then the system alert will not be raised.

The following topics describe how to work effectively with the system logs screens:

- Performing a Simple Search
- Performing an Advanced Search
- Working With Search Results
- Viewing Log Message Details
- Setting Up Logging

Performing a Simple Search

In a simple search, you can search for log messages based on the following criteria:

- Posted After date
The default value is today's date.
- Posted Before date
The default value is tomorrow's date.
- Component Application
- Component
- Module
- Level

Enter values into the fields as desired and click **Go** to perform a search.

Performing an Advanced Search

To run an advanced search, click the **Advanced Search** button. You can use any combination of the following search criteria:

- Logged From
- Logged To
The default time interval is from 12:00 AM today to 12:00 AM tomorrow.
- Application
- Responsibility
- User

- Log Level
- Module
- Message
- Host
- Java Virtual Machine
- Database Session ID
- Security Group
- Database Instance

On this page, the LOVs only display values that are reflected in existing log messages. For example, the User LOV only shows users who are specified in one or more log messages. It does not show the entire list of Oracle E-Business Suite users. Furthermore, the LOVs are also filtered by any other search criteria you have entered on the page.

Optionally, you can perform searches that depend on the Component Type. In the Component region, select a Type from the drop-down list. The page will refresh to offer additional search fields. For example, for Concurrent Programs, you can search by Concurrent Program Application or Concurrent Program Name.

Working with Search Results

Viewing Search Results

When you perform a search, the System Log Summary table shows how many log messages were returned and how many are at each log level.

Individual log messages are listed in the System Log Details table. For each log message, the sequence number, module, log level, user, and time are displayed. You can drill down on an individual message or on a user to view details.

Downloading Search Results

To download all returned log messages, click the Download All button. (This includes the full range of log messages, not only those displayed on the current page.) The downloadable file is a comma-delimited CSV file.

To download your choice of currently displayed log messages, select them in the table and click the Download button.

Additionally, you can save all search results by clicking the **Add to Support Cart** button.

Viewing Log Message Details

Summary

- **Module:** The unit of code specified in the FND_LOG API call. A module might be a PL/SQL stored procedure, a C file, or a Java class.
- **Level**
- **Time Posted**
- **Message Text**

Context

- **Component:** Name, Application
- **User and Responsibility:** User, Responsibility, Application
- **Database Session:** AUDSID, DB Instance
- **Others:** Session ID, Node, Security Group, Process ID, Thread ID, JVM ID

Request Summary

- **Request ID**
- **Request Log**
- **Program Name**
- **Output File**

Attachment

In the Attachment region, additional context information (such as environment variables or file versions) may be available in some cases.

Optionally, you can add this page to the Support Cart.

Setting Up Logging

Navigation: Site Map > Monitoring > Logs (under Current Activity) > Log Setup (button)

On the Log Setup screen, you can configure logging according to user, responsibility, application, or site. Additionally, you can view any Java System Property settings for the current JVM that may be active. Note that Java System Property settings override all other settings.

Setting Up Logging for Users, Responsibilities, or Applications

The following procedure explains how to set up logging for a particular user. The steps are the same for responsibilities or applications. Note that user settings override responsibility settings, responsibility settings override application settings, and application settings override site settings. In the table, null values indicate that the

setting is to be inherited from the next higher profile level.

1. If the User table is not currently displayed, then click the icon to show it.
2. If there is a blank User Name field, then click the flashlight icon to select a user name. If there is not a blank User Name field, then click the **Add Another Row** button to add an empty row to the table, then select a user name.
3. In the Log Enabled field, select null, Yes, or No. A null value means that the setting will be inherited from a higher level profile value.
4. In the Log Level field, select a log level. Log messages greater than or equal to the specified level will be stored.
5. (Optional) In the Midtier Log File Name field, type in a valid middle-tier file path. If this field is blank, then log messages will be stored in the database. Note: Server PL/SQL messages are always logged to the database.
6. (Optional) In the Module field, enter the module for which you want to enable logging. For example, "fnd%".
7. Click Apply to save your work.

Setting Up Logging for a Site

The following procedure explains how to set up logging for your entire site.

1. In the Log Enabled field, select null, Yes, or No. (A null value means that the setting will be inherited from a higher level profile value.)
2. In the Log Level field, select a log level. Log messages greater than or equal to the specified level will be stored. It is strongly recommended that you choose 4 - Exception, 5 - Event, or 6 - Unexpected. Significant system performance issues may arise if logging is enabled at less than 4 - Exception.
3. (Optional) In the Midtier Log File Name field, type in a valid middle-tier file path. If this field is blank, then log messages will be stored in the database. Note: Server PL/SQL messages are always logged to the database.
4. (Optional) In the Module field, enter the module for which you want to enable logging. For example, "fnd%".
5. Click **Apply** to save your work.

Diagnostics in Oracle Applications Manager

Oracle Applications Manager allows you to run diagnostic utilities from the Diagnostics and Repair tab on the OAM Site Map.

Debug Workbench

Navigation Path: Site Map > Diagnostics and Repair (tab) > Diagnostics (heading) > Debug Workbench (link)

Overview

The Debug Workbench enables you to centrally control and monitor the debugging of Oracle E-Business Suite components. Using the Debug Workbench, you can set up debug rules for system components and view the debug information that has been collected.

The Debug Workbench can be launched from Oracle Applications Manager and from the Standard Request Submission (SRS) form using the button **Debug Options...** By default, this button is disabled. To enable this button, set the Concurrent: Allow Debugging profile option to Y.

Using the Main Debug Workbench Screen

On the main Debug Workbench screen, a table lists summary information (Rule ID, Component Name, and so on) for the debug rules that exist on the system. On this screen, you can:

- Filter the table by component type (Concurrent Programs, Forms).
- Create debug rules.
- Search for past executions of debug rules.
- Delete a debug rule.

Creating Debug Rules

You create debug rules to collect debug information about specific system components.

To create a new debug rule, use the following procedure:

1. On the main Debug Workbench screen, click the **Create** button. This launches a multi-step flow of screens that guide you through the rule-making process.
2. Choose the component type that you want to debug. Optionally, you can enter a comment to describe the rule.
3. Choose the component instance.
4. You must set up at least one debug option. Debug options such as logging level, PL/SQL profiler, SQL trace, and Reports trace are available. For a given rule, you can select any combination of available debug option values.

5. Specify the context and schedule of the rule. You can set a rule to execute for a specific responsibility or user, and to execute either during a specific span of time or for a certain number of repetitions.
6. Review your work and click the **Finish** button to save the new rule.

The new rule will appear on the main Debug Workbench screen.

Client System Analyzer Data Collections

Main Navigation Path: Site Map > Diagnostics and Repair (subtab) > Diagnostics (heading) > Client System Analyzer Data Collections (link)

Overview

In Oracle Applications Manager, you can view the data that has been collected by the Client System Analyzer. For more information about using the Client System Analyzer from the Oracle E-Business Suite, see My Oracle Support Knowledge Document 277904.1.

Tasks

You can perform the following tasks on the main Client System Analyzer Data Collections page:

- Click the refresh icon to update the data displayed in the table.
- Filter the table by user name. To do so, select the desired operator (is, contains, starts with, ends with) from the drop-down list, type a search term into the text box, and click **Go**.
- Select one or more rows of data collections and add them to the Support Cart.
- Select one or more rows of data collections and delete them.
- Sort the table by user name by clicking the Applications User Name column header.
- Sort the table by collection date by clicking the Collection Date column header.
- Click an icon in the View column to see the details of a particular data collection.
- Click the **Add to Support Cart** button to add the page itself to the Support Cart.

Data Collection Details

The default set of collected data is organized into categories as follows.

- Client Identification Information

- OS user name
- Host name
- Domain
- IP address

- Network Configuration and Performance Information
 - Latency
 - Bandwidth
 - Subnet

- Browser and Java Information
 - Browser type
 - JVM vendor
 - JVM version
 - Proxy information

- Hardware Information
- CPU Information
- OS Information
 - OS name
 - OS vendor
 - Base version
 - Update level

- OS Components
- OS Properties
- OS-Registered Software

Troubleshooting Wizards

Oracle Applications Manager provides several wizards:

- Concurrent Manager Recovery
- Service Infrastructure
- Generic Collection Service (GCS) and Forms Monitoring Wizard
- CP Signature

Concurrent Manager Recovery

Navigation: Site Map - Diagnostics and Repair > Concurrent Manager Recovery (under Troubleshooting Wizards)

Use this feature when the Internal Concurrent Manager fails to start.

Click the **Run Wizard** button to start the recovery process. You cannot run this process if the Internal Concurrent Manager is currently running.

If you encounter any problems, each wizard screen can be added to the Support Cart.

Step 1- Active Managers with a Database Session

This screen lists all managers that must be stopped before proceeding with the recovery.

Listed for each manager are:

- CP ID - The Concurrent Program ID.
- Manager - The manager name.
- Node - The node on which the manager is running.
- DB Session ID - Drills down to the Database Session Details screen.
- Session Status
- OS ID
- Started At - The time at which the manager was started.
- Running Request - Drills down to display the request in the Advanced Search for Requests page.

You may want to wait for any requests that are running to complete before you execute the shutdown. Drill down on the Running Request to view it.

Click **Shutdown** to shut down all the listed managers, and then click the **Refresh** icon to verify that they were shut down. If a manager fails to shut down from this page, you can drill down to the **Database Session Details** page and use the **Terminate** button to end the session from there. Return to the **Concurrent Manager Recovery** screen and refresh the page to verify all managers have been shut down before proceeding to the next step.

Step 2 - Managers Deemed Active but Without Database Session

Any processes listed here must be terminated before continuing. Because these processes have lost their database sessions, they must be manually terminated from the command line. Refer to your operating system documentation for instructions on terminating a process from the command line.

After terminating the processes, click **Update** to mark the processes as no longer active in the database table. Click the **Refresh** icon to verify that all processes have been terminated.

Listed for each process are:

- CP ID
- Manager
- Node
- OS PID
- Started At

Step 3 - Reset Conflict Resolution

Click the **Reset** button to reset the listed requests for conflict resolution. This action changes requests that are in a Pending/Normal phase and status to Pending/Standby. Click the **Refresh** icon to verify that all requests have been reset.

You can drill down on the Request ID to view the request in the **Advanced Search for Requests** screen.

Listed for each request are:

- Request ID
- Program
- User

Step 4 - Requests that are Orphaned

This page lists the requests that do not have a manager. If any requests have Active Sessions listed, drill down on the session ID and terminate the session from the

Database Session Details screen. Return to the Concurrent Manager Recovery screen and click the **Refresh** icon to verify that the session is no longer active.

Listed for each request are:

Request ID - Drills down to display the request in the **Advanced Search for Requests** page.

- Parent ID
- Program
- User
- Phase
- Status
- Active Session

Concurrent Manager Recovery Summary

The summary page lists the information collected from the previous steps. After reaching this page, you should be able to restart your Internal Concurrent Manager. If you cannot, retry starting the Internal Concurrent Manager with DIAG=Y, refresh the summary page, add it to the Support Cart with the log files, and send them to Oracle Support.

Log Files Collected - Click on the log file name to view it. The log files can be added to the Support Cart.

Report Summary

- Active Managers with a Database Session
- Managers Deemed Active but Without a Database Session
- Reset Conflict Resolution
- Requests that are Orphaned

Service Infrastructure

Navigation: Site Map > Diagnostics and Repair > Service Infrastructure (under Troubleshooting Wizards)

Using the Service Infrastructure diagnostic wizard, you can examine existing Generic Service Management data to determine potential problems, and update the data to eliminate the issues.

Click **Run Wizard** to begin using the wizard.

Step 1: Active Nodes without a Service Manager

This screen lists any active nodes without a registered service manager. Concurrent processing requires a registered Service Manager on every registered node. If you need to register service managers for the listed nodes, you can click on the **Register** button to do so.

Step 2: Active Concurrent Processing Nodes without an Internal Monitor

This screen lists any concurrent processing nodes that need a registered Internal Monitor. Click the **Register** button to register Internal Monitors for any listed nodes.

Step 3: Service Managers without Active Nodes

This screen lists service managers and Internal Monitors that are registered for deactivated or nonexistent nodes. If you do not plan on using these nodes in the future, these managers, including the Internal Monitor, can be disabled. Click the **Disable** button to disable the managers for a node.

Step 4: Active Nodes with Inactive Service Managers

All active nodes should have active service managers. This screen lists active nodes without active service managers. Click the **Activate** button to activate service manager definitions for the listed nodes.

Step 5: Enabled Service Instances without Workshifts

This screen shows service instances without any workshifts defined. You can add the Standard workshift to the listed service instances using the **Add Workshifts** button.

Step 6: All Nodes should be Uppercased (for Service Instances)

This screen lists any service instances that are assigned to a node that does not have an uppercase name. Use the **Uppercase** button to change the names of the listed nodes to uppercase.

Step 7: All Nodes should be Uppercased (for Processes)

This screen lists any processes on nodes that do not have an uppercase name. Use the **Uppercase** button to change the names of the listed nodes to uppercase.

Service Infrastructure Summary

This screen shows a summary of the data found for each of the previous screens, as well as any changes you made.

Configuration and Log files are listed first. Two log files and two configuration files are listed for each node. You can click on the name of the file to view it and add it to the Support Cart. You can add all the files to the Support Cart using the **Add All Files to Support Cart** button.

Generic Collection Service (GCS) and Forms Monitoring Wizard

Navigation Path: Site Map > Diagnostics and Repair (tab) > Troubleshooting Wizards (heading) > GCS and Forms Monitoring (link)

Overview

The GCS and Forms Monitoring wizard helps you troubleshoot the OAM Generic Collection Service.

Prerequisites

The wizard cannot be launched unless the Internal Concurrent Manager (ICM) is up and running.

Running the Wizard

Click **Run Wizard** to start the wizard. The steps in the wizard are as follows:

1. If necessary, register the OAM Generic Collection Service on all listed nodes.
2. If necessary, enable the OAM Generic Collection Service on all listed nodes.
3. If necessary, activate the OAM Generic Collection Service on all listed nodes.
4. See the registration of the Forms Listener.
5. If necessary, enable the Forms Listener on all listed nodes.
6. If necessary, set the Sign-On Audit level to "FORM".
7. See a summary screen where you can view a log file and add files to the Support Cart.

CP Signature

The CP Signature Wizard collects information regarding the current status of concurrent processing on the system.

Navigation: Site Map > Diagnostics and Repair > CP Signature

This wizard collects information on the following:

- Configuration status for Parallel Concurrent Processing, Real Application Clusters, and Generic Service Management
- Registered nodes
- Concurrent processing package versions
- Concurrent processing package errors
- Concurrent processing profile options

- Service instances that could be managed by concurrent processing
- Concurrent processing processes
- Request processing manager specialization rules
- Request Conflict Resolution
- Concurrent request processing statistics
- Recent requests to run the Purge Concurrent Request and/or Manager Data program

Support Cart

The Support Cart feature allows you to save Oracle Applications Manager pages with their data and then zip them up in a file to send to Oracle Support. Oracle Support can then view your pages in the Oracle Applications Manager display format.

When you click the **Add to Support Cart** button, the page is added to the Support Cart. If you have filtered or sorted the data, your manipulated view is submitted.

For example, these are some of the pages with the Support Cart feature:

- Configuration Overview
- Site Level Profile Settings
- Recommended/Mandatory Initialization Parameters
- ICM Environment
- Products Installed
- Invalid Objects
- Concurrent Manager Recovery
- Report Results
- All log files

To view the contents of the Support Cart, click on the **Support Cart** global button.

Click **Save Cart** to save the contents to a zip file that you can send to Oracle Support.

Any contents of the cart that are not saved are automatically deleted when you log out of Oracle Applications Manager.

To restore a saved cart, click **Restore Cart** to browse your directory for the saved cart.

To restore a cart file, select a cart file from the list displayed, or use **Browse** to select a file from the directory. Then click **Restore**.

Support Cart Contents

Description

Enter a TAR Number and additional details for the Support Cart Contents.

Applications Signature

The Support Cart can collect a standard set of information regarding your E-Business Suite system. Oracle Support requires this information when logging a technical assistance request (TAR).

To collect this information, click **Collect**.

In the **Generic** region, information is collected on:

- Product information - For each product, the version, current patch level, and status (for example, "Installed") is shown.
- Database parameters - The init.ora parameter settings.
- Patches - For each individual patch applied, the patch number, type (for example, "Patch Set" or "Maintenance Pack"), and application timestamp is shown.
- Topology - This page includes data about all the nodes of the applications infrastructure. For each node, it collects information about the operating system and the different services running on that node.
- Database version

Click the **View** icon to view these pages. If you want to delete a page, select it and click the **Delete** button. Clicking **Collect** again will collect information for all four pages again.

In the **Nodes** region, you can specify to include or exclude output and log files for specific nodes as well.

Other Information Collected

Pages that you save using the **Add to Support Cart** button are listed under this tab.

Oracle Applications Manager Log

This page displays the log file generated by Oracle Applications Manager.

Navigation: Site Map > Administration > Applications Manager Log (under Others)

The current message level of the log is shown. To change the level, select the desired

option and click **Go**.

Note: Changing the log level from this page will only be effective until the servlet is restarted. For a persistent setting, the log level initialization parameter must be changed in zone.properties. The parameter is: oracle.apps.oam.logger.level

For example:

```
servlet.weboam.initArgs=oracle.apps.oam.logger.level=USER
```

Bounce Apache/Jserv for your changes to zone.properties to take effect.

The possible settings are:

- **USER** - includes messages related to Oracle Applications Manager initialization routines, trace information about the error message, and any diagnostic messages related to customizations or extensions that have been added.
- **SUPPORT** - includes the User level messages and additional information useful to support for diagnosing problems (for example, configuration setting details, prerequisite patch-related issues, and module-related information).
- **DEV** - (Development) includes trace information related to code paths (for example, "Inside method A") and any code-related information that could be useful to the developer to diagnose a problem. This level also includes performance-related log messages.

The default is **USER**.

The log can be added to the Support Cart.

Purging in Oracle Applications Manager

Navigation: Site Map > Maintenance > (Critical Activities) Setup and Monitor

Purge programs help reduce the amount of transient data stored in an Oracle E-Business Suite system. Periodically purging unneeded data helps to:

- Reduce system downtime for upgrades
- Decrease backup times
- Increase storage efficiency
- Improve system performance

Oracle E-Business Suite has several concurrent programs defined as purge programs. These programs can then be added to the Critical Activities by navigating to the Setup link. These features can then be run from the Critical Activities Monitor link.

Oracle Workflow Manager

Oracle Workflow Manager Overview

Oracle Workflow Manager is a component of Oracle Applications Manager that allows system administrators to manage Oracle Workflow for multiple Oracle E-Business Suite instances from a single console.

Using Oracle Workflow Manager, administrators can control Workflow system services, such as notification mailers, agent listeners, and other service components, background engines, purging obsolete Workflow data, and cleanup of the Workflow control queue. Administrators can also monitor work item processing by viewing the distribution of all work items by status and drilling down to additional information. Additionally, they can monitor event message processing for local Business Event System agents by viewing the distribution of event messages by status as well as queue propagation schedules. With this ability to monitor work items and event messages, a system administrator can identify possible bottlenecks easily.

To access Oracle Workflow Manager, log into Oracle Applications Manager and select an applications system. Then, you can follow one of the following navigation paths:

- Choose Workflow Manager from the pull-down menu in the Applications Dashboard page and click the Go button.
- Choose Site Map, choose the Administration tab, and then choose the Home link in the Workflow region of the Site Map page. You can also choose one of the other links in the Workflow region to navigate directly to the corresponding page within Oracle Workflow Manager.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go

You can also use other features to help manage Oracle Workflow.

- Use Oracle Diagnostics Framework to run diagnostic tests that check the setup of your Oracle Workflow installation and review debugging information.

- Use Oracle E-Business Suite Logging to review Oracle Workflow logs. Oracle Workflow uses the Oracle E-Business Suite Logging framework to standardize and centralize in the database logging activities related to the Oracle Workflow Business Event System and Oracle XML Gateway.

Note: The Java middle tier components of Oracle Workflow, including notification mailers and agent listeners, also use Oracle E-Business Suite Logging; however, due to the high volume of messages that pass through these components, their information is logged to the file system by default.

Gathering Oracle Workflow Statistics

Some Oracle Workflow Manager graphs and lists may summarize large volumes of data, depending on the level of activity in your Oracle E-Business Suite instance. To enhance performance in displaying these statistics, Oracle Workflow Manager periodically runs concurrent programs to gather the statistics and displays the graphs and lists based on the latest data from the concurrent programs.

- Workflow Agent Activity Statistics Concurrent Program (FNDWFAASTATCC) - Gathers statistics for the Agent Activity graph in the Workflow System status page and for the agent activity list in the Agent Activity page.
- Workflow Mailer Statistics Concurrent Program (FNDWFMLRSTATCC) - Gathers statistics for the throughput graph in the Notification Mailer Throughput page.
- Workflow Work Items Statistics Concurrent Program (FNDWFWITSTATCC) - Gathers statistics for the Work Items graph in the Workflow System status page, for the Completed Work Items list in the Workflow Purge page, and for the work item lists in the Active Work Items, Deferred Work Items, Suspended Work Items, and Errored Work Items pages.

These concurrent programs are scheduled to run every 24 hours by default. They do not require any parameters. You can optionally cancel the default scheduled requests and run the programs with a different schedule if you want to gather statistics at a different frequency.

Each of these graphs and lists displays the date and time when its statistics were last updated, as well as a refresh icon that you can select to refresh the statistics immediately if necessary. However, note that if your Oracle E-Business Suite instance contains very large volumes of workflow data, you may encounter delays or page timeouts when refreshing the data.

Note: Oracle Workflow Manager statistics that typically represent smaller volumes of data, such as work item details and work item

activity details, are queried directly rather than through the concurrent programs.

Oracle Workflow System Status

The Workflow System status page provides a high-level view of the status of your Oracle Workflow instance. The page displays the date and time when the system status information was last updated. To refresh this information, click the refresh icon. To add the information from this page to your support cart, click the Add to Support Cart button.

Note: The system status information is queried directly, separately from the concurrent programs that gather other Oracle Workflow statistics.

The Workflow System status page shows the up, down, or unavailable summary status of the following Workflow features:

- Notification Mailers - To manage notification mailer service components, click the Notification Mailers status icon.
- Agent Listeners- To manage agent listener service components, click the Agent Listeners status icon.
- Service Components - To manage all types of service components, click the Service Components status icon.
- Background Engines - To view Workflow Background Process concurrent requests, click the Background Engines status icon.
- Purge - To view summary information about Purge Obsolete Workflow Runtime Data concurrent requests and completed work items, click the Purge status icon.
- Control Queue Cleanup - To view Workflow Control Queue Cleanup concurrent requests, click the Control Queue Cleanup status icon.

For service component features, including notification mailer service components, agent listener service components, and all types of service components grouped together, the summary status icons represent the following statuses:

- Down - At least one service component of this type has a status of Stopped with Error or System Deactivated. You should investigate the error.
- Up - At least one service component of this type has a status of Running or Suspended, and no service components of this type have a status of Stopped with Error or System Deactivated.

- Unavailable - No service components of this type have a status of Running, Suspended, Stopped with Error, or System Deactivated. For example, if all service components of this type either have not yet been completely configured, or have stopped without errors, then the Unavailable summary status is displayed.

To submit a concurrent request for a feature that runs as a concurrent program, choose the program you want from the Submit Request For pull-down menu and click the Go button. You can submit requests for the following programs:

- Background Engines
- Purge
- Control Queue Cleanup

Related Database Parameters

This region displays information about database initialization parameters required for Oracle Workflow. For each parameter, the list shows the parameter name, actual parameter value, recommended value, and description. If the actual value does not match the recommended value, the recommended value is marked with a warning indicator icon.

The JOB_QUEUE_PROCESSES parameter defines the number of job queue processes for your instance. Oracle Workflow requires job queue processes to handle propagation of Business Event System event messages by AQ queues and for notification mailers. The recommended number of processes for Oracle Workflow is ten or more.

Note: In Oracle Database 10g and higher, you do not need to set the AQ_TM_PROCESSES parameter.

Workflow Metrics

This region displays summary information about work items and Business Event System agent activity.

Work Items

This graph displays the distribution of all work items with the following statuses: Active, Deferred, Suspended, and Error.

- To show this graph if it is hidden, click the Show link.
- To hide this graph if it is shown, click the Hide link.
- The graph header displays the date and time when the work item statistics were last updated. To refresh this information, click the refresh icon. See: Gathering Oracle Workflow Statistics, page 15-2.

- To view the distribution of item types within a status, either click the bar for that status in the graph, or click the status name link.
- To view the number of work items with a particular status, position the mouse pointer over the bar for that status in the graph.

Note: A work item can be counted in more than one status. For example, all work items that do not have an end date are counted as Active work items, including deferred, suspended, and errored work items as well as running work items. Also, if an activity within an item is deferred, and the work item as a whole is suspended, the work item is included in the count for both the Deferred and Suspended statuses. Consequently, the total of the counts for all the statuses is greater than the actual number of work items.

Agent Activity

This graph displays the distribution of all event messages on Business Event System agents with the following statuses: Ready, Waiting, Expired, Undeliverable, and Error.

Note: Messages are not explicitly assigned a status of Error. The Error bar in the graph represents messages of any status on the WF_ERROR agent.

- To show this graph if it is hidden, click the Show link.
- To hide this graph if it is shown, click the Hide link.
- The graph header displays the date and time when the agent activity statistics were last updated. To refresh this information, click the refresh icon. See: Gathering Oracle Workflow Statistics, page 15-2.
- To view the distribution of event messages with different statuses on different agents, either click the bar for a status in the graph, or click an event message status name link.
- To view the number of event messages with a particular status, position the mouse pointer over the bar for that status in the graph.

Related Links

This region provides links to other Oracle Workflow management features.

Configuration

Click the Service Components link to configure service components, including

notification mailers and agent listeners.

Click the Queue Propagation link to view database initialization parameters required for queue propagation and a list of propagation schedules for Business Event System agents.

Throughput

- Click the Work Items link to view the distribution of completed work items across different item types.
- Click the Notification Mailers link to view the notification mailer throughput. This graph shows the throughput of the notification mailers by displaying the distribution of notifications in the WF_NOTIFICATIONS table with the following statuses:
 - Processed - Outbound notifications for which an e-mail message has been sent by a notification mailer service component.
 - Waiting - Outbound notifications for which an e-mail message has not yet been sent.

The graph header displays the date and time when the notification mailer throughput statistics were last updated. To refresh this information, click the refresh icon. See: Gathering Oracle Workflow Statistics, page 15-2.

To view the number of notifications with a particular status, position the mouse pointer over the bar for that status in the graph.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Related Links > Throughput > Notification Mailers

- Click the Agent Activity link to view the distribution of event messages with different statuses on different agents.

Service Components

The Generic Service Component Framework helps to simplify and automate the management of background Java services. Service component containers and their service components are run through Generic Service Management (GSM), which you can control through Oracle Applications Manager (OAM).

A service component container is an instance of a service that manages the running of the individual service components that belong to it. The container monitors the status of its components and handles control events for itself and for its components. These actions are recorded in a log for the container.

A service component is an instance of a Java program which has been defined according to the Generic Service Component Framework standards so that it can be managed through this framework. Currently, Oracle Workflow provides four service

component types: Workflow Mailer, Workflow Agent Listener, Workflow Java Agent Listener, and Workflow Web Services Outbound.

Oracle Workflow provides several seeded service components of these types, within seeded containers, to perform standard processing. You can optionally create additional service components to perform custom processing. If you create custom service components, you can either assign them to the seeded containers, or, based on the volume to be handled by the seeded containers, you can also choose to create your own custom containers.

All service components have certain attributes required by the Generic Service Component Framework. General definition attributes for a component include the component name, startup mode, container type, inbound agent, outbound agent, and correlation ID. Detail attributes include the container that owns the component, the maximum idle time for an on-demand component, maximum error count, number of inbound and outbound processing threads, component log level, read timeout period, minimum sleep time, maximum sleep time, error sleep time, and whether to close connections when the read timeout period expires.

A service component can have one of three startup modes.

- Automatic - When a component container is started, it will automatically start its automatic service components. It will also monitor these components and restart them automatically when necessary.
- On-Demand - A component container will start its on-demand service components if those components have messages waiting to be processed. For example, an on-demand notification mailer service component will be started if there are messages waiting on the WF_NOTIFICATION_OUT queue. The container will stop an on-demand service component when that component's maximum idle time has been exceeded.
- Manual - You must manually start and stop the service component through Workflow Manager. The component container does not start or stop its manual service components.

All service components use the Oracle Applications GSM container type. A component can have either an inbound agent to process inbound messages, an outbound agent to process outbound messages, or both. An Oracle Advanced Queuing (AQ) correlation ID can be assigned to a component to limit its processing to only messages marked with that correlation ID.

Oracle Workflow provides three predefined containers in which you can create components, the Workflow Mailer Service, the Workflow Agent Listener Service, and the Workflow Document Web Services Service. For an on-demand service component, you can specify the maximum amount of time that the service component can remain idle before it is stopped by its container. A service component can have either one inbound processing thread, to enable inbound processing, or none, to disable inbound processing. A service component can have one or more outbound processing threads, to

enable outbound processing depending on the volume of outbound messages, or none, to disable outbound processing. Some types of service components perform only inbound processing or only outbound processing. For example, agent listeners only process inbound event messages and consequently should always have an outbound thread count of zero.

A diagnostic log is recorded for each component container, from the time the container starts to the time it stops. When a container is restarted, a new log is begun. You can view the log through Workflow Manager. Each log entry is marked with the container ID, and, if applicable, with the ID of the service component that generated it. You can specify the level of detail of the information you want to record for each component container. You can also specify a separate log level for an individual service component within the container. The log levels you can select, in order from most detailed to least detailed, are as follows:

- 1 - Statement
- 2 - Procedure
- 3 - Event
- 4 - Exception
- 5 - Error
- 6 - Unexpected

The default log level for both containers and service components is Error. This level is the recommended setting for normal usage.

A processing thread for a service component runs in a loop in which it reads messages from the queue associated with its assigned agent and then waits during a specified amount of sleep time before checking the queue for messages again. The read timeout period defines the amount of time the service component continues attempting to read messages from the queue, after the last message has been dequeued, before timing out. If another message is received before this time expires, that message is processed and the timeout period begins again. If the timeout period expires and no more messages have been received, the service component stops reading and its sleep time begins.

The minimum sleep time for a service component defines the minimum amount of time during which the service component waits, after its read timeout period expires, before it checks the queue for messages again. If a queue receives messages infrequently, you can choose to increase the sleep time between read attempts when no messages are received by setting a maximum sleep time greater than the minimum sleep time. In this case, the service component initially waits for the minimum sleep time after it finishes reading messages from its queue. If no messages are read in subsequent attempts, then the sleep time between read attempts gradually increases until the maximum sleep time is reached. Increasing the sleep time can help enhance performance if messages are received infrequently. You can also set the maximum sleep time parameter to 0 (zero) to

indicate that the sleep time should not be increased. In this case, the service component always waits for the minimum sleep time between read attempts.

The error sleep time for a service component defines the amount of time during which the service component waits, after an error occurs, before it attempts to begin processing again. Additionally, a service component processing thread can either close its connections after its read timeout period expires, when its sleep time begins, or the connections can remain open until the processing thread stops.

A service component may also have additional configuration parameters that are specific to the type of processing it performs. For example, a notification mailer service component has configuration parameters to specify the inbound and outbound e-mail servers it uses.

Among both the common and the type-specific configuration parameters, some parameters can be refreshed dynamically while a service component is running. These parameters are identified by a refresh icon in the configuration pages for the component. For example, the component log level, inbound thread count, and outbound thread count are refreshable parameters.

The control events you can perform for a service component include:

- Starting a service component
- Suspending a running service component, so that the threads stop processing but connections are not closed
- Resuming a suspended service component
- Refreshing a running service component with changed parameters
- Stopping a running or suspended service component

A service component may also have additional control commands that are specific to the type of processing it performs. For example, Workflow Mailer components include a command to launch summary notifications.

You can perform these control events manually at runtime by choosing the relevant command for the component in the Service Components page. You can also schedule single or repeating control events when you are configuring a service component.

A service component can have one of the following statuses.

- Not Configured - Some required configuration parameters for the component have not been completed. The component cannot be started until its configuration is complete.
- Starting - The component is preparing to run.
- Running - The component is running normally. You can choose to suspend processing for a component in this state, refresh the configuration parameters for

the component that are dynamically refreshable, or stop the component.

- Suspending - The component is preparing to suspend its processing.
- Suspended - The component's thread has stopped processing, but its connections remain open. When a component is suspended, you can either resume its processing or stop it altogether.
- Resuming - The component is preparing to resume processing and return to a Running status.
- Stopping - The component is preparing to stop running.
- Stopped - The component was stopped normally, without errors.
- Stopped with Error - The component reached the maximum number of errors specified in its Max Error Count parameter and has stopped. The component container will restart an automatic component in this status, or an on-demand component in this status that has messages waiting to be processed.
- System Deactivated - An automatic or on-demand component was deactivated automatically by its container because the component was stopped with an error the maximum number of times specified in the container's SVC_COMP_MAX_ERROR_COUNT service parameter. A component in this status will not be restarted automatically until the container is restarted.
- User Deactivated - An automatic or on-demand component was manually stopped by a user. It will not be restarted automatically. If you want to restart it, you must do so manually.

A component with a status of Starting, Running, Suspending, Suspended, Resuming, or Stopping is considered to be active. While a component is active, you cannot edit the component name, startup mode, container type, inbound agent, outbound agent, correlation ID, container, or, for an on-demand component, the maximum idle time. You must stop the component before you can change these attributes. However, you can edit the component's other configuration parameters while it is active. If you edit any refreshable parameters, the component will be dynamically refreshed with the new parameter values.

You can manually stop a component from any status. Also, if a container stops for any reason, all of its components are stopped as well.

If the status of a service component changes to Stopped with Error or System Deactivated, Oracle Workflow posts a system alert to the System Alerts and Metrics page in Oracle Applications Manager.

Viewing Service Components

The Service Components page shows the service components that are defined in your Oracle Workflow installation.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Service Components status icon

To add the information from this page to your support cart, click the Add to Support Cart button.

For each service component, the list displays the service component name, status, type, startup mode, container type, and container. Click any column heading to sort the list by that column.

- To filter the service components displayed in the list, select a service component property from the Filter pull-down menu, enter a filter value in the text field, and click the Go button. You can filter by the following properties:
 - Service component name
 - Service component status
 - Service component type display name
 - Service component type internal name
- To verify that the statuses displayed for the service components in the list are current, click the Verify All button.
- To create a new service component, click the Create button.
- To edit a service component's configuration, select the service component and click the Edit button. The steps to edit the configuration depend on the service component type.
- To view the diagnostic log of the service component container in which this service component is running, select the service component and click the View Log button. The log includes log messages for this component and any other component belonging to that container.
- To view details about a service component, either click the service component link in the Name column, or select the service component and click the View Details button. The information that is displayed depends on the service component type.
- To review the events that have been scheduled to control the running of the service component, click the View Event History button. For each event, the Event History page displays the event name, status, user who requested the event, component

status before the event was processed, date the event processing was completed, container for the service component, container type, and any event parameters for a refresh event. You can use this event history as an audit trail to review who scheduled control events for the service component. The status of an event may be Pending, Skipped, In Progress, Completed, or Error. In some cases, an event may be skipped if the component is not in an appropriate status at the time for which the event is scheduled. For example, a refresh event cannot be executed if the component is stopped at the scheduled time.

- To delete a service component, select the service component and click the Delete button. If the service component is currently active, you must stop it before you can delete it.

Note: Several of the seeded service components are required by Oracle Workflow and Oracle XML Gateway and cannot be deleted. If you want to disable them, you can stop them manually using the Stop command from the command pull-down menu. However, note that stopping these components disables the features they support. For example, stopping the Workflow Error Agent Listener and Workflow Java Error Agent Listener disables error handling for the Business Event System.

- To manually control the running of a service component, select the service component, choose the command you want from the command pull-down menu, and click the Go button. You can choose the following commands:
 - Refresh
 - Resume
 - Start
 - Stop
 - Suspend
 - Launch Summary Notifications (Workflow Mailer service components only)
- To manage the service instances for the container of a service component through GSM, click the container link in the Container column.

Creating Service Components

The Pick Component Type page lets you choose the type of service component you want to create. This page lists the name and description of each available type. Select the type that you want and click the Continue button. The steps to complete the service

component configuration depend on the type you select.

Oracle Workflow provides the following service component types.

- Workflow Mailer - Service components that perform send and respond e-mail processing for the Notification System.
- Workflow Agent Listener - Service components that process inbound messages on Business Event System agents in the database.
- Workflow Java Agent Listener - Service components that process inbound messages on Business Event System agents in the middle tier.
- Workflow Web Services Outbound - Service components that process outbound Web service messages.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Service Components status icon > Create

Reviewing Service Component Details

The Component Details page lets you review the configuration of a service component.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Service Components status icon > (B) View Details

The Component Details page displays the configuration parameters defined for the service component and any special status information, as well as the control events that are currently scheduled for the service component. For each event, the list shows the event name, initial start time, whether the event is currently running, the next scheduled execution time for a repeating event, the last previous execution time for a repeating event, the interval in minutes between executions of a repeating event, the number of times the event has failed, the job ID of the DBMS job used to schedule the event, and the PL/SQL API that DBMS job runs.

- To add the information from this page to your support cart, click the Add to Support Cart button.
- For Workflow Mailer service components only, to send test messages, click the Test Mailer button. In the Test Notification Mailer page, select the recipient role to which the messages should be sent, and click the Send Test Message button.

Note: To send a test message successfully, you must select a recipient role that either has a valid e-mail address defined, or that has members with valid e-mail addresses defined. The recipient role must also have a notification preference that includes individual e-mail notifications.

If you set an override e-mail address for the notification mailer, the

Test Notification Mailer page displays that address. In this case the test message is sent to the override address rather than the e-mail address of the recipient role. However, you must still select a recipient role to enable the notification mailer to send the test messages.

Oracle Workflow sends two test messages to the recipient role: one message with content built using PL/SQL and one message with Oracle Application Framework content. Check the e-mail account for the recipient role to view the test messages and reply to them with the Acknowledge response. If you did not implement inbound e-mail processing for this mailer, use the Worklist pages to respond to the test messages after viewing the outbound messages in e-mail. After you acknowledge both test messages, Oracle Workflow sends a confirmation message to the same recipient role to show that the notification mailer successfully processed the inbound response e-mails.

If you do not receive the test messages or the response confirmation message, or if the message content does not appear correctly, check the notification mailer setup, including the mail servers and the mailer configuration parameters. In particular, if the Oracle Application Framework content does not appear correctly, check the Application Framework Agent and WF: Workflow Mailer Framework Web Agent profile options, as well as the Framework User, Framework Responsibility, Framework Application ID, and Framework URL Timeout parameters in the advanced configuration wizard. See: Setting Up a Notification Mailer, page 15-22 and Message Generation, page 15-48.

Note: Oracle Workflow sends the test messages by launching the PLSQL/OAFwk Response Test Process in the System: Tests (WFTTESTS) item type. This item type is stored in a file called wftstmlr.wft in the \$FND_TOP/import/<lang> subdirectory. You can optionally use the Status Monitor to check the status of the test process.

- For Workflow Mailer service components only, to set an override address where you want to send all outgoing e-mail notifications, click the Set Override Address button. Use an override address when you are testing workflow definitions or mailer processing so that you can automatically receive all the test notifications at one e-mail address, instead of having to check or change each individual recipient's e-mail address. To ensure that the override address is accessible and that its use is authorized, you must verify the request before the notification mailer can use the address.

In the Set Override Address page, review the current override address, if any. Enter the e-mail address you want to set as the new override address, and choose Submit. Then check the e-mail account you specified for the verification e-mail message.

In the Verify Override Address page, enter the verification code shown in the e-mail message, and choose Apply. If necessary, you can use the link provided in the verification e-mail message to navigate back to the Verify Override Address page. You must log in to Oracle Applications Manager before you can access this page.

To remove the override address, navigate to the Set Override Address page and choose the Clear Override Address button. The notification mailer then resumes sending e-mail notifications to the individual recipients' e-mail addresses.

- To review the events that have been scheduled to control the running of the service component, click the View Event History button. For each event, the Event History page displays the event name, status, user who requested the event, component status before the event was processed, date the event processing was completed, container for the service component, container type, and any event parameters for a refresh event. You can use this event history as an audit trail to review who scheduled control events for the service component. The status of an event may be Pending, Skipped, In Progress, Completed, or Error. In some cases, an event may be skipped if the component is not in an appropriate status at the time for which the event is scheduled. For example, a refresh event cannot be executed if the component is stopped at the scheduled time.
- To view the diagnostic log of the Generic Service Management (GSM) service component container in which this component is running, click the View Log button. The log includes log messages for this component and any other component belonging to that container.
- To change the values of the configuration parameters or the scheduled events, click the Edit button and navigate to the appropriate page within the service component configuration wizard.
- To return to the Service Components page, click the OK button.

Service Instances for Service Component Containers

You can use Oracle Applications Manager to control service component containers as service instances of type Generic Service Component Container in GSM.

Editing Service Parameters for a Container

Among other properties, a GSM service instance can have work shifts assigned to it. A work shift in turn can have service parameters associated with it. For a service instance that is a service component container, these service parameters apply to the container as a whole to determine how the container manages the components that belong to it.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Service Components status icon > container link > (B) Edit > (B) Edit Service Parameters

The Edit Service Parameters page initially displays the service parameters that can be specified for a container in the Edit Service Parameters field, together with their seeded default values. In most cases, you do not need to change these values. However, you can optionally edit these values in the Edit Service Parameters field if you choose.

You can also optionally delete any of the service parameters from the Edit Service Parameters field. In this case, for all parameters except the proxy setting parameters, the parameter values are obtained from the global settings stored in the WF_RESOURCES table. The default values in the WF_RESOURCES table are the same as the initial default values in the Edit Service Parameters page.

In the Edit Service Parameters field, the service parameter names and values should be specified separated by colons, in the following format:

```
<name1>=<value1>:<name2>=<value2>: . . . <nameN>=<valueN>
```

The following service parameters can be specified for a container:

- SVC_WRITE_DIAG_TO_GSM_LOG - Specify Y if you want to write diagnostic information to the GSM log file in all cases. The default value is Y. Specify N if you want to let the FND: Debug Log Filename (AFLOG_FILENAME) profile option determine where to write the log, either to a specified file or to the database if no file is specified. For more information about FND: Debug Log profile options, refer to *Oracle E-Business Suite Setup Guide*.
- SVC_CONTAINER_LOOP_SLEEP - Specify the sleep time in seconds during which the container waits, after it finishes reading control messages from its GSM queue, before it checks that queue for messages again. The default sleep time is 10 seconds.
- SVC_CONTAINER_READ_TIMEOUT - Specify the maximum amount of time in seconds that the container continues to block on the GSM queue after processing the last message. If another message is received before this time expires, that message is processed and the timeout period begins again. If the timeout period expires and no more messages have been received, the container stops blocking on the queue and its sleep time begins. The default timeout period is 10 seconds.
- SVC_CONTAINER_LOG_LEVEL - Specify the level of detail to record for the container in its log. The default value is 5 (Error). The valid levels, in order from most detailed to least detailed, are:
 - 1 - Statement
 - 2 - Procedure
 - 3 - Event
 - 4 - Exception
 - 5 - Error

- 6 - Unexpected
- SVC_COMP_MONITOR_LOOP_SLEEP - Specify the sleep time in seconds during which the container waits, after it starts any automatic components that need to be started, before it checks its automatic components again. The default value is 60 seconds.
- SVC_COMP_MONITOR_ONDEMAND_FREQ - Specify the interval in seconds to determine how often the container checks whether its on-demand components need to be started or stopped. This activity is more costly than monitoring the automatic components and should usually be performed less frequently. The default value is 300 seconds.
- SVC_COMP_MAX_ERROR_COUNT - The container-level maximum error count. If any automatic or on-demand component in the container is stopped with an error the specified number of times, the component status will be set to System Deactivated, and the container will no longer automatically restart the component. The default value is 5.

You can also optionally specify the following service parameters for proxy settings. You should set these parameters if components in this container need to use a proxy server to access web content that is outside a firewall. For example, a mailer component may need to access outside web content that is to be included in an e-mail notification. The Generic Service Component Framework uses the values you set in these service parameters to set the relevant Java System Properties.

- SVC_PROXY_SET - Specify `true` to indicate that you want to use a proxy for your connections. The default value is `NONE`.
- SVC_PROXY_HOST - Specify the host machine for the proxy. The default value is `NONE`.
- SVC_PROXY_PORT - Specify the port on which the proxy is listening. The default value is `NONE`.
- SVC_NONPROXY_HOSTS - Specify any hosts that components in this container should access directly, rather than through the proxy server. When the service container starts, the Generic Service Component Framework uses the value of this parameter to set the `http.nonProxyHosts` System Property. Specify the list of hosts separated by vertical bars (`|`), without any spaces. You can use an asterisk (`*`) as a wildcard character within the host names. For example:
`*.us.example.com|*.example.org|*.example.net`

Note: If you use AutoConfig to specify proxy settings for your Oracle E-Business Suite instance, then you do not need to set the proxy-related service parameters here. In this case it is recommended that you

continue to use AutoConfig to manage your proxy settings.

Use the proxy-related service parameters only if you do not use a proxy setup elsewhere, but you do require it for service components such as workflow mailers or agent listeners.

Selecting the Log Level for a Container

You can use the Service Status page to control the running of a service component container, including changing the log level for the container. The log level controls how much information is recorded in the log. Note that the log level you select here applies only to the log messages for the container. You can assign separate log levels to the individual components within the container.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Service Components status icon > container link > (B) View Status

The log level with which the container starts is determined by the value of the SVC_CONTAINER_LOG_LEVEL service parameter. If no value is defined for that parameter, the log level is obtained from the default setting stored in the WF_RESOURCES table. The default container log level, which is also the recommended setting, is Error.

If the container is running, you can optionally specify a different container log level for the current session. To change the log level, select the level you want from the Change Log Level To pull-down menu and click the Go button. The log levels you can select, in order from most detailed to least detailed, are as follows:

- 1 - Statement
- 2 - Procedure
- 3 - Event
- 4 - Exception
- 5 - Error
- 6 - Unexpected

Note that the log level you set dynamically in the Service Status page applies only for the duration of the current container session, and does not change the log level stored for the container in the service parameters. To set the log level permanently, so that the container starts with that log level in each new session, edit the value of the SVC_CONTAINER_LOG_LEVEL service parameter in the Edit Service Parameters page. See: Editing Service Parameters for a Container, page 15-15.

If the log level has been changed dynamically for the current session, the Service Status page may not display the log level that is currently in effect for the container. However,

you can always review the current log level in the container log file by choosing View Log in the Service Components page or the Component Details page.

Creating Service Component Containers

If you create custom service components, you can choose to create custom containers to manage those service components. You create a container as a GSM service instance of type Generic Service Component Container in Oracle Applications Manager.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Service Components status icon > container link > (B) Create New

Among other properties, a GSM service instance can have work shifts assigned to it. A work shift in turn can have service parameters associated with it. For a service instance that is a service component container, these service parameters apply to the container as a whole to determine how the container manages the components that belong to it. If you create a custom container, you should specify service parameters for the work shifts for your new service instance in order to specify how to run the new container. To enter service parameters easily, copy the service parameters from one of the seeded Oracle Workflow containers to your new container.

After creating a customer container, you can assign service components to it using the appropriate service component configuration wizard. Ensure that your custom containers are running in order to run the service components belonging to them.

Notification Mailers

A notification mailer is a Java program that performs e-mail send and response processing for the Oracle Workflow Notification System, using the JavaMail API. You need to implement one or more notification mailers only if you want to have your workflow users receive their notifications by e-mail, as well as from the Worklist Web pages.

Managing Notification Mailers

The notification mailer program is defined as a service component type in the Generic Service Component Framework. This framework helps to simplify and automate the management of background Java services.

Oracle Workflow provides one seeded notification mailer service component, called Workflow Notification Mailer. Most of the configuration parameters for this mailer are set to default values. You can enter several of the remaining required parameters using AutoConfig. After installation, you then only need to enter the e-mail inbox password in order to complete the configuration of this mailer. Alternatively, if you only want to send outbound messages and do not need to receive inbound messages, you only need to disable inbound processing in order to complete the configuration of this mailer. If the mail servers and Business Event System components used by the notification mailers are set up, and the Workflow Mailer Service container to which the Workflow

Notification Mailer belongs is started, then the seeded notification mailer automatically starts running once its configuration is complete.

You cannot delete the seeded Workflow Notification Mailer or edit its name, assigned agents, correlation ID value, or container. However, if necessary you can optionally update other configuration parameters, schedule control events, or manually choose control commands to start, stop, suspend, resume, or refresh this notification mailer.

Note: Oracle Alert also uses the Workflow Notification Mailer to send and receive alert e-mail messages. If you use Oracle Alert, ensure that the configuration of the Workflow Notification Mailer meets your alert requirements. See: Setup Steps, *Oracle Alert User's Guide*.

You can also optionally create additional notification mailer service components. For example, you can create a notification mailer that processes only messages that belong to a particular workflow item type, or instances of a particular message from a particular item type. You can create additional mailers that process the same types of message to increase throughput.

The correlation ID for a notification mailer determines which messages it can process.

- To create a general notification mailer that can process any message from any item type that is not being handled by a dedicated notification mailer, leave the correlation ID blank. The seeded Workflow Notification Mailer has a blank correlation ID so that it can run as a general mailer.
- To dedicate a notification mailer to processing messages from a particular item type, set the correlation ID to the internal item type name followed by a colon and a percent sign.
- To dedicate a notification mailer to processing instances of a particular message from a particular item type, set the correlation ID to the internal item type name followed by a colon and then the internal message name.

Note: If you run a general notification mailer and a dedicated notification mailer at the same time, then the general notification mailer does not process messages that match the dedicated notification mailer's correlation ID, as long as the dedicated notification mailer has a status of `Running`, `Stopped With Error`, or `System Deactivated`. If the dedicated notification mailer has any other status, such as `User Deactivated` or `Suspended`, then the general notification mailer does process the messages that match the dedicated notification mailer's correlation ID.

To ensure consistency in message handling, all notification mailers that can process the same messages must share the same values for certain parameters. Multiple mailers can process the same messages in the following cases:

- A general mailer runs at the same time as any dedicated mailers.
- Multiple general mailers run at the same time.
- Multiple dedicated mailers for the same item type or message definition run at the same time.

In these cases, the notification mailers must share the same values for the following parameters:

- HTML Agent
- Attach Images to Outbound E-mails
- Attach Stylesheet to Outbound E-mail
- Autoclose FYI
- Direct Response
- Reset NLS
- Inline Attachments
- All message template parameters

However, these mailers can have different values for the From and Reply-to Address parameters. The headers of each notification e-mail message will contain the From and Reply-to Address values of the notification mailer that actually sent the message, unless the message itself has the special #WFM_FROM and #WFM_REPLYTO message attributes defined to override the notification mailer's parameters. See: Notification Mailer Attributes, *Oracle Workflow Developer's Guide*.

You can also configure any notification mailer service component to process only inbound messages, or only outbound messages. You associate inbound and outbound mailers with each other by assigning them the same mailer node name. The mailer node name indicates which inbound mailer can process incoming responses to outbound messages sent by a particular outbound mailer.

You can optionally assign the same node name to multiple mailers for load balancing purposes. However, each mailer that performs inbound processing for a node must have its own inbox.

- If you enable both outbound and inbound processing for the same mailer, that mailer will automatically use the same node name for both types of processing, enabling it to process incoming responses to the outbound messages it sent. You can optionally also create other notification mailers that share the same node name.
- If you create an outbound-only mailer, but you still want to perform response

processing for e-mail responses to the outbound messages it sends, you should create at least one other mailer with the same node name that does perform inbound message processing. Otherwise, there will be no inbound mailer that can process incoming responses to outbound messages sent by this outbound mailer.

- If you only want to implement outbound message processing, without inbound e-mail response processing, then you can configure an outbound-only mailer without creating a corresponding inbound mailer. In this case you should configure the mailer to use message templates for response-required notifications that do not request a response by e-mail, but instead direct recipients to respond from the Notification Details Web page. For example, you can configure the mailer to send response-required notifications using the Workflow View From UI message template, which is an alternative template provided by Oracle Workflow in the System: Mailer item type, or create your own custom message templates. The outbound-only mailer can still use the standard message templates to send outbound summary notifications or For Your Information (FYI) notifications that do not require a response.
- Create an inbound-only mailer only if you have also created at least one mailer with the same node name that performs outbound message processing. If no outbound mailer shares the same node name, no incoming response messages will be marked with that node name, and the inbound-only mailer will have no messages to process.

Dedicated mailers for different item types or message definitions should use different node names.

If you create custom notification mailer service components, you can either assign them to the seeded container for notification mailers, named Workflow Mailer Service, or, based on the volume to be handled by the seeded container, you can also choose to create your own custom containers.

Setting Up a Notification Mailer

Currently, Oracle Workflow supports the Simple Mail Transfer Protocol (SMTP) for outbound messages and the Internet Message Access Protocol (IMAP) for inbound messages. You must have an SMTP server set up in order to send Oracle Workflow notification e-mail messages, and an IMAP server set up if you want to receive e-mail notification responses. Users can receive e-mail notifications using various e-mail clients, although notifications may be displayed differently in different clients, depending on the features each client supports.

Note: Oracle Workflow supports IMAP version 4 (IMAP4) compliant mail servers. Ensure that your mail server uses this IMAP version.

To set up a notification mailer, you must perform the following steps.

1. Set up an SMTP mail server to send outbound messages.

You can optionally configure the SMTP server to require authentication for server connections through the Simple Authentication and Security Layer (SASL). The Oracle Workflow notification mailer supports the PLAIN, LOGIN, and DIGEST-MD5 authentication mechanisms. Additionally, if you have applied patch 9452181 for JavaMail version 1.4.x, then the notification mailer can also support the Microsoft NTLM authentication mechanism. If you configure your SMTP server to use one of these authentication mechanisms, set up a user name and password for the notification mailer to use in establishing an authenticated connection to the server.

If you configure your SMTP server to support more than one authentication mechanism, then the notification mailer uses the mechanism that appears first in the server's list of supported mechanisms. Consequently, if you want the notification mailer to use a particular mechanism, ensure that that mechanism appears first in the server's list. At a minimum, you should ensure that the first authentication mechanism listed for the server is one that the notification mailer supports.

Note: If you use the PLAIN or LOGIN authentication mechanisms, it is recommended to connect to the SMTP server through Secure Sockets Layer (SSL) to encrypt the user name and password that are sent to the server. See: *Connecting to Mail Servers Through SSL, Oracle Workflow Administrator's Guide*. If you use the DIGEST-MD5 or NTLM authentication mechanisms, the JavaMail API encrypts the user name and password before sending the data to the SMTP sever.

2. Set up an IMAP4 compliant mail server with an e-mail account for the notification mailer if you want to receive inbound messages.

The notification mailer requires three folders in this e-mail account: the inbox, a folder to store processed messages, and a folder to store discarded messages. If the e-mail account does not already include folders named PROCESS and DISCARD, Oracle Workflow automatically creates these two folders when you complete the basic notification mailer configuration. You can optionally specify other folders for the notification mailer using the advanced configuration wizard.

Note: If you create the PROCESS and DISCARD folders manually before configuring the notification mailer, use your e-mail client to create these folders. A notification mailer may not be able to access folders that were created using command line tools outside the e-mail client.

However, note that you must not use an e-mail client to access the

notification mailer's e-mail account while the notification mailer is running. Use the e-mail client only during setup.

3. You can use AutoConfig to enter the following configuration parameters for the seeded Workflow Notification Mailer service component during installation. For more information about running AutoConfig, see: Technical Configuration, *Oracle E-Business Suite Setup Guide*, and Technical Configuration Tools, *Oracle E-Business Suite Concepts*.
 - SMTP Server
 - IMAP Server (if you want to receive inbound messages)
 - Inbox Username (if you want to receive inbound messages)
 - Reply To E-mail Address (if you want to receive inbound messages)
 - HTML Agent Name - This parameter defaults to the value you enter for the Applications Servlet Agent parameter in AutoConfig. Use the following format:
`http://<server_name:port>/OA_HTML/`

Note: When you enter the SMTP Server and IMAP Server parameters, specify each server in the following format:
`<server_name>[:<port_number>]`

- For the IMAP Server parameter, specify the actual host name. Do not use `localhost` as the setting for this parameter.
- For the SMTP Server parameter, Oracle strongly recommends that you specify the actual host name. However, you can specify `localhost` as the setting for the SMTP Server parameter if you ensure that an SMTP server is configured to send e-mails to all valid domains on each host where concurrent managers run. If you have implemented Parallel Concurrent Processing to allow concurrent processing activities to be distributed across multiple nodes in a cluster system, then you must configure an SMTP server on every node. Otherwise, if a concurrent manager attempts to execute outbound notification mailer processing on a node without an SMTP server, the processing will fail.
- You can optionally specify the port number to use on each server. If you do not specify a port number, the notification mailer uses port 143 on the IMAP server and port 25 on the

SMTP server by default.

4. Ensure that the Business Event Local System status is set to Enabled in the Workflow Configuration page, and that the JOB_QUEUE_PROCESSES database initialization parameter, which is required for the Business Event System, is set to an appropriate value. The Business Event Local System status is set to Enabled by default, and usually you do not need to change this status. If notification processing is not being completed, however, you should check this preference value.
5. **(Recommended)** You can optionally set the WF: Workflow Mailer Framework Web Agent profile option to the host and port of the Web server that notification mailers should use to generate the content for Oracle Application Framework regions that are embedded in notifications. If this profile option is not set, notification mailers will use the same Web agent specified in the Application Framework Agent profile option. However, on a load-balanced Web server, notification mailers might not be able to render Oracle Application Framework content within a notification. In this case, set the WF: Workflow Mailer Framework Web Agent profile option to a physical host, instead of a virtual host. The WF: Workflow Mailer Framework Web Agent profile option should be set at site level. See: *Overview of Setting User Profiles, Oracle E-Business Suite Setup Guide*.
6. Before a service component can run, the container which manages it must first be started. The seeded Workflow Notification Mailer service component belongs to a container named Workflow Mailer Service, while the seeded agent listener service components that are also required for notification mailer processing belong to a container named Workflow Agent Listener Service. You should ensure that these two containers are running. If you create your own custom containers for custom service components, ensure that those containers are running as well. Use the Service Instances page to start the containers as service instances in Generic Service Management (GSM).
7. When the Workflow Agent Listener Service container is running, it automatically starts seeded agent listener service components named Workflow Deferred Notification Agent Listener, Workflow Error Agent Listener, and Workflow Inbound Notifications Agent Listener, which are required for notification mailer processing. Ensure that these agent listeners are running.
8. Use the notification mailer configuration wizard to configure your notification mailer service component. The Basic Configuration page lets you configure a notification mailer quickly by entering only the minimum required parameters, while the advanced configuration wizard lets you specify additional parameters to control how the notification mailer processes messages.

If you entered configuration parameters for the seeded Workflow Notification Mailer through AutoConfig, you only need to enter the password for the e-mail

inbox in order to complete the configuration for that mailer and begin running it. If you did not enter parameters for the seeded mailer through AutoConfig, then in order to complete the configuration for that mailer you need to enter only the SMTP server, IMAP server, e-mail inbox username, e-mail inbox password, and reply-to e-mail address. All other configuration parameters for the seeded Workflow Notification Mailer are initially set to default values and do not need to be changed, although you can optionally do so if you choose.

Note: The IMAP server, e-mail inbox username, e-mail inbox password, and reply-to e-mail address are required only if you want to receive inbound messages. Alternatively, if you only want to send outbound messages and do not need to receive inbound messages, you only need to disable inbound processing in order to complete the configuration of the Workflow Notification Mailer.

9. **(Optional)** By default, the seeded Workflow Notification Mailer has a Launch Summary Notifications event scheduled to send summary notifications once a day. You can optionally use the notification mailer configuration wizard to modify the start time and interval for this event's schedule, or to schedule the Launch Summary Notifications event at the interval you choose for any notification mailer service component. When this event is processed, a summary notification is sent to each role with a notification preference of SUMMARY or SUMHTML, listing all the notifications that are currently open for that role.
10. **(Optional)** You can configure a notification mailer to connect to the SMTP server and IMAP server through Secure Sockets Layer (SSL) to encrypt the data exchanged. See: *Connecting to Mail Servers Through SSL, Oracle Workflow Administrator's Guide*.
11. **(Optional)** You can optionally set the internal mailer parameter named `HTML_DELIMITER` to specify which characters the notification mailer uses to delimit response values in response templates for HTML-formatted e-mail notifications. Valid values for the `HTML_DELIMITER` parameter are:
 - `DEFAULT` - The notification mailer uses the default delimiters, currently set as the single quote (') for both the opening and the closing delimiter. The notification mailer also uses the default delimiters if the `HTML_DELIMITER` parameter value is left null.
 - `APOS` - The notification mailer uses the single quote, or apostrophe (') , as both the opening and the closing delimiter. This setting is currently the same as the default.
 - `QUOTE` - The notification mailer uses the double quote (") as both the opening and the closing delimiter.

- BRACKET - The notification mailer uses the left bracket ([) as the opening delimiter and the right bracket (]) as the closing delimiter.

Using single quotes as the delimiters accommodates e-mail applications that cannot process double quotes in the tag for the response template link, but can accept single quotes. However, if you want users to be able to use apostrophes or single quotes in their response values without entering an escape character, you can use double quotes or brackets as the delimiters, depending on what your e-mail application supports. See: To Respond to an HTML E-mail Notification, *Oracle Workflow User's Guide*.

Note: If the HTML_DELIMITER parameter is set to an invalid value, the notification mailer throws an exception at startup. Any notifications created during this time are rendered with the default delimiters instead.

By default, the HTML_DELIMITER parameter is set to the value DEFAULT. Use the `afsvcpup.sql` script to change the parameter value to specify the delimiters you want to use. See: To Set Internal Mailer Parameters, *Oracle Workflow Administrator's Guide*.

If a particular notification message has the special #WFM_HTML_DELIMITER message attribute defined, however, the notification mailer will use the #WFM_HTML_DELIMITER attribute value to determine which delimiters to use for that notification, instead of using the HTML_DELIMITER parameter value.

Note: The HTML_DELIMITER parameter only controls the response templates for HTML-formatted notifications. This parameter does not apply to plain text notifications.

12. **(Optional)** The seeded Workflow Notification Mailer uses the Automatic startup mode by default and will be started automatically when you complete its configuration. If you select the Manual startup mode for a notification mailer service component, use the Service Components page to start that notification mailer. You can also use this page to manage any notification mailer service component.

Outbound Notification Mailer Processing

When the Workflow Engine determines that a notification message must be sent, it raises an event in the Business Event System called `oracle.apps.wf.notification.send`. Oracle Workflow provides a seeded subscription to this event, which is defined to be deferred immediately so that the workflow process that owns the notification can continue. The event is placed on the standard WF_DEFERRED agent. Oracle Workflow provides a seeded agent listener named Workflow Deferred Notification Agent Listener

that runs on this agent to continue notification processing. This agent listener is dedicated solely to processing deferred notification events.

When the event is dequeued from WF_DEFERRED and the subscription is processed, the subscription requires the event data for the event, causing the generate function for the event to be executed. The generate function for this event performs the following actions:

- Resolves the notification recipient role to a single e-mail address, which itself can be a mail list.
- Checks the notification preference of the recipient to determine whether an e-mail notification is required, and in what type of format.
- Switches its database session to the recipient role's preferred language and territory as defined in the directory service.
- Generates an XML representation of the notification message and any optional attachments using the appropriate message template.

Finally, the subscription places the event message on the standard WF_NOTIFICATION_OUT agent.

A notification mailer service component polls the WF_NOTIFICATION_OUT agent for messages that must be sent by e-mail. When the notification mailer dequeues a message from this agent, it uses a Java-based notification formatter to convert the XML representation of the notification into a MIME (Multipurpose Internet Mail Extensions) encoded message and sends the message by the Simple Mail Transfer Protocol (SMTP).

The e-mail notifications are based on message templates defined in Oracle Workflow Builder. Oracle Workflow provides a set of standard templates in the System: Mailer item type, which are used by default. It is not recommended to modify the standard templates. However, you can customize the message templates used to send your e-mail notifications by creating your own custom message templates in a custom item type using the Workflow Builder. Then assign these templates to a particular notification in a workflow process by defining special message attributes. In this case the templates assigned to the notification override any other templates.

You can also create your own custom message templates in the System: Mailer item type using the Workflow Builder, and assign these templates to a particular notification mailer service component in the mailer configuration parameters. The templates assigned to a mailer override the default System: Mailer templates. However, if any notifications have templates specifically assigned to them through message attributes, the notification-level templates still override the templates assigned to the mailer.

If the notification mailer cannot deliver an e-mail notification to the recipient's e-mail address, it performs the following actions:

- Sets the mail status of the notification to FAILED. This mail status indicates that an exception prevented this e-mail notification from being delivered but does not

prevent the mailer from processing other notifications.

- Adds the e-mail address to its invalid e-mail address list. To avoid unnecessary processing, each notification mailer stores a list of e-mail addresses to which it could not deliver messages, and does not attempt to send any further messages to those addresses. Instead, for any subsequent notifications to the listed addresses, the notification mailer simply sets the mail status directly to `FAILED`.

Note: Each notification mailer can store up to 100 e-mail addresses in its invalid e-mail address list. If the notification mailer encounters additional invalid addresses when the list is already full, the notification mailer removes the oldest addresses from the list and adds the new addresses in their place. Also, the notification mailer clears the list by removing all addresses whenever you stop and restart the mailer.

- Changes the notification preference of the recipient to `DISABLED`. To further help avoid unnecessary processing, if a recipient has a notification preference of `DISABLED`, Oracle Workflow does not generate a complete XML representation of any notifications to that recipient, and a notification mailer does not attempt to send e-mail notifications to that recipient. Instead, the notification mailer simply sets the mail status of the notifications directly to `FAILED`. The change in notification preference also indicates to the user that e-mail notifications cannot be delivered. You or the user must correct the issue that caused the failure and then reset the notification preference in order for the user to receive e-mail notifications.
- Sends a notification to the `SYSADMIN` user with the information that an e-mail notification could not be sent to one or more recipients, that the notification preference for those recipients has been set to `DISABLED`, and that those recipients' original notification preferences, which are listed, should be reset after the issues that caused the failures are corrected. See: User Notification Preference Update Report Message, *Oracle Workflow Administrator's Guide*.

Individual users can reset their notification preference manually using the Preferences page in Oracle E-Business Suite. You can also run the Workflow Directory Services Bulk Reset `DISABLED` Notification Preference concurrent program to reset the notification preference for multiple users at once. See: Handling Mailer Errors, *Oracle Workflow Administrator's Guide*.

After correcting the e-mail issues and resetting `DISABLED` notification preferences, you can run the Resend Failed/Error Workflow Notifications concurrent program to retry open notifications that previously could not be sent. See: Handling Mailer Errors, *Oracle Workflow Administrator's Guide*.

Inbound Notification Mailer Processing

Notification mailers can also process e-mail responses from users, using the Internet Message Access Protocol (IMAP). A notification mailer uses a Java-based e-mail parser to interpret the text of each message and create an XML representation of it.

A notification mailer uses three folders in your response mail account for response processing: one to receive incoming messages, one to store processed messages, and one to store discarded messages.

A notification mailer does the following to process response messages:

- Logs into its IMAP e-mail account.
- Checks the inbox folder for messages. If a message exists, the notification mailer reads the message, checking for the notification ID (NID) and node identifier in the NID line.
- If the message is not a notification response, meaning it does not contain an NID line, the notification mailer moves the message to the discard folder and treats it as an unsolicited message. For the first unsolicited message from a particular e-mail address, the notification mailer also sends a warning message back to the sender of the message. However, to avoid sending unnecessary warnings due to bounced or auto-reply messages, each mailer node stores a list of e-mail addresses from which it has received unsolicited mail, and does not send any further warning messages to those addresses. Instead, if the node receives a second unsolicited message from a particular address, the notification mailer discards the message and raises the `oracle.apps.wf.mailer.unsolicited` event. You can optionally define a subscription to this event if you want to perform some other action in response to the second unsolicited message. For all subsequent unsolicited messages, the notification mailer simply discards the message.

Note: Each mailer node can store up to 100 e-mail addresses in its warned list. If the node receives unsolicited messages from additional addresses when the list is already full, the notification mailer removes the oldest addresses from the list and adds the new addresses in their place. Also, the notification mailer clears the list by removing all addresses when you start the mailer for the first time, and again whenever you stop and restart its container. In these cases, the mailer may send another warning message if it receives further unsolicited e-mail from an address that is no longer on the warned list.

Note: You can optionally use the `Send Warning for Unsolicited E-mail` mailer parameter to prevent notification mailers from

sending any warning messages at all. See: Notification Mailer Configuration Wizard, page 15-32.

- If the message is a notification response, but for a different node, the notification mailer leaves the message in the inbox and adds the e-mail's Unique Message ID (UID) to its ignore list.
- If the message is a notification response for the current node, meaning it contains an NID line including the node identifier of the current node, the notification mailer processes the message.

The notification mailer performs the following steps for messages that belong to its node.

- Retrieves the notification ID.
- Checks to see if the message bounced by referring to the tags specified in the configuration parameters, if any. If the message bounced, the notification mailer updates the notification's status and stops any further processing, based on the specifications of the tag list.
- Checks the Oracle Workflow database for this notification based on the NID line.
 - If the notification does not exist, meaning the notification ID or the access key in the NID line is invalid, the notification mailer moves the message to the discard folder. If the NID line is incorrectly formatted, the notification mailer moves the message to the discard folder and treats it as an unsolicited message.
 - If the notification exists, but is closed or canceled, the notification mailer moves the message to the processed folder and sends a Workflow Closed Mail or Workflow Canceled Mail message to the recipient role, respectively.

Note: You can optionally use the Send E-mails for Canceled Notifications mailer parameter to prevent notification mailers from sending any notification cancellation messages. See: Notification Mailer Configuration Wizard, page 15-32.

- If the inbound message is a response to a request for more information that has already been answered, or if the message is formatted as a more information response but no information was requested for that notification, then the notification mailer moves the message to the discard folder and sends a Workflow More Info Answered Mail message to the sender of the message.
- If the notification exists and is open, the notification mailer generates an XML representation of the message and places it on the standard

WF_NOTIFICATION_IN agent as an event called oracle.apps.wf.notification.receive.message. The notification mailer then moves the message for the completed notification to the processed folder.

Note: If the character encoding of the response message is not compatible with the database codeset, the notification mailer may not be able to parse the response and recognize the response values. Ensure that the character encoding of messages in your mail client is compatible with the codeset of your database.

Finally, if there are no more unprocessed messages in the inbox, the notification mailer logs out of the e-mail account.

Oracle Workflow provides a seeded agent listener named Workflow Inbound Notifications Agent Listener that runs on the WF_NOTIFICATION_IN agent to continue notification processing for the valid response messages placed on that agent. When an event message is dequeued from WF_NOTIFICATION_IN, Oracle Workflow executes a seeded subscription that calls the appropriate notification response function. This function verifies the response values with the definition of the notification message's response attributes in the database. If a response value is invalid, or if no response value is included, the notification mailer sends a Workflow Invalid Mail message to the recipient role, or, for an invalid response to a request for more information, the notification mailer sends a Workflow Invalid Open Mail (More Information Request) message to the recipient role. If the responses are valid, the notification response function records the response and completes the notification.

Notification Mailer Configuration Wizard

Use the notification mailer configuration wizard to configure a new notification mailer service component, or to edit the configuration of an existing notification mailer service component. The notification mailer configuration wizard begins with the Basic Configuration page, which lets you configure a notification mailer quickly by entering only the minimum required parameters.

From the Basic Configuration page, you can also navigate to the advanced configuration wizard to specify additional parameters that control how the notification mailer processes messages. The advanced configuration wizard lets you define general and detail attributes, define e-mail server and message generation parameters, schedule control events, and define tags to assign statuses to unusual messages.

Some parameters appear in both the Basic Configuration page and the advanced configuration wizard. Both the Basic Configuration page and the advanced configuration wizard also let you send test messages.

Note: If you are configuring the seeded Workflow Notification Mailer and you entered configuration parameters for this mailer through

AutoConfig, then you only need to enter the password for the e-mail inbox in order to complete the configuration for that mailer. If you did not enter parameters for the seeded mailer through AutoConfig, then in order to complete the configuration for that mailer you need to enter only the SMTP server, IMAP server, e-mail inbox username, e-mail inbox password, and reply-to e-mail address. All other configuration parameters for the seeded Workflow Notification Mailer are initially set to default values and do not need to be changed, although you can optionally do so if you choose.

Note that the IMAP server, e-mail inbox username, e-mail inbox password, and reply-to e-mail address are required only if you want to receive inbound messages. Alternatively, if you only want to send outbound messages and do not need to receive inbound messages, you only need to disable inbound processing in order to complete the configuration of the Workflow Notification Mailer.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Notification Mailers status icon > (B) Create > (B) Continue

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Notification Mailers status icon > (B) Edit

Basic Configuration

This page lets you configure a notification mailer quickly by entering only the minimum required parameters in a single page. You must set parameters marked with an asterisk (*) to appropriate values for your environment before you can run the notification mailer.

Details

- **Name** - The name of the service component. This name must be unique. The name of the seeded notification mailer service component is `Workflow Notification Mailer`, and you cannot change this value.

Outbound E-mail Account (SMTP)

- **Server Name** - The name of the outbound SMTP mail server. Oracle strongly recommends that you specify the actual host name for the SMTP server. However, you can specify `localhost` as the setting for this parameter if you ensure that an SMTP server is configured to send e-mails to all valid domains on each host where concurrent managers run. If you have implemented Parallel Concurrent Processing to allow concurrent processing activities to be distributed across multiple nodes in a cluster system, then you must configure an SMTP server on every node. Otherwise, if a concurrent manager attempts to execute outbound notification mailer processing on a node without an SMTP server, the processing will fail. Also, when

you save the configuration, Oracle Workflow Manager tests the connection to the SMTP server from within the Web tier host. Consequently, if you set the outbound server name to `localhost`, you should ensure that an SMTP server is configured on the Web tier host as well.

You can optionally specify the port number to use on that server. If you do not specify a port number, the notification mailer uses port 25 by default. Specify the server in the following format: `<server_name>[:<port_number>]`

For example: `mysmtpserver.mycompany.com:25`

- **Username** - If the outbound SMTP server is configured to require authentication, enter the user name of the account that the notification mailer uses to connect to the SMTP server.
- **Password** - If the outbound SMTP server is configured to require authentication, enter the password for the account specified in the Username parameter. The password value is masked as asterisks in the display and is stored in encrypted form.
- **Outbound SSL Enabled** - Select this parameter to enable the notification mailer to use Secure Sockets Layer (SSL) for connections to the SMTP server. Deselect this parameter to use non-SSL connections.

Note: When you enable SSL, the notification mailer connects to the SMTP server through port 465 by default. You can optionally specify a different port number along with the SMTP server name in the Outbound E-mail Account (SMTP): Server Name parameter.

Before you can use SSL, you must also complete additional setup steps. See: *Connecting to Mail Servers Through SSL, Oracle Workflow Administrator's Guide.*

Inbound E-mail Account (IMAP)

- **Inbound Processing** - Select this parameter to enable inbound e-mail processing with this notification mailer. Deselect this parameter to disable inbound e-mail processing for this notification mailer and dedicate the notification mailer solely to outbound processing.

If you disable inbound processing, you can leave the other inbound parameters blank.

- **Server Name** - The name of the inbound IMAP mail server. Note that you must specify the actual host name for the server. Do not use `localhost` as the setting for this parameter. You can optionally specify the port number to use on that server. If you do not specify a port number, the notification mailer uses port 143 by default. Specify the server in the following format: `<server_name>[:<port_number>]`

For example: `myimapserver.mycompany.com:143`

- **Username** - The user name of the mail account that the notification mailer uses to receive e-mail messages.
- **Password** - The password for the mail account specified in the Username parameter. The password value is masked as asterisks in the display and is stored in encrypted form.
- **Reply-To Address** - The address of the e-mail account that receives incoming messages, to which notification responses should be sent. This value must be a full RFC822-compliant e-mail address.

If a particular notification message has the special `#WFM_REPLYTO` message attribute defined, however, the notification mailer will use the `#WFM_REPLYTO` attribute value as the reply address for that message, instead of the Reply-To Address parameter value.

Note: If you enable inbound processing, Oracle Workflow by default sets the From parameter, which is displayed in the From field of the message headers, to the name portion of the reply-to address. For example, if the reply-to address is `Workflow@mycompany.com`, the notification mailer sets the From parameter to `Workflow`.

If you disable inbound processing, Oracle Workflow by default sets both the Reply-To Address parameter and the From parameter to `nobody@<server_name>`, where `<server_name>` is the name of the outbound SMTP mail server.

To specify a different From value, navigate to the advanced configuration wizard.

- **Inbound SSL Enabled** - Select this parameter to enable the notification mailer to use SSL for connections to the IMAP server. Deselect this parameter to use non-SSL connections.

Note: When you enable SSL, the notification mailer connects to the IMAP server through port 993 by default. You can optionally specify a different port number along with the IMAP server name in the Inbound E-mail Account (IMAP): Server Name parameter.

Before you can use SSL, you must also complete additional setup steps. See: *Connecting to Mail Servers Through SSL, Oracle Workflow Administrator's Guide*.

Note: The notification mailer requires three folders in the IMAP mail

account: the inbox, a folder to store processed messages, and a folder to store discarded messages. If you enable inbound processing and the mail account you specify in the Username parameter does not already include folders named PROCESS and DISCARD, Oracle Workflow automatically creates these two folders. To specify other folders for the notification mailer, navigate to the advanced configuration wizard.

Note: If you enable inbound processing, the notification mailer uses the Workflow Open Mail (Templated) message, which provides a response template for sending responses by e-mail, as the default message template for e-mail notifications that require a response. If you disable inbound processing, the notification mailer uses the Workflow Open Mail (Outlook Express) message, which provides a link in HTML notifications for entering responses in the Notification Details page, as the default message template for e-mail notifications that require a response. To specify other message templates, navigate to the advanced configuration wizard.

Note that the plain text version of the Workflow Open Mail (Outlook Express) message requests a response by e-mail. If you disable inbound processing, ensure that your users do not have a notification preference of MAILTEXT or MAILATTH. Alternatively, if you disable inbound processing and you want users to receive plain text notifications, use the advanced configuration wizard to specify a message template that directs recipients to respond from the Notification Details Web page, such as the standard Workflow View From UI message template or a custom message template.

To cancel any changes on this page, click the Cancel button.

To save this configuration, click the Apply button.

To send test messages, click the Test Mailer button. In the Test Notification Mailer page, select the recipient role to which the messages should be sent, and click the Send Test Message button.

Note: To send a test message successfully, you must select a recipient role that either has a valid e-mail address defined, or that has members with valid e-mail addresses defined. The recipient role must also have a notification preference that includes individual e-mail notifications.

If you set an override e-mail address for the notification mailer, the Test Notification Mailer page displays that address. In this case the test message is sent to the override address rather than the e-mail address of the recipient role. However, you must still select a recipient role to enable the notification mailer to send the test messages. See: Reviewing

Oracle Workflow sends two test messages to the recipient role: one message with content built using PL/SQL and one message with Oracle Application Framework content. Check the e-mail account for the recipient role to view the test messages and reply to them with the Acknowledge response. If you did not implement inbound e-mail processing for this mailer, use the Worklist pages to respond to the test messages after viewing the outbound messages in e-mail. After you acknowledge both test messages, Oracle Workflow sends a confirmation message to the same recipient role to show that the notification mailer successfully processed the inbound response e-mails.

If you do not receive the test messages or the response confirmation message, or if the message content does not appear correctly, check the notification mailer setup, including the mail servers and the mailer configuration parameters. In particular, if the Oracle Application Framework content does not appear correctly, check the Application Framework Agent and WF: Workflow Mailer Framework Web Agent profile options, as well as the Framework User, Framework Responsibility, Framework Application ID, and Framework URL Timeout parameters in the advanced configuration wizard. See: Setting Up a Notification Mailer, page 15-22 and Message Generation, page 15-48.

Note: Oracle Workflow sends the test messages by launching the PLSQL/OAFwk Response Test Process in the System: Tests (WFTESTS) item type. This item type is stored in a file called wftstmlr.wft in the \$FND_TOP/import/<lang> subdirectory. You can optionally use the Status Monitor to check the status of the test process.

To set additional parameters for this notification mailer in the advanced configuration wizard, click the Advanced button.

Define

This page lets you define general attributes for the service component. Some attributes are already set to required values and cannot be modified. You must set attributes marked with an asterisk (*) to appropriate values for your environment before you can run the service component.

- **ID** - The configuration wizard displays the identifier for the service component.
- **Status** - The configuration wizard displays the status of the service component.
- **Name** - The name of the service component. This name must be unique. You can only edit the name when the notification mailer is not running. The name of the seeded notification mailer service component is `Workflow Notification Mailer`, and you cannot change this value.
- **Startup Mode** - Select Automatic, Manual, or On-Demand as the startup mode for

the service component. You can only edit the startup mode when the notification mailer is not running. The seeded Workflow Notification Mailer is assigned the Automatic startup mode by default, but you can optionally change this value.

- **Container Type** - The container type to which this service component belongs, which is always Oracle Applications Generic Service Management (Oracle Applications GSM).
- **Inbound Agent** - The Business Event System agent for inbound processing. The inbound agent for a notification mailer service component is always WF_NOTIFICATION_IN.
- **Outbound Agent** - The Business Event System agent for outbound processing. The outbound agent for a notification mailer service component is always WF_NOTIFICATION_OUT.
- **Correlation ID** - Enter a correlation ID value to determine which messages this notification mailer can process.
 - To create a general notification mailer that can process any message from any item type that is not being handled by a dedicated notification mailer, leave the correlation ID blank. The seeded Workflow Notification Mailer has a blank correlation ID so that it can run as a general mailer to process all messages; you cannot change this setting.

- To dedicate a notification mailer to processing messages from a particular item type, set the correlation ID to the internal item type name followed by a colon and a percent sign, in the following format:

```
<item_type_name>:%
```

For example:

```
WFDEMO:%
```

- To dedicate a notification mailer to processing instances of a particular message from a particular item type, set the correlation ID to the internal item type name followed by a colon and then the internal message name, in the following format:

```
<item_type_name>:<message_name>
```

For example:

```
WFDEMO:APPROVE_REQUISITION
```

By dedicating a notification mailer to a particular item type or message definition, you can increase throughput for the associated messages.

Both dedicated and general notification mailer components are compatible with each other. You can run several dedicated and general notification mailers at the same time if you choose. In this case, note that a general notification mailer does not

process messages that match a dedicated notification mailer's correlation ID, as long as the dedicated notification mailer has a status of `Running`, `Stopped With Error`, or `System Deactivated`. If the dedicated notification mailer has any other status, such as `User Deactivated` or `Suspended`, then the general notification mailer does process the messages that match the dedicated notification mailer's correlation ID.

To cancel any changes on this page, click the `Cancel` button.

To save these settings and proceed to the next step of the configuration wizard, click the `Next` button.

Details

This page lets you define detail attributes for the service component. You must set attributes marked with an asterisk (*) to appropriate values for your environment before you can run the service component. A refresh icon identifies attributes that can be refreshed dynamically while the service component is running.

- **ID** - The configuration wizard displays the identifier for the service component.
- **Status** - The configuration wizard displays the status of the service component.
- **Name** - The configuration wizard displays the name defined for the service component.
- **Container** - The container to which the service component will belong. Oracle Workflow provides a container called `Workflow Mailer Service` for notification mailer service components.
- **Maximum Idle Time** - If you selected the `On-Demand` startup mode for the service component, enter the maximum time in minutes that the service component can remain idle before it is stopped. An on-demand component that is stopped in this way will be restarted by its container when it is needed again to process new messages.
- **Max Error Count** - The number of consecutive errors the service component can encounter before its container stops it and changes its status to `Stopped with Error`. If an error is resolved and processing can continue, the error count is reset. The default value for the maximum error count is 10.
- **Inbound Thread Count** - Set the inbound processing thread count to 1 (one) to enable inbound message processing with this notification mailer. Select 0 (zero) to disable inbound message processing for this notification mailer and dedicate the notification mailer solely to outbound processing. If you selected the `Inbound Processing` parameter in the `Basic Configuration` page, the inbound thread count is set to 1; if you deselected the `Inbound Processing` parameter, the inbound thread count is set to 0.

The inbound thread count cannot be greater than 1, because only one thread can access the e-mail inbox at a time. If you disable inbound message processing for this notification mailer, but you still want to perform e-mail response processing, you should create at least one other notification mailer with the same node name that does perform inbound message processing. Otherwise, there will be no inbound mailer that can process incoming responses to outbound messages sent by this outbound mailer.

- **Outbound Thread Count** - Specify the number of outbound processing threads you want to execute simultaneously with this notification mailer. You can set the outbound thread count to 1 (one) or more depending on the volume of outbound messages you need to send. Specify 0 (zero) to disable outbound message processing for this notification mailer and dedicate the notification mailer solely to inbound processing. If you disable outbound message processing for this notification mailer, you should create at least one outbound notification mailer with the same node name. Otherwise, no inbound response messages will be marked with that node name and this inbound mailer will have no messages to process. The default value for the outbound thread count is 1.
- **Log Level** - Select the level of detail for the information you want to record in the service component container log. The recommended log level, which is also the default value, is Error. Usually the log level only needs to be changed if you want to record additional detailed information for debugging purposes. You can choose the following levels:
 - 1 - Statement
 - 2 - Procedure
 - 3 - Event
 - 4 - Exception
 - 5 - Error
 - 6 - Unexpected
- **Processor Read Wait Timeout** - Specify the amount of time in seconds that the service component's processing thread continues to wait, after reading the last message from its assigned queue, before timing out. If another message is received before this time expires, that message is processed and the timeout period begins again. If the timeout period expires and no more messages have been received, the service component stops reading and its sleep time begins. The default read timeout period for a notification mailer is 10 seconds.
- **Processor Min Loop Sleep** - Specify the minimum sleep time in seconds during which the service component waits, after its read timeout period expires, before it

checks its queue for messages again. The default minimum sleep time for a notification mailer is 5 seconds.

- **Processor Max Loop Sleep** - Specify the maximum sleep time in seconds if you want to increase the sleep time between read attempts when no messages are received. If you specify a maximum sleep time that is greater than the minimum sleep time, then the service component initially waits for the minimum sleep time after it finishes reading messages from its queue. If no messages are read in subsequent attempts, then the sleep time between read attempts gradually increases until the maximum sleep time is reached. Increasing the sleep time can help enhance performance if messages are received infrequently. You can also specify 0 (zero) for this parameter to indicate that the sleep time should not be increased. In this case, the service component always waits for the minimum sleep time between read attempts. The default maximum sleep time for a notification mailer is 60 seconds.
- **Processor Error Loop Sleep** - Specify the sleep time in seconds during which the service component waits, after an error occurs, before it attempts to begin processing again. The default error sleep time for a notification mailer is 60 seconds.
- **Processor Close on Read Timeout** - Select this parameter to specify that the service component should close its connections after its read timeout period expires, when its sleep time begins. Deselect this parameter to specify that the connections should remain open until the processing thread stops.

Selecting this parameter lets the notification mailer close its session with the IMAP server or SMTP server if it could not read a message from the IMAP inbox or from the database, respectively, before the read timeout period ended. For example, if an external process is accessing the IMAP inbox, the notification mailer may not be able to read or access the inbox for some time. In this case it may be advantageous for the notification mailer to close the existing connection, wait for a while, and then try to re-establish a new connection. Additionally, some IMAP servers may cause an idle session to time out and become invalid. In this case also, it is advantageous for the notification mailer to close the existing connection and re-establish a new one.

To cancel any changes on this page, click the Cancel button.

To return to the previous step of the configuration wizard, click the Back button.

To save these settings and proceed to the next step of the configuration wizard, click the Next button.

E-mail Servers

This page lets you define e-mail server parameters for the notification mailer. Some parameters are already set to required values and cannot be modified. You must set parameters marked with an asterisk (*) to appropriate values for your environment before you can run the notification mailer. A refresh icon identifies attributes that can be

refreshed dynamically while the service component is running. If the notification mailer is currently running, then parameters marked with a refresh icon will be refreshed immediately when you select the Next button.

General

- **Mailer Node Name** - The node identifier name used by this notification mailer. The maximum length for a node name is eight characters. The node name cannot include any spaces or any of the following characters: left bracket ([), right bracket (]), slash (/), or at sign (@). The node name is included with the outgoing notification ID in outbound messages, and is used in inbound messages to identify the notification mailer that should process the messages. If you use the inbound and outbound thread count parameters to create notification mailers that are dedicated to either inbound or outbound processing, you should ensure that you assign the same node name to at least one outbound mailer and one inbound mailer, so that inbound mailer can process responses to messages sent by the outbound mailer. You can optionally assign the same node name to multiple mailers for load balancing purposes. However, each mailer that performs inbound processing for a node must have its own inbox. The default node name is WFMAIL.

Note: The node name for each node must be unique. However, multiple mailers can share the same node.

If a particular notification message has the special #WFM_NODENAME message attribute defined, however, an outbound notification mailer will include the #WFM_NODENAME attribute value when sending the message, instead of the Mailer Node Name mailer parameter value.

- **Email Parser** - The Java class used to parse an incoming notification response e-mail formatted according to the templated response method and to create an XML document for the response. The notification mailer uses this parser when the Direct Response parameter is deselected. The default standard e-mail parser provided by Oracle Workflow is named oracle.apps.fnd.wf.mailer.TemplatedEmailParser. Usually you do not need to change this value.

If you are not implementing inbound e-mail processing for this mailer, leave the default as a placeholder value.

Note: You do not need to change the value of the Email Parser parameter if you select the Direct Response parameter. The notification mailer automatically switches to the alternate e-mail parser when the Direct Response parameter is selected.

- **Alternate Email Parser** - The Java class used to parse an incoming notification response e-mail formatted according to the direct response method and to create an XML document for the response. The notification mailer uses this parser when the

Direct Response parameter is selected. The default alternate e-mail parser provided by Oracle Workflow is named `oracle.apps.fnd.wf.mailer.DirectEmailParser`. Usually you do not need to change this value.

If you are not implementing inbound e-mail processing for this mailer, leave the default as a placeholder value.

Note: You do not need to change the value of the Alternate Email Parser parameter if you deselect the Direct Response parameter. The notification mailer automatically switches to the standard e-mail parser when the Direct Response parameter is deselected.

- **Expunge Inbox on Close** - Select this parameter to purge deleted messages from the inbox folder when the notification mailer closes this folder. If you do not select this parameter, copies of messages that were moved to the discard or processed folders remain in the inbox, in a deleted state, until you manually expunge them using your e-mail application.

Inbound E-mail Account

- **Inbound Protocol** - Oracle Workflow currently supports the IMAP protocol for inbound e-mail.
- **Inbound Server Name** - The name of the inbound mail server. Note that you must specify the actual host name for the server. Do not use `localhost` as the setting for this parameter. You can optionally specify the port number to use on that server. If you do not specify a port number, the notification mailer uses port 143 by default. Specify the server in the following format: `<server_name>[:<port_number>]`

For example: `myimapserver.mycompany.com:143`

If you are not implementing inbound e-mail processing for this mailer, enter a placeholder value.

- **Username** - The user name of the mail account that the notification mailer uses to receive e-mail messages.

If you are not implementing inbound e-mail processing for this mailer, enter a placeholder value.

- **Password** - The password for the mail account specified in the Username parameter. The password value is masked as asterisks in the display and is stored in encrypted form.

If you are not implementing inbound e-mail processing for this mailer, enter a placeholder value.

- **Inbox Folder** - The name of the folder from which the notification mailer receives

inbound messages. This value is case-insensitive. The default value is `INBOX`. The inbox must be separate from the processed and discard folders. Each notification mailer that performs inbound processing should have its own separate inbox.

Note: Usually, you use a dedicated mail account for notification mailer processing. If you want to use a mail account for the notification mailer that you also use for other purposes, you should create a folder in that account where you will place inbound messages destined for the notification mailer and specify that folder in the `Inbox Folder` parameter. Otherwise, the notification mailer will attempt to process all messages in the regular inbox and discard any messages that are not notification responses. If you do specify a separate folder to use as the notification mailer inbox folder, however, you must move messages from the regular inbox to that separate folder yourself. Depending on your mail program, you may be able to create a filter in the mail account to move such messages automatically. Use your e-mail client to create the separate folder. A notification mailer may not be able to access folders that were created using command line tools outside the e-mail client.

If you are not implementing inbound e-mail processing for this mailer, leave the default as a placeholder value.

- **Inbound Connection Timeout** - The maximum amount of time, in seconds, that the notification mailer will wait to establish a connection to the inbound server before timing out. The default inbound connection timeout period for a notification mailer is 120 seconds.
- **Inbound Message Fetch Size** - The maximum number of messages that the notification mailer can fetch from the inbox at one time. The default inbound message fetch size is 100 messages.
- **Maximum Ignore List Size** - The maximum number of notification IDs that the notification mailer can store in its ignore list, indicating that this notification mailer will make no further attempts to process them. For example, if the mailer encountered a connection error while processing a notification, that notification ID is temporarily added to the ignore list, and is then removed from the list the next time the inbox folder is successfully closed. The default maximum ignore list size is 1000. Usually you do not need to change this value.

Note: If the notification mailer finds additional messages to be ignored in the inbox when the ignore list is already full, the notification mailer removes the oldest notification IDs from the list and adds the new notification IDs instead.

- **Inbound SSL Enabled** - Select this parameter to enable the notification mailer to use SSL for connections to the IMAP server. Deselect this parameter to use non-SSL connections.

Note: When you enable SSL, the notification mailer connects to the IMAP server through port 993 by default. You can optionally specify a different port number along with the IMAP server name in the Inbound Server Name parameter.

Before you can use SSL, you must also complete additional setup steps. See: *Connecting to Mail Servers Through SSL, Oracle Workflow Administrator's Guide.*

Outbound E-mail Account

- **Outbound Protocol** - Oracle Workflow currently supports the SMTP protocol for outbound e-mail.
- **Outbound Server Name** - The name of the outbound mail server. Oracle strongly recommends that you specify the actual host name for the SMTP server. However, you can specify `localhost` as the setting for this parameter if you ensure that an SMTP server is configured to send e-mails to all valid domains on each host where concurrent managers run. If you have implemented Parallel Concurrent Processing to allow concurrent processing activities to be distributed across multiple nodes in a cluster system, then you must configure an SMTP server on every node. Otherwise, if a concurrent manager attempts to execute outbound notification mailer processing on a node without an SMTP server, the processing will fail. Also, when you save the configuration, Oracle Workflow Manager tests the connection to the SMTP server from within the Web tier host. Consequently, if you set the outbound server name to `localhost`, you should ensure that an SMTP server is configured on the Web tier host as well.

You can optionally specify the port number to use on that server. If you do not specify a port number, the notification mailer uses port 25 by default. Specify the server in the following format: `<server_name>[:<port_number>]`

For example: `mysmtpserver.mycompany.com:25`

If you are not implementing outbound e-mail processing for this mailer, enter a placeholder value.

- **Username** - If the outbound SMTP server is configured to require authentication, enter the user name of the account that the notification mailer uses to connect to the SMTP server.
- **Password** - If the outbound SMTP server is configured to require authentication, enter the password for the account specified in the Username parameter. The password value is masked as asterisks in the display and is stored in encrypted form.

- **Test Address** - This parameter has been replaced by the override e-mail address, which is available through the Component Details page for a notification mailer. See: Reviewing Service Component Details, page 15-13.
- **Outbound Connection Timeout** - The maximum amount of time, in seconds, that the notification mailer will wait to establish a connection to the outbound server before timing out. The default outbound connection timeout period for a notification mailer is 120 seconds.
- **Outbound SSL Enabled** - Select this parameter to enable the notification mailer to use Secure Sockets Layer (SSL) for connections to the SMTP server. Deselect this parameter to use non-SSL connections.

Note: When you enable SSL, the notification mailer connects to the SMTP server through port 465 by default. You can optionally specify a different port number along with the SMTP server name in the Outbound Server Name parameter.

Before you can use SSL, you must also complete additional setup steps. See: Connecting to Mail Servers Through SSL, *Oracle Workflow Administrator's Guide*.

E-mail Processing

- **Processed Folder** - The name of the mail folder where the notification mailer places successfully processed notification messages. This value is case-insensitive. The processed folder must be separate from the inbox and the discard folder.

The default value for this parameter is `PROCESS`. If you enabled inbound processing in the Basic Configuration page and the mail account you specified did not already include a folder named `PROCESS`, Oracle Workflow automatically created a folder with this name in that account when you completed the basic notification mailer configuration.

You can optionally specify the name of a different folder in this parameter. In this case, ensure that you use your e-mail client to create the folder. A notification mailer may not be able to access folders that were created using command line tools outside the e-mail client.

Note: The notification mailer does not perform any further operations on messages in the processed folder. You can review, back up, or delete these messages through your e-mail application if necessary.

If you are not implementing inbound e-mail processing for this mailer, leave the default as a placeholder value.

- **Discard Folder** - The name of the mail folder where the notification mailer places

incoming messages that are not recognized as notification messages. This value is case-insensitive. The discard folder must be separate from the inbox and the processed folder.

The default value for this parameter is `DISCARD`. If you enabled inbound processing in the Basic Configuration page and the mail account you specified did not already include a folder named `DISCARD`, Oracle Workflow automatically created a folder with this name in that account when you completed the basic notification mailer configuration.

You can optionally specify the name of a different folder in this parameter. In this case, ensure that you use your e-mail client to create the folder. A notification mailer may not be able to access folders that were created using command line tools outside the e-mail client.

Note: The notification mailer does not perform any further operations on messages in the discard folder. You can review, back up, or delete these messages through your e-mail application if necessary.

If you are not implementing inbound e-mail processing for this mailer, leave the default as a placeholder value.

- **Allow Forwarded Response** - Indicate whether to allow a user to respond by e-mail to an e-mail notification that has been forwarded from another role. This parameter is selected by default.
 - If Allow Forwarded Response is selected, the notification mailer never checks the "From" e-mail address of the notification response and always allows the response to be processed.

Note: Note that there are limitations when you deselect Allow Forwarded Response. For example, suppose a notification is sent to a distribution list mail alias that does not have a user/role relationship in the Oracle Workflow directory service. If any user from the distribution list responds to the notification, the notification mailer will always treat that notification response as unsolicited mail, because the "From" e-mail address, which is an individual user's e-mail address, will never match the distribution list mail alias.

- If Allow Forwarded Response is deselected, the notification mailer will check whether the "From" e-mail address of the notification response exactly matches the e-mail address of the recorded recipient role or the e-mail address of a user in that role. If the two e-mail addresses match exactly, meaning the notification was not forwarded or was forwarded according to a valid routing rule, the

notification mailer treats the response as a valid response. If the two e-mail addresses do not match exactly, meaning the notification was simply forwarded using the e-mail Forward command, the notification mailer does not process the response and treats it as unsolicited mail.

Note: To enhance security, Oracle Workflow does not allow a notification to be reassigned to the process owner who initiated the workflow, nor to the from role for the notification, when the reassignment is attempted through the Worklist pages or through a vacation rule. However, if you select the Allow Forwarded Response parameter, then a user specified as the process owner or the from role can still respond to the notification through e-mail if the original recipient forwards it through e-mail. To prevent this possibility, deselect the Allow Forwarded Response parameter.

To cancel any changes on this page, click the Cancel button.

To return to the previous step of the configuration wizard, click the Back button.

To save these settings and proceed to the next step of the configuration wizard, click the Next button.

Note: When you click the Next button, the configuration wizard validates the parameters you entered. If the inbound thread count is set to 1, the configuration wizard also verifies that it can connect to the e-mail account on the specified inbound mail server with the specified user name and password, and that the folders specified in the Processed Folder and Discard Folder parameters exist in that e-mail account. If the parameters are successfully validated, and the notification mailer is currently running, then Oracle Workflow Manager immediately refreshes the notification mailer with the new parameters.

Message Generation

This page lets you define message generation parameters for the notification mailer. Some parameters are already set to required values and cannot be modified. You must set parameters marked with an asterisk (*) to appropriate values for your environment before you can run the notification mailer. A refresh icon identifies attributes that can be refreshed dynamically while the service component is running. If the notification mailer is currently running, parameters marked with a refresh icon will be refreshed immediately when you select the Next button or the Finish button.

Send

- **From** - A value that appears in the From field of the message header of a notification e-mail. You can specify the From parameter value either as a display name only, or as a full RFC822-compliant address.
 - If you specify a display name only, the notification mailer adds the e-mail address from the Reply-to Address parameter to create a full RFC822-compliant address for the From message header. The full address is created in the following format: "*Display Name*" <*reply_to_address*>
 - If you specify a full RFC822-compliant address, the notification mailer uses only that From parameter value in the From message header, and does not include the Reply-to Address value.

If a particular notification message has the special #WFM_FROM message attribute defined, however, the notification mailer will use the #WFM_FROM attribute value in the From field for that message, instead of the From parameter value.

The default From parameter value for the seeded notification mailer service component is `Workflow Mailer`. For other notification mailers, if you selected the Inbound Processing parameter in the Basic Configuration page, Oracle Workflow by default sets the From parameter to the name portion of the reply-to address specified in the Basic Configuration page. For example, if the reply-to address is `Workflow@mycompany.com`, Oracle Workflow sets the From parameter to `Workflow`.

If you deselected the Inbound Processing parameter in the Basic Configuration page, Oracle Workflow by default sets the From parameter to `nobody@<server_name>`, where `<server_name>` is the name of the outbound SMTP mail server specified in the Basic Configuration page.

If you are not implementing outbound e-mail processing for this mailer, leave the default as a placeholder value.

- **Reply-to Address** - The address of the e-mail account that receives incoming messages, to which notification responses should be sent. This value must be a full RFC822-compliant e-mail address.

If a particular notification message has the special #WFM_REPLYTO message attribute defined, however, the notification mailer will use the #WFM_REPLYTO attribute value as the reply address for that message, instead of the Reply-to Address parameter value.

Note: If the From parameter value is specified as a display name only, then the notification mailer also uses the reply-to e-mail address together with that display name to create a full RFC822-compliant address for the From field of the message

header.

If you deselected the Inbound Processing parameter in the Basic Configuration page, Oracle Workflow by default sets the Reply-to Address parameter to `nobody@<server_name>`, where `<server_name>` is the name of the outbound SMTP mail server specified in the Basic Configuration page. If you are not implementing inbound e-mail processing for this mailer, leave the default as a placeholder value.

- **HTML Agent** - The base URL that identifies the HTML agent that handles HTML notification responses. This URL is required to support e-mail notifications with HTML attachments. Usually the HTML agent specified here can match the value of the Applications Servlet Agent profile option; however, you can optionally specify a different HTML agent for a particular notification mailer. The HTML agent should be specified in the following format:

```
http://<server_name:port>/OA_HTML/
```

where `<server_name:port>` represents the server and TCP/IP port number on which your servlet agent accepts requests.

Note: The notification mailer can also still handle an HTML agent value in the previous format:

```
http://<server_name:port>/pls/wf
```

If a particular notification message has the special `#WFM_HTMLAGENT` message attribute defined, however, the notification mailer will use the `#WFM_HTMLAGENT` attribute value as the HTML agent for that message, instead of the HTML Agent mailer parameter value.

- **Message Formatter** - Oracle Workflow uses the `oracle.apps.fnd.wf.mailer.NotificationFormatter` Java class to generate notification messages.
- **Framework User** - The numerical user ID for the user through which a notification mailer accesses Oracle Application Framework content for inclusion in e-mail notifications. The Framework user must have workflow administrator privileges in order to access the content for every user's notifications.

The default value for this parameter is `0`, which is the user ID for the `SYSADMIN` user. This setting lets the notification mailer access Oracle Application Framework content through the `SYSADMIN` user, which is also the default workflow administrator role. If you change the Workflow System Administrator preference, check the Framework User parameter to ensure that the user accessed by the notification mailer has workflow administrator privileges. Set the Framework User parameter to a user that is a member of the Workflow System Administrator role, or, if you set the Workflow System Administrator preference to a responsibility, set

the Framework User parameter to a user that has that responsibility. See: Setting Global User Preferences, *Oracle Workflow Administrator's Guide*.

Note: You can use the Workflow Mailer URL Access Tester page to test whether Oracle Application Framework content can be generated correctly for e-mail notifications. See: Testing Mailer URL Access, *Oracle Workflow Administrator's Guide*.

- **Framework Responsibility** - The numerical responsibility ID for the responsibility through which a notification mailer accesses Oracle Application Framework content for inclusion in e-mail notifications. The user specified in the Framework User parameter must have this responsibility assigned. The default value for this parameter is 20420, which is the responsibility ID for the System Administrator responsibility.
- **Framework Application ID** - The numerical application ID for the application through which a notification mailer accesses Oracle Application Framework content for inclusion in e-mail notifications. The responsibility specified in the Framework Responsibility parameter must be assigned to this application. The default value for this parameter is 1, which is the application ID for the System Administration application.
- **Framework URL Timeout** - The maximum amount of time, in seconds, that the notification mailer will wait to access a URL for Oracle Application Framework content before timing out. The default Framework URL timeout period for a notification mailer is 30 seconds.
- **Attach Images to Outbound Emails** - Select this parameter to attach any images referenced in HTML content included in a message, such as Oracle Application Framework content, to outbound notification e-mail messages. Deselect this parameter to display the image references as off-page URLs instead of attaching the images.
- **Attach Stylesheet to Outbound Email** - Select this parameter to attach any stylesheet referenced in HTML content included in a message, such as Oracle Application Framework content, to outbound notification e-mail messages. Deselect this parameter to display the stylesheet reference as a URL instead of attaching the stylesheet.

Note: E-mail clients vary in their support for stylesheet references within HTML content in the body of an e-mail. Some e-mail clients do not support references to a stylesheet that is attached to the e-mail, while others do not support any form of stylesheet references within HTML content at all. Consequently, attaching a

stylesheet may not have the same effect in all e-mail clients.

- **Autoclose FYI** - Indicate whether this notification mailer automatically closes notifications that do not require a response, such as FYI (For Your Information) notifications, after sending the notifications by e-mail. This parameter is selected by default. If Autoclose FYI is deselected, all FYI notifications will remain open in the Worklist until users manually close these notifications.
- **Direct Response** - By default, notification mailers require a response format for plain text notifications called the templated response method. Select this parameter to use the direct response method instead.
 - With the templated response method, a notification mailer sends plain text notifications requiring a templated response to users with a notification preference of MAILTEXT or MAILATTH. Users must reply using a template of response prompts and enter their response values between the quotes following each prompt.
 - With the direct response method, a notification mailer sends plain text notifications requiring a direct response to users with a notification preference of MAILTEXT or MAILATTH. Users must enter their response values directly as the first lines of a reply.

Note: Responses that are generated automatically from an HTML-formatted notification or attachment must always use a response template, regardless of which response method you select.

See: Workflow Open Mail (Templated) Message, *Oracle Workflow Administrator's Guide*, Workflow Open Mail (Direct) Message, *Oracle Workflow Administrator's Guide*, To Respond to a Plain Text E-mail Notification Using Templated Response, *Oracle Workflow User's Guide*, To Respond to a Plain Text E-mail Notification Using Direct Response, *Oracle Workflow User's Guide*, and Example 'Respond' Message Attributes, *Oracle Workflow Developer's Guide*.

- **Reset NLS** - Select this parameter if you want the notification mailer to encode each notification message with character encoding according to the notification recipient's preferred language. Deselect this parameter if you want the notification mailer to use the same character encoding for all notification messages. This parameter is deselected by default.

If a particular notification message has the special #WFM_RESET-NLS message attribute defined, however, then the notification mailer will use the #WFM_RESET-NLS attribute value to determine whether to encode the message with character encoding for the preferred language, instead of using the Reset NLS parameter value. Additionally, you can use the Character Encoding Configuration

page in the Workflow administrator Web pages to specify the character encoding that you want to use under either Reset NLS setting, overriding the default logic for determining the character encoding.

- If the Reset NLS parameter is deselected at the notification mailer level and is not overridden at the message level, or if the #WFM_RESET-NLS message attribute is set to N at the message level, then the notification mailer uses the same character encoding for all notification messages.
 - By default, the notification mailer uses the default character encoding for the database.
 - If you want to use different character encoding instead, then you can specify the override character encoding in the Character Encoding Configuration page.
- If the Reset NLS parameter is selected at the notification mailer level and is not overridden at the message level, or if the #WFM_RESET-NLS message attribute is set to Y at the message level, then the notification mailer encodes each notification message with character encoding according to the notification recipient's preferred language.
 - By default, the notification mailer uses the following logic to determine the character encoding for the message.
 - If the notification recipient has specified both a preferred language and a preferred territory, then the notification mailer uses the character encoding listed in the WF_LANGUAGES table for that language and territory.
 - If no preferred territory is specified, then the notification mailer uses the character encoding associated with the first entry it encounters in the WF_LANGUAGES table for the user's preferred language.
 - If no preferred language is specified, then the notification mailer uses the character set listed in WF_LANGUAGES for the language AMERICAN and territory AMERICA.
 - If you want to use different character encoding instead, then you can use the Character Encoding Configuration page to specify the override character encoding for each language installed in your database. In this case the notification mailer uses the override character encoding configured for the notification recipient's preferred language.

Note: You can also review and update the Reset NLS parameter setting for your notification mailers in the Character Encoding

Configuration page Any changes you make in that page will be reflected in the notification mailer configuration wizard as well.

See: *Configuring Character Encoding for Notification Mailers, Oracle Workflow Administrator's Guide.*

- **Inline Attachments** - Select this parameter to set the Content-Disposition MIME header to `inline` for attachments to notification messages, including the Notification Detail Link, HTML Message Body, Notification References containing attached URLs, and attached PL/SQL documents. Deselect this parameter to set the Content-Disposition MIME header to `attachment` for these attachments. For example, if your e-mail application cannot display HTML content such as the Notification Detail Link inline, deselect this parameter to display this link as an attachment instead. Note, however, that some e-mail clients may not support the Content-Disposition header, or may support it in varying ways. Consequently, the Inline Attachment setting may not always have the desired effect, depending on the e-mail clients with which users read their e-mail messages.
- **Send Warning for Unsolicited E-mail** - Select this parameter to allow the notification mailer to send back a warning message the first time it receives an unsolicited e-mail message from a particular e-mail address. Deselect this parameter to prevent the notification mailer from sending warning messages.
- **Send E-mails for Canceled Notifications** - Select this parameter to allow the notification mailer to send cancellation messages to users when previously sent notifications are canceled. Deselect this parameter to prevent the notification mailer from sending cancellation messages.

If you set up multiple notification mailers in the same Oracle E-Business Suite instance, you must set this parameter to the same setting for all the notification mailers.

Templates

This region lets you specify the message templates you want to use to generate e-mail notifications. The template for a particular type of e-mail notification determines the basic format of the notification, including what header information to include, and whether and where to include details such as the message due date and priority.

Oracle Workflow provides a set of standard templates in the System: Mailer item type, which are used by default. It is not recommended to modify the standard templates. However, you can customize the message templates used to send your e-mail notifications by creating your own custom message templates in the System: Mailer item type using the Workflow Builder, and assigning these templates to a particular notification mailer service component in this region. The templates assigned to a mailer override the default System: Mailer templates.

Additionally, you can create your own custom message templates in a custom item type using the Workflow Builder, and assign these templates to a particular notification in a workflow process by defining special message attributes. In this case the templates assigned to the notification override both the templates assigned to a mailer and the default System: Mailer templates.

If you are not implementing outbound e-mail processing for this mailer, leave the default templates as placeholder values.

- **Attached URLs** - The notification mailer uses this template to create the Notification References attachment for HTML-formatted notification messages that include URL attributes with Attach Content checked. The template must include a list with links to each URL.
- **Outbound Closed Notification** - The notification mailer uses this template to inform the recipient that a previously sent notification is now closed.
- **Outbound Cancelled Notification** - The notification mailer uses this template to inform the recipient that a previously sent notification is canceled. You can optionally use the Send E-mails for Canceled Notifications parameter to specify whether or not the notification mailer should send Outbound Cancelled Notification messages.
- **Invalid Response Notification** - The notification mailer uses this template to inform a user that the user responded incorrectly to a notification. For example, if a response message from a user contains a valid notification ID (NID) line matching it with a notification, but does not contain any response value or contains an invalid response value, the notification mailer sends an Invalid Response Notification message to the user. This template must describe how to respond to the notification correctly.
- **Open Notification** - If you are using the default response method, which is templated response, the notification mailer uses this template to send open notifications that require a response. This message template must provide a response template for the recipient as well as instructions on how to use the response template.

Note: In addition to the default Workflow Open Mail (Templated) message template, Oracle Workflow also provides a predefined template called Workflow Open Mail (Outlook Express). This template is provided to accommodate e-mail applications such as Microsoft Outlook Express or other e-mail clients that cannot process the response links included in the HTML bodies of the Workflow Open Mail (Templated) and Workflow Open Mail (Direct) templates. If you use one of these e-mail clients, you can select the Workflow Open Mail (Outlook Express) message

template to have HTML e-mail notifications include a link to the Notification Details Web page which lets users respond to the notification there.

If you are configuring this notification mailer for outbound message processing only and you are not implementing any corresponding inbound e-mail response processing, then you should set the Open Notification parameter to a message template that does not request a response by e-mail, but instead directs recipients to respond from the Notification Details Web page. For example, you can select the Workflow View From UI message template provided by Oracle Workflow, or create your own custom message template.

If you selected the Inbound Processing parameter in the Basic Configuration page, the Open Notification parameter is set to the Workflow Open Mail (Templated) message template by default. If you deselected the Inbound Processing parameter, the Open Notification parameter is set to the Workflow Open Mail (Outlook Express) message template by default.

Note: The plain text version of the Workflow Open Mail (Outlook Express) message requests a response by e-mail. If you disable inbound processing, ensure that your users do not have a notification preference of MAILTEXT or MAILATTH. Alternatively, if you disable inbound processing and you want users to receive plain text notifications, specify a message template that directs recipients to respond from the Notification Details Web page.

- **Open Notification (Direct Response Parsing)** - If you select the Direct Response parameter, the notification mailer uses this template to send open notifications that require a response. The response instructions in the plain text message body must describe how to reply using the direct response method. This message is used for notifications sent to performers with a notification preference of MAILTEXT or MAILATTH. The response instructions in the HTML-formatted message body must describe how to reply using the automatically generated response template. This message is used for notifications sent to performers with a notification preference of MAILHTML or MAILHTML2, and is also attached to notifications sent to performers with a notification preference of MAILATTH.

Note: Responses that are generated automatically from an HTML-formatted notification or attachment must always use a response template, regardless of which response method you select.

Note: If you are configuring this notification mailer for outbound

message processing only and you are not implementing any corresponding inbound e-mail response processing, then you should set the Open Notification (Direct Response Parsing) parameter to a message template that does not request a response by e-mail, but instead directs recipients to respond from the Notification Details Web page. For example, you can select the Workflow View From UI message template provided by Oracle Workflow, or create your own custom message template.

See: Workflow Open Mail (Templated) Message, *Oracle Workflow Administrator's Guide*, Workflow Open Mail (Direct) Message, *Oracle Workflow Administrator's Guide*, To Respond to a Plain Text E-mail Notification Using Templated Response, *Oracle Workflow User's Guide*, To Respond to a Plain Text E-mail Notification Using Direct Response, *Oracle Workflow User's Guide*, and Example 'Respond' Message Attributes, *Oracle Workflow Developer's Guide*.

- **Open FYI Notification** - The notification mailer uses this template to send notifications that do not require a response. The template must indicate that the notification is for your information (FYI) and does not require a response.
- **Outbound Summary Notification** - This template is no longer used.
- **Outbound Warning Notification** - The notification mailer uses this template to send a message to a user the first time it receives unsolicited mail from that user. For example, if a message from a user does not contain a notification ID (NID) line matching it with a notification, or contains an incorrectly formatted NID line, the notification mailer sends an Outbound Warning Notification message to the user. You can optionally use the Send Warning for Unsolicited E-mail parameter to specify whether or not the notification mailer should send Outbound Warning Notification messages.
- **Open Notification (More Information Request)** - The notification mailer uses this template to send a request for more information about a notification from one user to another user.

Note: If you use an e-mail application such as Microsoft Outlook Express that cannot process the response link included in the default Workflow Open Mail (More Information Request) message template, you can select an alternative template named Workflow More Information Request (Outlook Express) instead. In particular, if you set the Open Notification parameter to use the Workflow Open Mail (Outlook Express) message, then you should also set the Open Notification (More Information Request) parameter to use the Workflow More Information Request (Outlook Express) message.

- **Outbound HTML Summary Notification** - The notification mailer uses this template to send a summary of currently open workflow notifications to users and roles that have their notification preference set to SUMMARY or SUMHTML in the Oracle Workflow directory service.

To cancel any changes on this page, click the Cancel button.

To return to the previous step of the configuration wizard, click the Back button.

To save these settings and proceed to the next step of the configuration wizard, click the Next button.

To save these settings and proceed to the last step of the configuration wizard, click the Finish button.

Note: When you click the Next or Finish button, the configuration wizard validates the parameters you entered. If the parameters are successfully validated, and the notification mailer is currently running, then Oracle Workflow Manager immediately refreshes the notification mailer with the new parameters.

Scheduling Events

This page lets you schedule events to control the running of the service component. The events are raised at the scheduled time by DBMS jobs. For a notification mailer service component, you can schedule the following events:

- Start
- Refresh
- Suspend
- Resume
- Stop
- Launch Summary Notifications

For each event, the list displays the event name, date and time when the event is first scheduled to be raised, the interval in minutes at which the event is reraised, and, for a Refresh event, any parameters to be refreshed. You can specify the following refreshable parameters, using the parameters' internal names, when you refresh the notification mailer.

- `PROCESSOR_IN_THREAD_COUNT` - Inbound Thread Count
- `PROCESSOR_OUT_THREAD_COUNT` - Outbound Thread Count

- COMPONENT_LOG_LEVEL - Log Level, specified as a numerical value
 - 1 - Statement
 - 2 - Procedure
 - 3 - Event
 - 4 - Exception
 - 5 - Error
 - 6 - Unexpected
- EXPUNGE_ON_CLOSE - Expunge Inbox on Close
- ALLOW_FORWARDED_RESPONSE - Allow Forwarded Response
- FROM - From
- REPLYTO - Reply-to Address
- HTMLAGENT - HTML Agent
- ATTACH_IMAGES - Attach Images to Outbound E-mails
- ATTACH_STYLESHEET - Attach Stylesheet to Outbound E-mail
- AUTOCLOSE_FYI - Autoclose FYI
- RESET-NLS - Reset NLS
- INLINE_ATTACHMENT - Inline Attachments
- SEND_UNSOLICITED_WARNING - Send Warning for Unsolicited E-mail
- ATTACHED_URLS - Attached URLs
- CLOSED - Outbound Closed Notification
- CANCELED - Outbound Cancelled Notification
- OPEN_INVALID - Invalid Response Notification
- OPEN_MAIL - Open Notification
- OPEN_MAIL_DIRECT - Open Notification (Direct Response Parsing)

- OPEN_MAIL_FYI - Open FYI Notification
- SUMMARY - Outbound Summary Notification
- WARNING - Outbound Warning Notification
- OPEN_MORE_INFO - Open Notification (More Information Request)
- SUMHTML - Outbound HTML Summary Notification

To schedule events:

- If no events are currently scheduled, click the Add a Row button to add a new row to the list of events and enter the information for the event.
 - Select the event for the command you want to schedule.
 - Select the date when you want the event to be raised first.
 - Select the hour and minute to specify the time on the specified date when you want the event to be raised first. The hour values are in a twenty-four hour format. For example, select 00 for midnight, or 23 for 11 PM.
 - If you want to raise the event periodically, enter the time interval in minutes at which you want to raise the event. If you do not specify a repeating interval, the event is raised only once.
 - If you choose the refresh event, you can optionally enter any parameters you want to include with the event in order to refresh the notification mailer configuration parameters with those values when the event is raised. Specify the parameter names and values in the following format, separating the parameters with a colon (:):
`internal_parameter_name=parameter_value`
 For example: `PROCESSOR_OUT_THREAD_COUNT=3`
 If a parameter value itself contains a colon (:), then precede the colon with a backslash (\) as an escape character, as follows:
`\:`
 For example:
`OPEN_MAIL_DIRECT=WFMAIL\ :OPEN_MAIL_DIRECT`
- To schedule another event, click the Add Another Row button and enter the information for the event.
- To remove an event, select the event and click the Remove button.

To cancel any changes on this page, click the Cancel button.

To return to the previous step of the configuration wizard, click the Back button.

To save these settings and proceed to the next step of the configuration wizard, click the Next button.

To save these settings and proceed to the last step of the configuration wizard, click the Finish button.

Note: The configuration wizard verifies that an event is specified for every row in the list when you click the Next or Finish button. If you do not want to schedule another event, remove any empty rows before proceeding.

Tags

This page lets you enter patterns of text found in unusual messages and the status you want to assign to an inbound message if it contains any of those patterns. For example, unusual messages include bounced or returned messages and auto-reply messages such as those sent by vacation daemons, mass mailing lists, and so on. Since different mail systems vary in how they identify bounced, undeliverable, or otherwise invalid messages, you can use notification mailer tags to specify how your mail system identifies those stray messages and how you want the notification mailer to handle those messages should it come across them.

Oracle Workflow provides several predefined tags for text commonly found in undeliverable or auto-reply messages. For each tag, the list displays the pattern, which is the string of text to look for in the From line, Subject line, or body of the message, and the action, which is the mail status to assign to the message if that pattern is found. The notification mailer handles messages according to these mail status values, as follows:

- UNDELVRD - Moves the message to the discard folder and updates the notification's mail status to FAILED. Additionally, the notification preference of the recipient of the notification is updated to DISABLED. No error process is initiated for this notification activity. However, after correcting the issues that prevented the e-mail from being sent, you can reset the user's notification preference and then run the Resend Failed/Error Workflow Notifications program to re-enqueue failed notifications on the notification mailer's outbound queue. See: *Handling Mailer Errors, Oracle Workflow Administrator's Guide*.
- Unavailable - Moves the message to the discard folder and continues waiting for a reply to the notification since the notification's status is still OPEN, but its mail status is updated to UNAVAIL. This status is purely informative, as no further processing occurs with this notification.
- Ignore - Moves the message to the discard folder and continues waiting for a valid reply to the open notification. The notification's status is still OPEN and its mail status is still SENT.

You can define additional tags for other patterns you want the notification mailer to handle automatically.

- To add a new tag, click the Add Another Row button, enter the text pattern in the Pattern column, and select the status you want to assign to messages containing that pattern in the Action column.
- To remove a tag, select the tag and click the Remove button. You can only remove custom tags that you defined. You cannot remove predefined tags provided by Oracle Workflow.

Note: It is important that you uniquely identify bounced messages and auto-replies by defining tags to distinguish them from normal responses. If you do not identify bounced and auto-reply messages, the notification mailer can mistake these as invalid responses, send an Invalid Response Notification message, and continue to wait for a reply. In both cases a perpetual loop would occur where the notification mailer continues sending out an 'Invalid' message and the 'Invalid' message bounces back or is auto-replied each time.

Note: Only a message response that contains a notification ID can be handled through the FAILED and UNAVAIL mail statuses. If the notification mailer receives a message response that does not contain a notification ID, it moves the message response to the discard folder. If the Send Warning for Unsolicited E-mail parameter is selected, then for the first such message from a particular e-mail address, the notification mailer also sends an Outbound Warning Notification message to the sender to warn that it received unsolicited mail.

Note: If a message response matches more than one pattern in the list of tags, the message is tagged with the status of the first tag it matches. That is, the notification mailer performs a top to bottom comparison against the tag list. Due to this behavior, you should prioritize your patterns listing the UNDELVRD tags first, followed by the Unavailable and then Ignore tags.

To cancel any changes on this page, click the Cancel button.

To return to the previous step of the configuration wizard, click the Back button.

To save these settings and proceed to the next step of the configuration wizard, click the Next button.

To save these settings and proceed to the last step of the configuration wizard, click the Finish button.

Test

This page lets you test the configuration for a notification mailer that performs outbound e-mail processing by sending sample notification messages. Select the recipient role to which the messages should be sent, and click the Send Test Message button.

Note: To send a test message successfully, you must select a recipient role that either has a valid e-mail address defined, or that has members with valid e-mail addresses defined. The recipient role must also have a notification preference that includes individual e-mail notifications.

If you set an override e-mail address for the notification mailer, the Test page displays that address. In this case the test message is sent to the override address rather than the e-mail address of the recipient role. However, you must still select a recipient role to enable the notification mailer to send the test messages. See: *Reviewing Service Component Details*, page 15-13.

Oracle Workflow sends two test messages to the recipient role: one message with content built using PL/SQL and one message with Oracle Application Framework content. Check the e-mail account for the recipient role to view the test messages and reply to them with the Acknowledge response. If you did not implement inbound e-mail processing for this mailer, use the Worklist pages to respond to the test messages after viewing the outbound messages in e-mail. After you acknowledge both test messages, Oracle Workflow sends a confirmation message to the same recipient role to show that the notification mailer successfully processed the inbound response e-mails.

If you do not receive the test messages or the response confirmation message, or if the message content does not appear correctly, check the notification mailer setup, including the mail servers and the mailer configuration parameters. In particular, if the Oracle Application Framework content does not appear correctly, check the Application Framework Agent and WF: Workflow Mailer Framework Web Agent profile options, as well as the Framework User, Framework Responsibility, Framework Application ID, and Framework URL Timeout parameters in the advanced configuration wizard. See: *Setting Up a Notification Mailer*, page 15-22 and *Message Generation*, page 15-48.

Note: Oracle Workflow sends the test messages by launching the PLSQL/OAFwk Response Test Process in the System: Tests (WFTESTS) item type. This item type is stored in a file called wftstmlr.wft in the `$FND_TOP/import/<lang>` subdirectory. You can optionally use the Status Monitor to check the status of the test process.

To exit the advanced configuration wizard, click the Cancel button.

To return to the previous step of the configuration wizard, click the Back button.

To proceed to the next step of the configuration wizard, click the Next button.

To proceed to the last step of the configuration wizard, click the Finish button.

Review

This page lets you review the configuration parameter values that you set, the events that you scheduled, and the tags that you defined for this notification mailer service component.

- If you want to change any of these settings, return to the appropriate step in the configuration wizard to make your changes. To return to the previous step, click the Back button.
- To save these settings and finish the configuration, click the Finish button.

Agent Listeners

The Oracle Workflow Business Event System requires agent listeners to be scheduled to receive inbound event messages. An agent listener monitors a Business Event System agent for incoming messages and dequeues messages using the agent's queue handler. You should run agent listeners for your local inbound agents. Run PL/SQL agent listeners to process event subscriptions with a PL/SQL rule function in the database, and run Java agent listeners to process event subscriptions with a Java rule function in the middle tier.

When an event message is dequeued, the Event Manager begins subscription processing for the event. The Event Manager searches for and executes any active subscriptions by the local system to that event with a source type of External, and also any active subscriptions by the local system to the Any event with a source type of External. The agent listener exits after all event messages on the agent's queue have been dequeued.

The PL/SQL agent listener program is defined as a service component type in the Generic Service Component Framework. This framework helps to simplify and automate the management of background Java services.

Oracle Workflow provides several seeded agent listener service components to process messages on standard agents.

- Workflow Deferred Agent Listener - Handles messages on WF_DEFERRED to support deferred subscription processing. This service component is started automatically by its container.
- Workflow Deferred Notification Agent Listener - Handles notification messages on WF_DEFERRED to support outbound notification processing. This service component is started automatically by its container.
- Workflow Error Agent Listener - Handles messages on WF_ERROR to support

error handling for the Business Event System. This service component is started automatically by its container.

- Workflow Inbound Notifications Agent Listener - Handles messages on WF_NOTIFICATION_IN to support inbound e-mail notification processing. This service component is started automatically by its container.
- ECX Inbound Agent Listener - Handles message on ECX_INBOUND to support Oracle XML Gateway processing. This service component must be started manually. For more information, see the *Oracle XML Gateway User's Guide*.
- ECX Transaction Agent Listener - Handles message on ECX_TRANSACTION to support Oracle XML Gateway processing. This service component must be started manually. For more information, see the *Oracle XML Gateway User's Guide*.

You cannot delete the seeded agent listeners or edit their names, assigned agents, correlation ID values, or containers. However, if necessary you can update other configuration parameters, schedule control events, or manually choose control commands to start, stop, suspend, resume, or refresh the agent listeners.

You can also optionally create additional agent listener service components. For example, you can configure agent listeners for other inbound agents that you want to use for event message propagation, such as the standard WF_IN and WF_JMS_IN agents, or any custom agents. You can also configure an agent listener that only processes messages on a particular agent that are instances of a specific event.

In addition to the parameters in the configuration wizard, for both seeded and custom PL/SQL agent listeners, you can optionally set the following internal agent listener parameters.

- LISTENER_PROCESS_EVT_COUNT - Lets you specify the maximum number of event messages that the agent listener can process each time it runs, before returning control to its service component container.
- SQL_TRACE_LEVEL - Lets you enable SQL tracing at various levels or disable SQL tracing for the agent listener.
- NAVIGATION_RESET_THRESHOLD - Lets you reset the agent listener's navigation through waiting messages to include newly arrived messages, so that new high priority messages are processed sooner.

Use the `afsvcpup.sql` script to set these parameters. See: Scheduling Listeners for Local Inbound Agents, *Oracle Workflow Administrator's Guide* and To Set Internal Agent Listener Parameters, *Oracle Workflow Administrator's Guide*.

If you create custom agent listener service components, you can either assign them to the seeded container for agent listeners, named Workflow Agent Listener Service, or, based on the volume to be handled by the seeded container, you can also choose to create your own custom containers.

Before the seeded agent listener service components can run, the Workflow Agent Listener Service container which manages them must be first be started. You should ensure that this container is running. If you create your own custom containers for custom service components, ensure that those containers are running as well. Use the Service Instances page to start each container as a service instance in Generic Service Management (GSM). When the Workflow Agent Listener Service container is running, it automatically starts the Workflow Deferred Agent Listener, Workflow Deferred Notification Agent Listener, Workflow Error Agent Listener, and Workflow Inbound Notifications Agent Listener.

Agent Listener Configuration Wizard

The agent listener configuration wizard lets you configure an agent listener service component by defining general and detail attributes and scheduling control events. You can use the configuration wizard to configure a new agent listener service component, or to edit the configuration of an existing agent listener service component.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Service Components status icon > (B) Create > (B) Continue

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Service Components status icon > (B) Edit

Define

This page lets you define general attributes for the service component. Some attributes are already set to required values and cannot be modified. You must set attributes marked with an asterisk (*) to appropriate values for your environment before you can run the service component.

- **ID** - When you edit a previously created service component, the configuration wizard displays the identifier for the service component.
- **Status** - When you edit a previously created service component, the configuration wizard displays the status of the service component.
- **Name** - The name of the service component. This name must be unique.
- **Startup Mode** - Select Automatic, Manual, or On-Demand as the startup mode for the service component.
- **Container Type** - The container type to which this service component belongs, which is always Oracle Applications Generic Service Management (Oracle Applications GSM).
- **Inbound Agent** - The Business Event System agent that you want to monitor for inbound event messages.

- **Outbound Agent** - Leave this field blank. Agent listener service components do not use an outbound agent.
- **Correlation ID** - Optionally specify the Oracle Advanced Queuing (AQ) correlation ID of the event messages that you want the agent listener to process. The AQ correlation ID for an event message in the Business Event System is usually specified in the following format:

<event name>

By specifying a correlation ID in this attribute, you can dedicate the agent listener to listen only for messages that are instances of the specified event. You can also specify a partial value to listen for messages that are instances of any event whose name begins with the specified value.

Both dedicated and general agent listeners are compatible with each other. You can run several dedicated and general agent listeners for the same agent at the same time if you choose. The behavior of a general agent listener with a blank correlation ID depends on its agent.

- For the WF_DEFERRED agent only, a general agent listener does not process messages that match a dedicated agent listener's correlation ID, as long as the dedicated agent listener has a status of `Running`, `Stopped With Error`, or `System Deactivated`. If the dedicated agent listener has any other status, such as `User Deactivated` or `Suspended`, then the general agent listener does process the messages on WF_DEFERRED that match the dedicated agent listener's correlation ID.

For example, the seeded Workflow Deferred Notification Agent Listener has an AQ correlation ID of `oracle.apps.wf.notification.%`, meaning that this agent listener handles only notification event messages (those whose event name begins with `oracle.apps.wf.notification.`) on the WF_DEFERRED agent. However, the seeded Workflow Deferred Agent Listener does not have any correlation ID specified, so that it can process all event messages on the WF_DEFERRED agent that are not being handled by a dedicated agent listener. As long as the dedicated Workflow Deferred Notification Agent Listener has a status of `Running`, `Stopped With Error`, or `System Deactivated`, then the general Workflow Deferred Agent Listener processes all messages on the WF_DEFERRED agent except those whose event name begins with `oracle.apps.wf.notification.`, leaving those to be processed by the Workflow Deferred Notification Agent Listener.

- For all other agents, a general agent listener can process all messages on the agent. Even if you have configured a dedicated listener for a particular agent, a message that matches the dedicated agent listener's correlation ID may still be processed by a general listener if that listener is the first to access the message.

For example, the seeded Workflow Error Agent Listener and Workflow Inbound Notifications Agent Listener do not have any correlation ID specified

so that they can process all event messages on their respective agents.

Note: The AQ correlation ID is different than the correlation ID contained within the WF_EVENT_T event message structure.

To cancel the configuration without saving any changes, click the Cancel button.

To save these settings and proceed to the next step of the configuration wizard, click the Next button.

Details

This page lets you define detail attributes for the service component. You must set attributes marked with an asterisk (*) to appropriate values for your environment before you can run the service component. A refresh icon identifies attributes that can be refreshed dynamically while the service component is running.

- **ID** - When you edit a previously created service component, the configuration wizard displays the identifier for the service component.
- **Status** - When you edit a previously created service component, the configuration wizard displays the status of the service component.
- **Name** - The configuration wizard displays the name defined for the service component.
- **Container** - The container to which the service component will belong. Oracle Workflow provides a container called Workflow Agent Listener Service for agent listener service components.
- **Maximum Idle Time** - If you selected the On-Demand startup mode for the service component, enter the maximum time in minutes that the service component can remain idle before it is stopped. An on-demand component that is stopped in this way will be restarted by its container when it is needed again to process new messages.
- **Max Error Count** - The number of consecutive errors the service component can encounter before its container stops it and changes its status to Stopped with Error. If an error is resolved and processing can continue, the error count is reset. The default value for the maximum error count is 10.
- **Inbound Thread Count** - Set the inbound processing thread count to 1 (one) or higher to enable inbound message processing with this agent listener. Set the inbound thread count to 0 (zero) to disable this agent listener. The default value is 1. If this agent listener receives a high volume of inbound messages, you can set the inbound thread count to a higher value to increase throughput.

- **Outbound Thread Count** - Leave this parameter set to the default value of 0 (zero). Agent listener service components do not perform outbound message processing.
- **Log Level** - Select the level of detail for the information you want to record in the service component container log. The recommended log level, which is also the default value, is Error. Usually the log level only needs to be changed if you want to record additional detailed information for debugging purposes. You can choose the following levels:
 - 1 - Statement
 - 2 - Procedure
 - 3 - Event
 - 4 - Exception
 - 5 - Error
 - 6 - Unexpected
- **Processor Read Wait Timeout** - Specify the amount of time in seconds that the service component's processing thread continues to wait, after reading the last message from its assigned queue, before timing out. If another message is received before this time expires, that message is processed and the timeout period begins again. If the timeout period expires and no more messages have been received, the service component stops reading and its sleep time begins. The default read timeout period for an agent listener is 0 (zero) seconds.
- **Processor Min Loop Sleep** - Specify the minimum sleep time in seconds during which the service component waits, after its read timeout period expires, before it checks its queue for messages again. The default minimum sleep time for an agent listener is 120 seconds.
- **Processor Max Loop Sleep** - Specify the maximum sleep time in seconds if you want to increase the sleep time between read attempts when no messages are received. If you specify a maximum sleep time that is greater than the minimum sleep time, then the service component initially waits for the minimum sleep time after it finishes reading messages from its queue. If no messages are read in subsequent attempts, then the sleep time between read attempts gradually increases until the maximum sleep time is reached. Increasing the sleep time can help enhance performance if messages are received infrequently. You can also specify 0 (zero) for this parameter to indicate that the sleep time should not be increased. In this case, the service component always waits for the minimum sleep time between read attempts. The default value for an agent listener is 0 (zero).
- **Processor Error Loop Sleep** - Specify the sleep time in seconds during which the

service component waits, after an error occurs, before it attempts to begin processing again. The default error sleep time for an agent listener is 60 seconds.

- **Processor Close on Read Timeout** - Select this parameter to specify that the service component should close its connections after its read timeout period expires, when its sleep time begins. Deselect this parameter to specify that the connections should remain open until the processing thread stops.

To cancel any changes on this page, click the Cancel button.

To return to the previous step of the configuration wizard, click the Back button.

To save these settings and proceed to the next step of the configuration wizard, click the Next button.

To save these settings and proceed to the last step of the configuration wizard, click the Finish button.

Scheduling Events

This page lets you schedule events to control the running of the service component. The events are raised at the scheduled time by DBMS jobs. For an agent listener service component, you can schedule the following events:

- Start
- Refresh
- Suspend
- Resume
- Stop

For each event, the list displays the event name, date and time when the event is first scheduled to be raised, the interval in minutes at which the event is reraised, and, for a refresh event, any parameters to be refreshed. You can specify the following refreshable parameters, using the parameters' internal names, when you refresh the agent listener.

- `PROCESSOR_IN_THREAD_COUNT` - Inbound Thread Count
- `COMPONENT_LOG_LEVEL` - Log Level, specified as a numerical value
 - 1 - Statement
 - 2 - Procedure
 - 3 - Event
 - 4 - Exception

- 5 - Error
- 6 - Unexpected

To schedule events:

- If no events are currently scheduled, click the Add a Row button to add a new row to the list of events and enter the information for the event.
 - Select the event for the command you want to schedule. Oracle Workflow provides events to let you start, stop, refresh, suspend, or resume the service component.
 - Select the date when you want the event to be raised first.
 - Select the hour and minute to specify the time on the specified date when you want the event to be raised first. The hour values are in a twenty-four hour format. For example, select 00 for midnight, or 23 for 11 PM.
 - If you want to raise the event periodically, enter the time interval in minutes at which you want to raise the event. If you do not specify a repeating interval, the event is raised only once.
 - If you choose the refresh event, you can optionally enter any parameters you want to include with the event in order to refresh the agent listener configuration parameters with those values when the event is raised. Specify the parameter names and values in the following format, separating the parameters with a colon (:):
`internal_parameter_name=parameter_value`
 For example: `PROCESSOR_IN_THREAD_COUNT=1`
 If a parameter value itself contains a colon (:), then precede the colon with a backslash (\) as an escape character, as follows:
`\:`
- To schedule another event, click the Add Another Row button and enter the information for the event.
- To remove an event, select the event and click the Remove button.

To cancel any changes on this page, click the Cancel button.

To return to the previous step of the configuration wizard, click the Back button.

To save these settings and proceed to the next step of the configuration wizard, click the Next button.

To save these settings and proceed to the last step of the configuration wizard, click the Finish button.

Note: The configuration wizard verifies that an event is specified for every row in the list when you click the Next or Finish button. If you do not want to schedule another event, you should remove any empty rows before proceeding.

Review

This page lets you review the configuration parameter values that you set and the events that you scheduled for this service component.

- If you want to change any of these settings, return to the appropriate step in the configuration wizard to make your changes. To return to the previous step, click the Back button.
- To save these settings and finish the configuration, click the Finish button.

Java Agent Listeners

The Oracle Workflow Business Event System requires agent listeners to be scheduled to receive inbound event messages. An agent listener monitors a Business Event System agent for incoming messages and dequeues messages using the agent's queue handler. You should run agent listeners for your local inbound agents. Run PL/SQL agent listeners to process event subscriptions with a PL/SQL rule function in the database, and run Java agent listeners to process event subscriptions with a Java rule function in the middle tier.

When an event message is dequeued, the Event Manager begins subscription processing for the event. The Event Manager searches for and executes any active subscriptions by the local system to that event with a source type of External, and also any active subscriptions by the local system to the Any event with a source type of External. The agent listener exits after all event messages on the agent's queue have been dequeued.

The Java agent listener program is defined as a service component type in the Generic Service Component Framework. This framework helps to simplify and automate the management of background Java services.

Oracle Workflow provides several seeded Java agent listener service components to process messages on standard agents.

- Workflow Java Deferred Agent Listener - Handles messages on WF_JAVA_DEFERRED to support deferred subscription processing in the middle tier. This service component is started automatically by its container.
- Workflow Java Error Agent Listener - Handles messages on WF_JAVA_ERROR to support error handling for the Business Event System in the middle tier. This service component is started automatically by its container.

- Web Services IN Agent - Handles messages on WF_WS_JMS_IN to support inbound Web service message processing. This service component must be started manually.

You can optionally update the configuration of the Web Services IN Agent listener or delete this service component if necessary. You cannot delete the Workflow Java Deferred Agent Listener and Workflow Java Error Agent Listener or edit their names, assigned agents, correlation ID values, or containers. However, if necessary you can update other configuration parameters, schedule control events, or manually choose control commands to start, stop, suspend, resume, or refresh these Java agent listeners.

You can also optionally create additional Java agent listener service components. For example, you can configure Java agent listeners for other inbound agents that you want to use for event message propagation in the middle tier, such as custom agents. You can also configure a Java agent listener that only processes messages on a particular agent that are instances of a specific event.

In addition to the parameters in the configuration wizard, for both seeded and custom Java agent listeners, you can optionally set an internal agent listener parameter named `NAVIGATION_RESET_THRESHOLD`. This parameter lets you reset the agent listener's navigation through waiting messages to include newly arrived messages, so that new high priority messages are processed sooner. Use the `afsvcpup.sql` script to set this parameter. See: *Scheduling Listeners for Local Inbound Agents, Oracle Workflow Administrator's Guide* and *To Set Internal Agent Listener Parameters, Oracle Workflow Administrator's Guide*.

If you create custom Java agent listener service components, you can either assign them to the seeded container for agent listeners, named Workflow Agent Listener Service, or, based on the volume to be handled by the seeded container, you can also choose to create your own custom containers.

Before the seeded Java agent listener service components can run, the Workflow Agent Listener Service container which manages them must be first be started. You should ensure that this container is running. If you create your own custom containers for custom service components, ensure that those containers are running as well. Use the Service Instances page to start each container as a service instance in Generic Service Management (GSM). When the Workflow Agent Listener Service container is running, it automatically starts the Workflow Java Deferred Agent Listener and Workflow Java Error Agent Listener.

Java Agent Listener Configuration Wizard

The Java agent listener configuration wizard lets you configure a Java agent listener service component by defining general and detail attributes and scheduling control events. You can use the configuration wizard to configure a new Java agent listener service component, or to edit the configuration of an existing Java agent listener service component.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Service

Components status icon > (B) Create > (B) Continue

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Service Components status icon > (B) Edit

Define

This page lets you define general attributes for the service component. Some attributes are already set to required values and cannot be modified. You must set attributes marked with an asterisk (*) to appropriate values for your environment before you can run the service component.

- **ID** - When you edit a previously created service component, the configuration wizard displays the identifier for the service component.
- **Status** - When you edit a previously created service component, the configuration wizard displays the status of the service component.
- **Name** - The name of the service component. This name must be unique.
- **Startup Mode** - Select Automatic, Manual, or On-Demand as the startup mode for the service component.
- **Container Type** - The container type to which this service component belongs, which is always Oracle Applications Generic Service Management (Oracle Applications GSM).
- **Inbound Agent** - The Business Event System agent that you want to monitor for inbound event messages.
- **Outbound Agent** - Leave this field blank. Java agent listener service components do not use an outbound agent.
- **Correlation ID** - Optionally specify the Oracle Advanced Queuing (AQ) correlation ID of the event messages that you want the Java agent listener to process. The AQ correlation ID for an event message in the Business Event System is usually specified in the following format:

`<event name>`

By specifying a correlation ID in this attribute, you can dedicate the Java agent listener to listen only for messages that are instances of the specified event. You can also specify a partial value to listen for messages that are instances of any event whose name begins with the specified value.

Both dedicated and general Java agent listeners are compatible with each other. You can run several dedicated and general Java agent listeners for the same agent at the same time if you choose. The behavior of a general Java agent listener with a blank correlation ID depends on its agent.

- For the WF_JAVA_DEFERRED agent only, a general Java agent listener does not process messages that match a dedicated Java agent listener's correlation ID, as long as the dedicated Java agent listener has a status of `Running`, `Stopped With Error`, or `System Deactivated`. If the dedicated Java agent listener has any other status, such as `User Deactivated` or `Suspended`, then the general Java agent listener does process the messages on WF_JAVA_DEFERRED that match the dedicated Java agent listener's correlation ID.

For example, the seeded Workflow Java Deferred Agent Listener does not have any correlation ID specified, so that it can process all event messages on the WF_JAVA_DEFERRED agent that are not being handled by a dedicated agent listener. If you configure a dedicated Java agent listener for WF_JAVA_DEFERRED, then as long as that dedicated listener has a status of `Running`, `Stopped With Error`, or `System Deactivated`, the general Workflow Java Deferred Agent Listener processes all messages on the WF_JAVA_DEFERRED agent except those whose event name matches the dedicated listener's correlation ID, leaving those to be processed by the dedicated listener.

- For all other agents, a general Java agent listener can process all messages on the agent. Even if you have configured a dedicated listener for a particular agent, a message that matches the dedicated Java agent listener's correlation ID may still be processed by a general listener if that listener is the first to access the message.

For example, the seeded Workflow Java Error Agent Listener and Web Services IN Agent do not have any correlation ID specified so that they can process all event messages on their respective agents.

Note: The AQ correlation ID is different than the correlation ID contained within the WF_EVENT_T event message structure.

To cancel the configuration without saving any changes, click the Cancel button.

To save these settings and proceed to the next step of the configuration wizard, click the Next button.

Details

This page lets you define detail attributes for the service component. You must set attributes marked with an asterisk (*) to appropriate values for your environment before you can run the service component. A refresh icon identifies attributes that can be refreshed dynamically while the service component is running.

- **ID** - When you edit a previously created service component, the configuration wizard displays the identifier for the service component.

- **Status** - When you edit a previously created service component, the configuration wizard displays the status of the service component.
- **Name** - The configuration wizard displays the name defined for the service component.
- **Container** - The container to which the service component will belong. Oracle Workflow provides a container called Workflow Agent Listener Service for Java agent listener service components.
- **Maximum Idle Time** - If you selected the On-Demand startup mode for the service component, enter the maximum time in minutes that the service component can remain idle before it is stopped. An on-demand component that is stopped in this way will be restarted by its container when it is needed again to process new messages.
- **Max Error Count** - The number of consecutive errors the service component can encounter before its container stops it and changes its status to Stopped with Error. If an error is resolved and processing can continue, the error count is reset. The default value for the maximum error count is 10.
- **Inbound Thread Count** - Set the inbound processing thread count to 1 (one) or higher to enable inbound message processing with this Java agent listener. Set the inbound thread count to 0 (zero) to disable this Java agent listener. The default value is 1. If this Java agent listener receives a high volume of inbound messages, you can set the inbound thread count to a higher value to increase throughput.
- **Outbound Thread Count** - Leave this parameter set to the default value of 0 (zero). Java agent listener service components do not perform outbound message processing.
- **Log Level** - Select the level of detail for the information you want to record in the service component container log. The recommended log level, which is also the default value, is Error. Usually the log level only needs to be changed if you want to record additional detailed information for debugging purposes. You can choose the following levels:
 - 1 - Statement
 - 2 - Procedure
 - 3 - Event
 - 4 - Exception
 - 5 - Error

- 6 - Unexpected
- **Processor Read Wait Timeout** - Specify the amount of time in seconds that the service component's processing thread continues to wait, after reading the last message from its assigned queue, before timing out. If another message is received before this time expires, that message is processed and the timeout period begins again. If the timeout period expires and no more messages have been received, the service component stops reading and its sleep time begins. The default read timeout period for a Java agent listener is 10 seconds.
- **Processor Min Loop Sleep** - Specify the minimum sleep time in seconds during which the service component waits, after its read timeout period expires, before it checks its queue for messages again. The default minimum sleep time for a Java agent listener is 5 seconds.
- **Processor Max Loop Sleep** - Specify the maximum sleep time in seconds if you want to increase the sleep time between read attempts when no messages are received. If you specify a maximum sleep time that is greater than the minimum sleep time, then the service component initially waits for the minimum sleep time after it finishes reading messages from its queue. If no messages are read in subsequent attempts, then the sleep time between read attempts gradually increases until the maximum sleep time is reached. Increasing the sleep time can help enhance performance if messages are received infrequently. You can also specify 0 (zero) for this parameter to indicate that the sleep time should not be increased. In this case, the service component always waits for the minimum sleep time between read attempts. The default maximum sleep time for a Java agent listener is 60 seconds.
- **Processor Error Loop Sleep** - Specify the sleep time in seconds during which the service component waits, after an error occurs, before it attempts to begin processing again. The default error sleep time for a Java agent listener is 60 seconds.
- **Processor Close on Read Timeout** - Select this parameter to specify that the service component should close its connections after its read timeout period expires, when its sleep time begins. Deselect this parameter to specify that the connections should remain open until the processing thread stops.

To cancel any changes on this page, click the Cancel button.

To return to the previous step of the configuration wizard, click the Back button.

To save these settings and proceed to the next step of the configuration wizard, click the Next button.

To save these settings and proceed to the last step of the configuration wizard, click the Finish button.

Scheduling Events

This page lets you schedule events to control the running of the service component. The events are raised at the scheduled time by DBMS jobs. For a Java agent listener service component, you can schedule the following events:

- Start
- Refresh
- Suspend
- Resume
- Stop

For each event, the list displays the event name, date and time when the event is first scheduled to be raised, the interval in minutes at which the event is reraised, and, for a refresh event, any parameters to be refreshed. You can specify the following refreshable parameters, using the parameters' internal names, when you refresh the Java agent listener.

- `PROCESSOR_IN_THREAD_COUNT` - Inbound Thread Count
- `COMPONENT_LOG_LEVEL` - Log Level, specified as a numerical value
 - 1 - Statement
 - 2 - Procedure
 - 3 - Event
 - 4 - Exception
 - 5 - Error
 - 6 - Unexpected

To schedule events:

- If no events are currently scheduled, click the Add a Row button to add a new row to the list of events and enter the information for the event.
 - Select the event for the command you want to schedule. Oracle Workflow provides events to let you start, stop, refresh, suspend, or resume the service component.
 - Select the date when you want the event to be raised first.

- Select the hour and minute to specify the time on the specified date when you want the event to be raised first. The hour values are in a twenty-four hour format. For example, select 00 for midnight, or 23 for 11 PM.
- If you want to raise the event periodically, enter the time interval in minutes at which you want to raise the event. If you do not specify a repeating interval, the event is raised only once.
- If you choose the refresh event, you can optionally enter any parameters you want to include with the event in order to refresh the Java agent listener configuration parameters with those values when the event is raised. Specify the parameter names and values in the following format, separating the parameters with a colon (:):

`internal_parameter_name=parameter_value`

For example: `PROCESSOR_IN_THREAD_COUNT=1`

If a parameter value itself contains a colon (:), then precede the colon with a backslash (\) as an escape character, as follows:

`\:`

- To schedule another event, click the Add Another Row button and enter the information for the event.
- To remove an event, select the event and click the Remove button.

To cancel any changes on this page, click the Cancel button.

To return to the previous step of the configuration wizard, click the Back button.

To save these settings and proceed to the next step of the configuration wizard, click the Next button.

To save these settings and proceed to the last step of the configuration wizard, click the Finish button.

Note: The configuration wizard verifies that an event is specified for every row in the list when you click the Next or Finish button. If you do not want to schedule another event, you should remove any empty rows before proceeding.

Review

This page lets you review the configuration parameter values that you set and the events that you scheduled for this service component.

- If you want to change any of these settings, return to the appropriate step in the configuration wizard to make your changes. To return to the previous step, click the Back button.

- To save these settings and finish the configuration, click the Finish button.

Web Services Outbound

You can use Web services in Oracle Workflow to initiate outbound Web service requests and to accept inbound Web service requests.

When Web service messages are dequeued by the Oracle E-Business Suite, they are transmitted by the Web service outbound component.

The Web services outbound program is defined as a service component type in the Generic Service Component Framework. This framework helps to simplify and automate the management of background Java services.

Oracle Workflow provides a seeded Web services outbound component named Web Services OUT Agent to process messages on the standard WF_WS_JMS_OUT queue, which is a Business Event System agent. This service component must be started manually. You can optionally update its configuration if necessary.

You can also optionally create additional Web services outbound components. For example, you can configure a Web services outbound component that only processes messages on a particular agent or queue.

If you create custom Web services outbound components, you can either assign them to the seeded container for Web services outbound components, named Workflow Document Web Services Service, or, based on the volume to be handled by the seeded container, you can also choose to create your own custom containers.

Before the seeded Web services outbound component can run, the Workflow Document Web Services Service container which manages it must be first be started. You should ensure that this container is running. If you create your own custom containers for custom service components, ensure that those containers are running as well. Use the Service Instances page to start each container as a service instance in Generic Service Management (GSM).

Note: Inbound Web service messages are processed by a seeded service component of type Java agent listener, named Workflow Web Services In.

Web Services Outbound Configuration Wizard

The Web services outbound configuration wizard lets you configure a Web services outbound service component by defining general and detail attributes and scheduling control events. You can use the configuration wizard to configure a new Web services outbound service component, or to edit the configuration of an existing Web services outbound service component.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Service

Components status icon > (B) Create > (B) Continue

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Service Components status icon > (B) Edit

Define

This page lets you define general attributes for the service component. Some attributes are already set to required values and cannot be modified. You must set attributes marked with an asterisk (*) to appropriate values for your environment before you can run the service component.

- **ID** - When you edit a previously created service component, the configuration wizard displays the identifier for the service component.
- **Status** - When you edit a previously created service component, the configuration wizard displays the status of the service component.
- **Name** - The name of the service component. This name must be unique.
- **Startup Mode** - Select Automatic, Manual, or On-Demand as the startup mode for the service component.
- **Container Type** - The container type to which this service component belongs, which is always Oracle Applications Generic Service Management (Oracle Applications GSM).
- **Inbound Agent** - Leave this field blank. Web services outbound components do not use an inbound agent.
- **Outbound Agent** - The agent/queue that you want to monitor for outbound Web services messages.

To cancel the configuration without saving any changes, click the Cancel button.

To save these settings and proceed to the next step of the configuration wizard, click the Next button.

Details

This page lets you define detail attributes for the service component. You must set attributes marked with an asterisk (*) to appropriate values for your environment before you can run the service component. A refresh icon identifies attributes that can be refreshed dynamically while the service component is running.

- **ID** - When you edit a previously created service component, the configuration wizard displays the identifier for the service component.
- **Status** - When you edit a previously created service component, the configuration

wizard displays the status of the service component.

- **Name** - The configuration wizard displays the name defined for the service component.
- **Container** - The container to which the service component will belong. Oracle Workflow provides a container called Workflow Document Web Services Service for Web services outbound components.
- **Maximum Idle Time** - If you selected the On-Demand startup mode for the service component, enter the maximum time in minutes that the service component can remain idle before it is stopped. An on-demand component that is stopped in this way will be restarted by its container when it is needed again to process new messages.
- **Max Error Count** - The number of consecutive errors the service component can encounter before its container stops it and changes its status to Stopped with Error. If an error is resolved and processing can continue, the error count is reset. The default value for the maximum error count is 10.
- **Inbound Thread Count** - Leave this parameter set to the default value of 0 (zero). Web services outbound components do not perform inbound message processing.
- **Outbound Thread Count** - Specify the number of outbound processing threads you want to execute simultaneously with this Web services outbound component, depending on the volume of outbound messages you need to send. Specify 0 (zero) to disable this Web services outbound component. The default value is 1 (one).
- **Log Level** - Select the level of detail for the information you want to record in the service component container log. The recommended log level, which is also the default value, is Error. Usually the log level only needs to be changed if you want to record additional detailed information for debugging purposes. You can choose the following levels:
 - 1 - Statement
 - 2 - Procedure
 - 3 - Event
 - 4 - Exception
 - 5 - Error
 - 6 - Unexpected
- **Processor Read Wait Timeout** - Specify the amount of time in seconds that the

service component's processing threads continue to wait, after reading the last message from the assigned queue, before timing out. If another message is received before this time expires, that message is processed and the timeout period begins again. If the timeout period expires and no more messages have been received, the service component stops reading and its sleep time begins. The default read timeout period for a Web services outbound component is 10 seconds.

- **Processor Min Loop Sleep** - Specify the minimum sleep time in seconds during which the service component waits, after its read timeout period expires, before it checks its queue for messages again. The default minimum sleep time for a Web services outbound component is 5 seconds.
- **Processor Max Loop Sleep** - Specify the maximum sleep time in seconds if you want to increase the sleep time between read attempts when no messages are received. If you specify a maximum sleep time that is greater than the minimum sleep time, then the service component initially waits for the minimum sleep time after it finishes reading messages from its queue. If no messages are read in subsequent attempts, then the sleep time between read attempts gradually increases until the maximum sleep time is reached. Increasing the sleep time can help enhance performance if messages are received infrequently. You can also specify 0 (zero) for this parameter to indicate that the sleep time should not be increased. In this case, the service component always waits for the minimum sleep time between read attempts. The default maximum sleep time for a Web services outbound component is 60 seconds.
- **Processor Error Loop Sleep** - Specify the sleep time in seconds during which the service component waits, after an error occurs, before it attempts to begin processing again. The default error sleep time for a Web services outbound component is 60 seconds.
- **Processor Close on Read Timeout** - Select this parameter to specify that the service component should close its connections after its read timeout period expires, when its sleep time begins. Deselect this parameter to specify that the connections should remain open until the processing thread stops.

To cancel any changes on this page, click the Cancel button.

To return to the previous step of the configuration wizard, click the Back button.

To save these settings and proceed to the next step of the configuration wizard, click the Next button.

To save these settings and proceed to the last step of the configuration wizard, click the Finish button.

Scheduling Events

This page lets you schedule events to control the running of the service component. The events are raised at the scheduled time by DBMS jobs. For a Web services outbound

component, you can schedule the following events:

- Start
- Refresh
- Suspend
- Resume
- Stop

For each event, the list displays the event name, date and time when the event is first scheduled to be raised, the interval in minutes at which the event is reraised, and, for a refresh event, any parameters to be refreshed. You can specify the following refreshable parameters, using the parameters' internal names, when you refresh the Web services outbound component.

- `PROCESSOR_OUT_THREAD_COUNT` - Outbound Thread Count
- `COMPONENT_LOG_LEVEL` - Log Level, specified as a numerical value
 - 1 - Statement
 - 2 - Procedure
 - 3 - Event
 - 4 - Exception
 - 5 - Error
 - 6 - Unexpected

To schedule events:

- If no events are currently scheduled, click the Add a Row button to add a new row to the list of events and enter the information for the event.
 - Select the event for the command you want to schedule. Oracle Workflow provides events to let you start, stop, refresh, suspend, or resume the service component.
 - Select the date when you want the event to be raised first.
 - Select the hour and minute to specify the time on the specified date when you want the event to be raised first. The hour values are in a twenty-four hour format. For example, select 00 for midnight, or 23 for 11 PM.

- If you want to raise the event periodically, enter the time interval in minutes at which you want to raise the event. If you do not specify a repeating interval, the event is raised only once.
- If you choose the refresh event, you can optionally enter any parameters you want to include with the event in order to refresh the Web services outbound configuration parameters with those values when the event is raised. Specify the parameter names and values in the following format, separating the parameters with a colon (:):

`internal_parameter_name=parameter_value`

For example: `PROCESSOR_OUT_THREAD_COUNT=3`

If a parameter value itself contains a colon (:), then precede the colon with a backslash (\) as an escape character, as follows:

`\:`

- To schedule another event, click the Add Another Row button and enter the information for the event.
- To remove an event, select the event and click the Remove button.

To cancel any changes on this page, click the Cancel button.

To return to the previous step of the configuration wizard, click the Back button.

To save these settings and proceed to the next step of the configuration wizard, click the Next button.

To save these settings and proceed to the last step of the configuration wizard, click the Finish button.

Note: The configuration wizard verifies that an event is specified for every row in the list when you click the Next or Finish button. If you do not want to schedule another event, you should remove any empty rows before proceeding.

Review

This page lets you review the configuration parameter values that you set and the events that you scheduled for this service component.

- If you want to change any of these settings, return to the appropriate step in the configuration wizard to make your changes. To return to the previous step, click the Back button.
- To save these settings and finish the configuration, click the Finish button.

Background Engines

Background engine processes serve three purposes in Oracle Workflow: to handle activities deferred by the Workflow Engine, to handle timed out notification activities, and to handle stuck processes.

When the Workflow Engine initiates and performs a process, it completes all necessary activities before continuing to the next eligible activity. In some cases, an activity can require a large amount of processing resource or time to complete. Oracle Workflow lets you manage the load on the Workflow Engine by setting up supplemental engines to run these costly activities as background tasks. In these cases, the costly activity is deferred by the Workflow Engine and run later by a background engine. The main Workflow Engine can then continue to the next available activity, which may occur on some other parallel branch of the process.

A background engine must also be set up to handle timed out notification activities. When the Workflow Engine comes across a notification activity that requires a response, it calls the Notification System to send the notification to the appropriate performer, and then sets the notification activity to a status of 'NOTIFIED' until the performer completes the notification activity. Meanwhile, a background engine set up to handle timed out activities periodically checks for 'NOTIFIED' activities and whether these activities have time out values specified. If a 'NOTIFIED' activity does have a time out value, and the current date and time exceeds that time out value, the background engine marks that activity as timed out and calls the Workflow Engine. The Workflow Engine then resumes by trying to execute a <timeout> transition activity.

Additionally, a background engine must be set up to handle stuck processes. A process is identified as stuck when it has a status of ACTIVE, but cannot progress any further. For example, a process could become stuck in the following situations:

- A thread within a process leads to an activity that is not defined as an End activity but has no other activity modeled after it, and no other activity is active.
- A process with only one thread loops back, but the pivot activity of the loop has the On Revisit property set to Ignore.
- An activity returns a result for which no eligible transition exists. For instance, if the function for a function activity returns an unexpected result value, and no default transition is modeled after that activity, the process cannot continue.

The background engine sets the status of a stuck process to ERROR:#STUCK and executes the error process defined for it.

You can define and start up as many background engines as you like to check for deferred and timed out activities.

You run a background engine by submitting the Workflow Background Process concurrent program (FNDWFBG). Background engines can be restricted to handle activities associated with specific item types, and within specific cost ranges. A

background engine runs until it completes all eligible activities at the time it was initiated. Generally, you should set the background engine up to run periodically.

Ensure that you have at least one background engine that can check for timed out activities, one that can process deferred activities, and one that can handle stuck processes. At a minimum, you need to set up one background engine that can handle both timed out and deferred activities as well as stuck processes. Generally, you should run a separate background engine to check for stuck processes at less frequent intervals than the background engine that you run for deferred activities, normally not more often than once a day. Run the background engine to check for stuck processes when the load on the system is low.

Note: If you implement workflow RAC affinity, then you should also run background engines using the Workflow Background Process for RAC concurrent program (FNDWFBGRAC). This program runs background engines that each process only the RAC-enabled workflows that were launched in a specific RAC instance. Running background engines with RAC affinity provides faster access to the workflow runtime data and helps avoid contention. However, you cannot submit the Workflow Background Process for RAC concurrent program through Oracle Workflow Manager. You must submit this program through the standard request submission UI. See: *Setting Up Workflow RAC Affinity, Oracle Workflow Administrator's Guide* and *Setting Up Background Workflow Engines, Oracle Workflow Administrator's Guide*.

You should run the Workflow Background Process for RAC program for deferred activities, timed out activities, and stuck processes as needed depending on the requirements of your RAC-enabled workflows. If the RAC-enabled workflows run on a particular schedule, then you should run the Workflow Background Process for RAC program on a corresponding schedule. You should also continue running the Workflow Background Process program to handle workflows that are not RAC-enabled. To ensure that RAC-enabled workflows are processed using RAC affinity, schedule the Workflow Background Process for RAC program to run before the Workflow Background Process program, particularly if you run the Workflow Background Process program without specifying an item type.

Running Background Engines

To run a background engine, submit the Workflow Background Process concurrent program (FNDWFBG). When you start a new background engine, you can restrict the engine to handle activities associated with specific item types, and within specific cost ranges. You can submit the Workflow Background Process concurrent program several times to schedule different background engines to run at different times.

- To submit a request for the Workflow Background Process concurrent program, choose Background Engines from the Submit Request For pull-down menu in the Workflow System status page and click the Go button.
- To view Workflow Background Process concurrent requests, click the Background Engines status icon in the Workflow System status page.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go

Parameters

When you submit the Workflow Background Process concurrent program, specify the following parameters.

- **Item Type** - Specify an item type to restrict this engine to activities associated with that item type. If you do not specify an item type, the engine processes any activity regardless of its item type.

Note: If you implemented workflow RAC affinity, then the following conditions apply.

- To obtain the performance benefits of workflow RAC affinity, you should run the Workflow Background Process for RAC program for the item types that include RAC-enabled workflow processes.
- If any item types include both RAC-enabled and non-RAC workflow processes, then you should also run the normal Workflow Background Process program for those item types in order to handle the non-RAC workflow processes. Note that in this case the Workflow Background Process program executes eligible activities from all workflow processes in the specified item type, whether the processes are non-RAC or RAC-enabled, without respect to RAC affinity.
- If you run the Workflow Background Process program without specifying an item type, then it executes eligible activities from all workflow processes in all item types, whether the processes are non-RAC or RAC-enabled, without respect to RAC affinity.
- Consequently, to ensure that RAC-enabled workflows are processed using RAC affinity, schedule the Workflow Background Process for RAC program to run before the Workflow Background Process program, particularly if you run the Workflow Background Process program without specifying an item type.

- **Minimum Threshold** - Specify the minimum cost that an activity must have for this background engine to execute it, in hundredths of a second.
- **Maximum Threshold** - Specify the maximum cost that an activity can have for this background engine to execute it, in hundredths of a second. By using Minimum Threshold and Maximum Threshold you can create multiple background engines to handle very specific types of activities. The default values for these arguments are null so that the background engine runs activities regardless of cost.
- **Process Deferred** - Specify whether this background engine checks for deferred activities. Setting this parameter to Yes allows the engine to check for deferred activities.
- **Process Timeout** - Specify whether this background engine checks for activities that have timed out. Setting this parameter to Yes allows the engine to check for timed out activities.
- **Process Stuck** - Specify whether this background engine checks for stuck processes. Setting this parameter to Yes allows the engine to check for stuck processes.

Note: Make sure you have a least one background engine that can check for timed out activities, one that can process deferred activities, and one that can handle stuck processes. At a minimum, you need to set up one background engine that can handle both timed out and deferred activities as well as stuck processes.

Viewing Concurrent Requests

When you view the Workflow Background Process concurrent requests, the Background Engines page shows standard request detail information for these requests. For each request, the list displays the request ID, program short name, description, application short name, phase, status, requester, duration, wait time, and submission date. Click a column heading to sort the list by that column.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Background Engines status icon

- To show the details for a request if they are hidden, click the Show link in the Details column. Oracle Applications Manager displays details about the request depending on the status of the request. You can also perform actions, such as placing a hold on a request, canceling a request, viewing diagnostic information, viewing manager details, viewing logs, or viewing request output, by clicking the corresponding button. The actions that are available depend on the status of the request.
- To hide the details for a request if they are shown, click the Hide link in the Details

column.

- To search for concurrent requests with different criteria, click the New Search button or click one of the Quick Search links.
- To modify the search criteria from this search, click the Modify Search button.
- To add the information from this page to your support cart, click the Add to Support Cart button.

Purging Workflow Data

The Oracle Applications Manager console helps you easily maintain the Oracle Workflow and Oracle XML Gateway database tables. Oracle Workflow and Oracle XML Gateway access several tables that can grow quite large with obsolete workflow information that is stored for all completed workflow processes, as well as obsolete information for XML transactions. The size of these tables and indexes can adversely affect performance. These tables should be purged on a regular basis, using the Purge Obsolete Workflow Runtime Data concurrent program.

This program purges obsolete runtime information associated with work items, including status information, any associated notifications, and, if the ECX: Purge ECX data with WF profile option is set to `Y`, any associated Oracle XML Gateway transactions. By default, it also purges obsolete design information, such as activities that are no longer in use and expired ad hoc users and roles, and obsolete runtime information not associated with work items, such as notifications that were not handled through a workflow process and, if the ECX: Purge ECX data with WF profile option is set to `Y`, Oracle XML Gateway transactions that were not handled through a workflow process. You can optionally choose to purge only core runtime information associated with work items for performance gain during periods of high activity, and purge all obsolete information as part of your routine maintenance during periods of low activity.

Note: This program does not delete ad hoc users or roles whose expiration date is null. To ensure that ad hoc users and roles are purged in a timely fashion after they are no longer needed, estimate how long they should be active and specify an appropriate expiration date when you call `WF_DIRECTORY.CreateAdHocUser()`, `WF_DIRECTORY.CreateAdHocRole()`, or `WF_DIRECTORY.CreateAdHocRole2()` to create them.

To preserve electronic signature evidence for future reference, this program by default does not delete any notifications that required signatures or their associated signature information. If you do not need to maintain signature evidence, you can choose to delete signature-related information as well.

Note: You can also use the Purge Obsolete ECX Data concurrent program to purge Oracle XML Gateway transactions according to Oracle XML Gateway-specific parameters. For information about this program and about the ECX: Purge ECX data with WF profile option, see: Purge Obsolete ECX Data Concurrent Program, *Oracle XML Gateway User's Guide* and Purge Obsolete Workflow Runtime Data Concurrent Program, *Oracle XML Gateway User's Guide*.

Workflow Purge

The Workflow Purge page shows summary information about the next scheduled and last completed purge requests and about completed work items.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Purge status icon

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Related Links > Throughput > Work Items

To view work items with a different status, choose the status you want from the View pull-down menu and click the Go button. You can view items with the following statuses:

- Completed Work Items
- Active Work Items
- Deferred Work Items
- Suspended Work Items
- Errored Work Items

Requests Summary

This region displays summary information about the next scheduled and last completed Purge Obsolete Workflow Runtime Data concurrent requests.

- To show information in this region if it is hidden, click the Show link.
- To hide information in this region if it is shown, click the Hide link.

Next Scheduled

For the next scheduled Purge Obsolete Workflow Runtime Data concurrent request, Oracle Workflow Manager displays the request ID, requestor, status, requested start time, wait time, and parameters.

Last Completed

For the last completed Purge Obsolete Workflow Runtime Data concurrent request, Oracle Workflow Manager displays the request ID, requestor, status, completed time, duration, and parameters.

To view the log file for the request, click the Request Log link.

Completed Work Items

This region displays the distribution of completed work items across different item types.

- To show information in this region if it is hidden, click the Show link
- To hide information in this region if it is shown, click the Hide link.
- This region displays the date and time when the work item statistics were last updated. To refresh this information, click the refresh icon. See: Gathering Oracle Workflow Statistics, page 15-2.

For each work item type in the Completed Work Items list, Oracle Workflow Manager displays the work item type name, the persistence type, the retention period in days, the number of completed work items of that type, and the number of items of that type that are available for purging. Click any column heading to sort the list by that column.

- To filter the item types displayed in the list, select an item type property and an operator from the Filter pull-down menus, enter a filter value in the text field, and click the Go button. You can filter by the following properties:
 - Work item type display name
 - Work item type internal name
 - Persistence type
 - Retention period
 - Number of completed work items of this type
 - Number of items of this type available for purging
- To view details for work items of a particular item type, either click the item type link in the Work Item Type column, or select the item type and click the View Details button.

Submitting the Purge Program

You perform purging by submitting the Purge Obsolete Workflow Runtime Data

concurrent program (FNDWFPR). You can enter restrictions to specify the data that you want to purge.

- To submit a request for the Purge Obsolete Workflow Runtime Data concurrent program, either click the Purge button in the Completed Work Items region of the Workflow Purge page, or choose Purge from the Submit Request For pull-down menu in the Workflow System status page and click the Go button.
- To view Purge Obsolete Workflow Runtime Data concurrent requests, click the View Purge Requests button in the Completed Work Items region of the Workflow Purge page.

Parameters

When you submit the Purge Obsolete Workflow Runtime Data concurrent program, specify the following parameters.

- **Item Type** - Specify the item type to purge. Leave this field blank to purge the runtime data for all item types.
- **Item Key** - Specify the item key to purge. The item key is a unique identifier for an item within an item type. Leave this field blank to purge the runtime data for all items of the specified item type.
- **Age** - Specify the minimum age of data to purge, in days, if you are purging items with a Temporary persistence type. The default is 0 days.
- **Persistence Type** - Specify the persistence type of the data you want to purge, either Permanent or Temporary. The default is Temporary.
- **Core Workflow Only** - Enter 'Y' to purge only obsolete runtime data associated with work items, or 'N' to purge all obsolete runtime data as well obsolete design data. The default is 'N'.
- **Commit Frequency** - Enter the number of records to purge before the program commits data. To reduce rollback size and improve performance, set this parameter to commit data after a smaller number of records. The default is 500 records.

Note: After performing a commit, the program resumes purging work items with the next subsequent begin date. In some cases, if additional items have the same begin date as the last item that was purged before a commit, the program may not purge all eligible items. To purge these remaining work items, simply rerun the program.

- **Signed Notifications** - Enter 'N' to preserve signature evidence, including notifications that required electronic signatures and their associated signature

information. Enter 'Y' to purge signature-related information. The default is 'N'.

Viewing Concurrent Requests

When you view the Purge Obsolete Workflow Runtime Data concurrent requests, the Workflow Purge page shows standard request detail information for these requests. For each request, the list displays the request ID, program short name, description, application short name, phase, status, requestor, duration, wait time, and submission date. Click a column heading to sort the list by that column.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Purge status icon > (B) View Purge Requests

- To show the details for a request if they are hidden, click the Show link in the Details column. Oracle Applications Manager displays details about the request depending on the status of the request. You can also perform actions, such as placing a hold on a request, canceling a request, viewing diagnostic information, viewing manager details, viewing logs, or viewing request output, by clicking the corresponding button. The actions that are available depend on the status of the request.
- To hide the details for a request if they are shown, click the Hide link in the Details column.
- To search for concurrent requests with different criteria, click the New Search button or click one of the Quick Search links.
- To modify the search criteria from this search, click the Modify Search button.
- To add the information from this page to your support cart, click the Add to Support Cart button.

Completed Work Item Details

This page shows details about completed work items of a particular item type.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Purge status icon > (B) View Details

To view work items with a different status, choose the status you want from the View pull-down menu and click the Go button. You can view items with the following statuses:

- Completed Work Items
- Active Work Items
- Deferred Work Items

- Suspended Work Items
- Errored Work Items

Completed Work Items Stage Summary

This region displays the distribution of completed work items that ended at various activity stages within the workflow process. For each activity stage, the list displays the activity internal name and result, and the number of completed work items that ended at that stage. Click any column heading to sort the list by that column.

- By default, the list shows completed work items that ended within the last 30 days. To view completed work items that ended within a different period, enter a number of days in the Filter: End Date Within Last _ Days option and click the Go button.
- To view details about the work items that ended at a particular activity stage, either click the activity stage link in the Work Item Activity Stage column, or select the activity stage and click the View Details button.

Completed Work Item Activity Details

This page shows details about completed work items that ended at a particular activity stage within a particular item type.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Purge status icon > (B) View Details > (B) View Details

To view work items with a different status, choose the status you want from the View pull-down menu and click the Go button. You can view items with the following statuses:

- Completed Work Items
- Active Work Items
- Deferred Work Items
- Suspended Work Items
- Errored Work Items

Oracle Workflow Manager displays a list of all completed work items of the selected item type that ended at the selected activity stage. By default, the list shows completed work items that ended within the last 30 days. For each work item, the list displays the internal name of the activity at which the work item ended, the activity start date, end date, user assigned to perform the activity, and item key. Click any column heading to sort the list by that column.

- To filter the work items displayed in the list, select an activity property from the

Filter pull-down menu, enter a filter value in the text field, and click the Go button. You can filter by the following properties:

- Internal name of the activity at which the work item ended
 - Start date within a specified number of days
 - End date within a specified number of days
 - User assigned to perform the activity
 - Item key of the work item
- To launch the Workflow Monitor for a work item, select the work item and click the Launch Workflow Monitor button.

Note: If you perform an action in the Workflow Monitor that changes the status of the work item, then you must refresh your Oracle Workflow Manager web page in order to see the updated information.

Workflow Control Queue Cleanup

Oracle Workflow contains a standard Business Event System agent named WF_CONTROL, which is associated with a standard queue that is also named WF_CONTROL. This queue has a payload type of JMS Text message. The WF_CONTROL agent is used for internal processing only, and is not meant for customer use. You should not place custom event messages on this queue.

The Generic Service Component Framework uses WF_CONTROL to handle control events for containers and service components, such as notification mailer or agent listener service components. WF_CONTROL is also used for other Oracle E-Business Suite internal processing.

You do not need to schedule propagation for the WF_CONTROL agent, because the middle tier processes that use WF_CONTROL dequeue messages directly from its queue. However, the subscribers to the WF_CONTROL queue need to be cleaned up periodically. A concurrent program named Workflow Control Queue Cleanup is automatically scheduled to perform this cleanup for you.

When a middle tier process for Oracle E-Business Suite starts up, it creates a JMS subscriber to the queue. Then, when an event message is placed on the queue, a copy of the event message is created for each subscriber to the queue. If a middle tier process dies, however, the corresponding subscriber remains in the database. For more efficient processing, you should ensure that WF_CONTROL is periodically cleaned up by removing the subscribers for any middle tier processes that are no longer active. The Workflow Control Queue Cleanup concurrent program sends an event named oracle.apps.wf.bes.control.ping to check the status of each subscriber to the

WF_CONTROL queue. If the corresponding middle tier process is still alive, it sends back a response. The next time the cleanup program runs, it checks whether responses have been received for each ping event sent during the previous run. If no response was received from a particular subscriber, that subscriber is removed.

The recommended frequency for performing cleanup is every twelve hours. In order to allow enough time for subscribers to respond to the ping event, the minimum wait time between two cleanup runs is thirty minutes. If you run the procedure again less than thirty minutes after the previous run, it will not perform any processing.

Running Workflow Control Queue Cleanup

You perform Workflow control queue cleanup by submitting the Workflow Control Queue Cleanup concurrent program (FNDWFBES_CONTROL_QUEUE_CLEANUP). This program does not require any parameters. This concurrent program is scheduled to run every twelve hours by default, which is the recommended frequency for performing cleanup. You can optionally submit this program with a different schedule if you want to perform cleanup at a different frequency.

- To submit a request for the Workflow Control Queue Cleanup concurrent program, choose Control Queue Cleanup from the Submit Request For pull-down menu in the Workflow System status page and click the Go button.
- To view Workflow Control Queue Cleanup concurrent requests, click the Control Queue Cleanup status icon in the Workflow System status page.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go

Viewing Concurrent Requests

When you view the Workflow Control Queue Cleanup concurrent requests, the Control Queue Cleanup page shows standard request detail information for these requests. For each request, the list displays the request ID, program short name, description, application short name, phase, status, requester, duration, wait time, and submission date. Click a column heading to sort the list by that column.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Control Queue Cleanup status icon

- To show the details for a request if they are hidden, click the Show link in the Details column. Oracle Applications Manager displays details about the request depending on the status of the request. You can also perform actions, such as placing a hold on a request, canceling a request, viewing diagnostic information, viewing manager details, viewing logs, or viewing request output, by clicking the corresponding button. The actions that are available depend on the status of the request.
- To hide the details for a request if they are shown, click the Hide link in the Details

column.

- To search for concurrent requests with different criteria, click the New Search button or click one of the Quick Search links.
- To modify the search criteria from this search, click the Modify Search button.
- To add the information from this page to your support cart, click the Add to Support Cart button.

Active Work Items

The Active Work Items page shows the distribution of active work items across different item types. All work items that do not have an end date are counted as Active work items, including deferred, suspended, and errored work items as well as running work items.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Workflow Metrics > Work Items > Active

To view work items with a different status, choose the status you want from the View pull-down menu and click the Go button. You can view items with the following statuses:

- Completed Work Items
- Active Work Items
- Deferred Work Items
- Suspended Work Items
- Errored Work Items

The page displays the date and time when the work item statistics were last updated. To refresh this information, click the refresh icon. See: Gathering Oracle Workflow Statistics, page 15-2.

For each work item type, the Active Work Items page displays the work item type name and the number of active work items of that type. Click any column heading to sort the list by that column.

- To filter the item types displayed in the list, select an item type property and an operator from the Filter pull-down menus, enter a filter value in the text field, and click the Go button. You can filter by the following properties:
 - Work item type display name
 - Work item type internal name

- Number of active work items of this type

To view details about active work item activities within a particular item type, either click the item type link in the Work Item Type column, or select the item type and click the View Details button.

Active Work Item Activities

This page shows details about active work item activities within a particular item type. Active work item activities include only activities with a status of Active, Waiting, or Notified.

Note: Only activities with a status of Active, Waiting, or Notified are included in this page. Activities with a status of Deferred, Suspended, or Error are not included in this page, although the work items to which they belong are counted as Active work items. You can use the View pull-down menu to view details for activities with a status of Deferred, Suspended, or Error.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Workflow Metrics > Work Items > Active > (B) View Details

To view work items with a different status, choose the status you want from the View pull-down menu and click the Go button. You can view items with the following statuses:

- Completed Work Items
- Active Work Items
- Deferred Work Items
- Suspended Work Items
- Errored Work Items

Active Work Items Stage Summary

This region displays the distribution of active work items that are currently at various activity stages within the workflow process, if the activity has a status of Active, Waiting, or Notified. For each activity stage, the list displays the activity internal name and the number of active work items at that stage. Click any column heading to sort the list by that column.

- By default, the list shows active work items that started within the last 30 days. To view active work items that started within a different period, enter a number of days in the Filter: Start Date Within Last _ Days option and click the Go button.

- To view details about the work items at a particular activity stage, either click the activity stage link in the Work Item Activity Stage column, or select the activity stage and click the View Details button.

Active Work Item Activity Details

This page shows details about active work item activities of a particular activity stage within a particular item type. Active work item activities include only activities with a status of Active, Waiting, or Notified.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Workflow Metrics > Work Items > Active > (B) View Details > (B) View Details

To view work items with a different status, choose the status you want from the View pull-down menu and click the Go button. You can view items with the following statuses:

- Completed Work Items
- Active Work Items
- Deferred Work Items
- Suspended Work Items
- Errored Work Items

Oracle Workflow Manager displays a list of all active activities of the selected stage for work items of the selected item type. Active work item activities include only activities with a status of Active, Waiting, or Notified. By default, the list shows active work items that started within the last 30 days. For each activity, the list displays the activity internal name, start date, due date, user assigned to perform the activity, and item key of the work item. Click any column heading to sort the list by that column.

- To filter the work items displayed in the list, select an activity property from the Filter pull-down menu, enter a filter value in the text field, and click the Go button. You can filter by the following properties:
 - Internal name of the active activity
 - Start date within a specified number of days
 - Due date within a specified number of days
 - User assigned to perform the activity
 - Item key of the work item
- To abort all work items in the list, click the Abort All button. If you have filtered the

list, only the work items currently displayed in the list are aborted.

- To suspend all activities in the list, click the Suspend All button. If you have filtered the list, only the work items currently displayed in the list are suspended.
- To abort a single work item, select the activity you want and click the Abort button.
- To suspend a single activity, select the activity you want and click the Suspend button.
- To launch the Workflow Monitor for a work item, select the activity you want and click the Launch Workflow Monitor button.

Note: If you perform an action in the Workflow Monitor that changes the status of the work item, such as aborting the work item, then you must refresh your Oracle Workflow Manager web page in order to see the updated information.

Deferred Work Items

The Deferred Work Items page shows the distribution of deferred work items across different item types. An abnormal number of activities with a deferred status may indicate that there are not enough background engines available.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Workflow Metrics > Work Items > Deferred

To view work items with a different status, choose the status you want from the View pull-down menu and click the Go button. You can view items with the following statuses:

- Completed Work Items
- Active Work Items
- Deferred Work Items
- Suspended Work Items
- Errored Work Items

The page displays the date and time when the work item statistics were last updated. To refresh this information, click the refresh icon. See: Gathering Oracle Workflow Statistics, page 15-2.

For each work item type, the Deferred Work Items page displays the work item type name and the number of deferred work items of that type. Click any column heading to sort the list by that column.

- To filter the item types displayed in the list, select an item type property and an operator from the Filter pull-down menus, enter a filter value in the text field, and click the Go button. You can filter by the following properties:
 - Work item type display name
 - Work item type internal name
 - Number of deferred work items of this type
- To view details for work items of a particular item type, either click the item type link in the Work Item Type column, or select the item type and click the View Details button.

Deferred Work Item Details

This page shows details about deferred work items of a particular item type.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Workflow Metrics > Work Items > Deferred > (B) View Details

To view work items with a different status, choose the status you want from the View pull-down menu and click the Go button. You can view items with the following statuses:

- Completed Work Items
- Active Work Items
- Deferred Work Items
- Suspended Work Items
- Errored Work Items

Deferred Work Items Stage Summary

This region displays the distribution of deferred work items that are currently at various activity stages within the workflow process. For each activity stage, the list displays the activity internal name and the number of deferred work items at that stage. Click any column heading to sort the list by that column.

- By default, the list shows active work items that started within the last 30 days. To view deferred work items that started within a different period, enter a number of days in the Filter: Start Date Within Last _ Days option and click the Go button.
- To view details about the work items at a particular activity stage, either click the activity stage link in the Work Item Activity Stage column, or select the activity

stage and click the View Details button.

Deferred Work Item Activity Details

This page shows details about deferred work items that are currently at a particular activity stage within a particular item type.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Workflow Metrics > Work Items > Deferred > (B) View Details > (B) View Details

To view work items with a different status, choose the status you want from the View pull-down menu and click the Go button. You can view items with the following statuses:

- Completed Work Items
- Active Work Items
- Deferred Work Items
- Suspended Work Items
- Errored Work Items

Oracle Workflow Manager displays a list of all deferred activities of the selected stage for work items of the selected item type. By default, the list shows deferred work items that started within the last 30 days. For each activity, the list displays the activity internal name, start date, due date, user assigned to perform the activity, and item key of the work item. Click any column heading to sort the list by that column.

- To filter the work items displayed in the list, select an activity property from the Filter pull-down menu, enter a filter value in the text field, and click the Go button. You can filter by the following properties:
 - Internal name of the deferred activity
 - Start date within a specified number of days
 - Due date within a specified number of days
 - User assigned to perform the activity
 - Item key of the work item
- To abort all work items in the list, click the Abort All button. If you have filtered the list, only the work items currently displayed in the list are aborted.
- To suspend all activities in the list, click the Suspend All button. If you have filtered the list, only the work items currently displayed in the list are suspended.

- To abort a single work item, select the activity you want and click the Abort button.
- To suspend a single activity, select the activity you want and click the Suspend button.
- To launch the Workflow Monitor for a work item, select the activity you want and click the Launch Workflow Monitor button.

Note: If you perform an action in the Workflow Monitor that changes the status of the work item, such as aborting the work item, then you must refresh your Oracle Workflow Manager web page in order to see the updated information.

Suspended Work Items

The Suspended Work Items page shows the distribution of suspended work items across different item types.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Workflow Metrics > Work Items > Suspended

To view work items with a different status, choose the status you want from the View pull-down menu and click the Go button. You can view items with the following statuses:

- Completed Work Items
- Active Work Items
- Deferred Work Items
- Suspended Work Items
- Errored Work Items

The page displays the date and time when the work item statistics were last updated. To refresh this information, click the refresh icon. See: Gathering Oracle Workflow Statistics, page 15-2.

For each work item type, the Suspended Work Items page displays the work item type name and the number of suspended work items of that type. Click any column heading to sort the list by that column.

- To filter the item types displayed in the list, select an item type property and an operator from the Filter pull-down menus, enter a filter value in the text field, and click the Go button. You can filter by the following properties:
 - Work item type display name

- Work item type internal name
- Number of suspended work items of this type
- To view details for an item type, either click the item type link in the Work Item Type column, or select the item type and click the View Details button.

Suspended Work Item Details

This page shows details about all suspended work items of a particular item type.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Workflow Metrics > Work Items > Suspended > (B) View Details

To view work items with a different status, choose the status you want from the View pull-down menu and click the Go button. You can view items with the following statuses:

- Completed Work Items
- Active Work Items
- Deferred Work Items
- Suspended Work Items
- Errored Work Items

Suspended Work Items Stage Summary

This region displays the distribution of suspended work items that are currently at various activity stages within the workflow process. For each activity stage, the list displays the activity internal name and the number of suspended work items at that stage. Click any column heading to sort the list by that column.

- To view suspended work items that started within a specific period, enter a number of days in the Filter: Start Date Within Last _ Days option and click the Go button.
- To view details about the work items at a particular activity stage, either click the activity stage link in the Work Item Activity Stage column, or select the activity stage and click the View Details button.

Suspended Work Item Activity Details

This page shows details about all suspended work items at a particular activity stage within a particular item type.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go >

Workflow Metrics > Work Items > Suspended > (B) View Details > (B) View Details

To view work items with a different status, choose the status you want from the View pull-down menu and click the Go button. You can view items with the following statuses:

- Completed Work Items
- Active Work Items
- Deferred Work Items
- Suspended Work Items
- Errored Work Items

Oracle Workflow Manager displays a list of all suspended activities of the selected stage for work items of the selected item type. For each activity, the list displays the activity internal name, start date, due date, user assigned to perform the activity, and item key of the work item. Click any column heading to sort the list by that column.

- To filter the work items displayed in the list, select an activity property from the Filter pull-down menu, enter a filter value in the text field, and click the Go button. You can filter by the following properties:
 - Internal name of the suspended activity
 - Start date within a specified number of days
 - Due date within a specified number of days
 - User assigned to perform the activity
 - Item key of the work item
- To abort all work items in the list, click the Abort All button. If you have filtered the list, only the work items currently displayed in the list are aborted.
- To resume all activities in the list, click the Resume All button. If you have filtered the list, only the work items currently displayed in the list are resumed.
- To abort a single work item, select the activity you want and click the Abort button.
- To resume a single activity, select the activity you want and click the Resume button.
- To launch the Workflow Monitor for a work item, select the activity you want and click the Launch Workflow Monitor button.

Note: If you perform an action in the Workflow Monitor that changes the status of the work item, such as aborting the work item, then you must refresh your Oracle Workflow Manager web page in order to see the updated information.

Errored Work Items

The Errored Work Items page shows the distribution of errored work items across different item types.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Workflow Metrics > Work Items > Error

To view work items with a different status, choose the status you want from the View pull-down menu and click the Go button. You can view items with the following statuses:

- Completed Work Items
- Active Work Items
- Deferred Work Items
- Suspended Work Items
- Errored Work Items

The page displays the date and time when the work item statistics were last updated. To refresh this information, click the refresh icon. See: Gathering Oracle Workflow Statistics, page 15-2.

For each work item type, the Errored Work Items page displays the work item type name and the number of errored work items of that type. Click any column heading to sort the list by that column.

- To filter the item types displayed in the list, select an item type property and an operator from the Filter pull-down menus, enter a filter value in the text field, and click the Go button. You can filter by the following properties:
 - Work item type display name
 - Work item type internal name
 - Number of errored work items of this type
- To view details for an item type, either click the item type link in the Work Item Type column, or select the item type and click the View Details button.

Errored Work Item Details

This page shows details about all errored work items of a particular item type.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Workflow Metrics > Work Items > Error > (B) View Details

To view work items with a different status, choose the status you want from the View pull-down menu and click the Go button. You can view items with the following statuses:

- Completed Work Items
- Active Work Items
- Deferred Work Items
- Suspended Work Items
- Errored Work Items

Errored Work Items Stage Summary

This region displays the distribution of errored work items that are currently at various activity stages within the workflow process. For each activity stage, the list displays the activity internal name and the number of errored work items at that stage. Click any column heading to sort the list by that column.

- To view errored work items that started within a specific period, enter a number of days in the Filter: Start Date Within Last _ Days option and click the Go button.
- To view details about the work items at a particular activity stage, either click the activity stage link in the Work Item Activity Stage column, or select the activity stage and click the View Details button.

Errored Work Item Activity Details

This page shows details about all errored work items at a particular activity stage within a particular item type.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Workflow Metrics > Work Items > Error > (B) View Details > (B) View Details

To view work items with a different status, choose the status you want from the View pull-down menu and click the Go button. You can view items with the following statuses:

- Completed Work Items

- Active Work Items
- Deferred Work Items
- Suspended Work Items
- Errored Work Items

Oracle Workflow Manager displays a list of all errored activities of the selected stage for work items of the selected item type. For each activity, the list displays the activity internal name, start date, due date, user assigned to perform the activity, and item key of the work item. Click any column heading to sort the list by that column.

- To filter the work items displayed in the list, select an activity property from the Filter pull-down menu, enter a filter value in the text field, and click the Go button. You can filter by the following properties:
 - Internal name of the errored activity
 - Start date within a specified number of days
 - Due date within a specified number of days
 - User assigned to perform the activity
 - Item key of the work item
- To abort all work items in the list, click the Abort All button. If you have filtered the list, only the work items currently displayed in the list are aborted.
- To retry all activities in the list, click the Retry All button. If you have filtered the list, only the work items currently displayed in the list are retried.
- To abort a single work item, select the activity you want and click the Abort button.
- To retry a single activity, select the activity you want and click the Retry button.
- To launch the Workflow Monitor for a work item, select the activity you want and click the Launch Workflow Monitor button.

Note: If you perform an action in the Workflow Monitor that changes the status of the work item, such as aborting the work item, then you must refresh your Oracle Workflow Manager web page in order to see the updated information.

Note: You can also use the Retry Errored Workflow Activities

concurrent program to retry multiple errored activities for a particular item type at once. See: [Retry Errored Workflow Activities \(FNDWFRET\)](#), *Oracle Workflow Administrator's Guide*.

Agents

The Agent Activity page shows the distribution of event messages with different statuses on different Business Event System agents in your instance of Oracle Workflow.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Workflow Metrics > Agent Activity

The page displays the date and time when the agent activity statistics were last updated. To refresh this information, click the refresh icon. See: [Gathering Oracle Workflow Statistics](#), page 15-2.

For each agent, the list displays the agent name as well as the number of event messages on that agent with the following statuses: Ready, Waiting, Processed, Expired, and Undeliverable. Click any column heading to sort the list by that column.

- To view queue details for an agent, click the agent link in the Agent column.
- To view details about the messages being held on an agent, select the agent and click the Search Agent Entry Details button.

Note: The Agent Activity page displays event messages on the WF_ERROR agent according to their explicitly assigned status on the WF_ERROR queue, unlike the Agent Activity graph in the Workflow System Status page which summarizes all messages on the WF_ERROR agent in an Error status.

If an inbound agent has an abnormally large number of messages with a status of Ready, you may need to check the status of the agent listener processing message for that agent, or create a new agent listener service component for that agent. Similarly, if an outbound agent has an abnormally large number of messages with a status of Ready, you may need to check the status of the propagation schedule for that agent's queue, or schedule propagation if necessary.

Agent Queue Details

The Agent Details page displays the following details for the queue associated with an agent:

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Workflow Metrics > Agent Activity > agent link

- Owner - The owner of the queue.
- Name - The name of the queue.
- Queue Table - The name of the table in which the queue data resides.
- Queue ID - The object number of the queue.
- Queue Type - The type of the queue.
- Maximum Retries - The maximum number of attempts that is allowed when dequeuing a message from the queue.
- Retry Delay - The time interval between retry attempts, when dequeuing a message from the queue.
- Enqueue Enabled - Whether the queue is enabled for enqueueing.
- Dequeue Enabled - Whether the queue is enabled for dequeuing.
- Retention - The time interval during which processed messages are retained in the queue.
- User Comments - Descriptive comments about the queue.

After reviewing the agent queue details, choose the OK button to return to the Agent Activity page.

Message Details

The Search Queue page lets you search for messages being held on a particular agent and review details about those messages. This page displays different message details depending on the payload type of the agent's queue.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Workflow Metrics > Agent Activity > (B) Search Agent Entry Details

WF_EVENT_T and SYS.AQ\$_JMS_TEXT_MESSAGE

This page lets you review messages on queues with a payload type of WF_EVENT_T, such as the standard WF_ERROR or WF_DEFERRED queues, or SYS.AQ\$_JMS_TEXT_MESSAGE, such as the standard WF_CONTROL queue.

Enter filter criteria to locate the messages you want to review and click the Go button. You can filter by the following message properties:

- Internal event name
- Event key

- Correlation ID used to associate a message with other related messages
- Enqueue date either within the last seven days or prior to the last seven days
- Dequeue date either within the last seven days, prior to the last seven days, or on any date
- Status

Oracle Workflow Manager displays the event messages on the queue for the selected agent that match your filter criteria. For each message, the list displays the event name, event key, correlation ID, event parameters, From System that sent the message, To System that received the message, date the message was sent, error message, error stack, and the message status.

The list also includes any messages on the exception queue associated with the selected queue. Messages are transferred from a user queue to the associated exception queue if Oracle Advanced Queuing cannot retrieve or process them for some reason. For more information, see: Oracle Streams AQ Exception Handling, *Oracle Streams Advanced Queuing User's Guide and Reference*.

Note: Each queue table contains one default exception queue that is shared by all the user queues in that queue table. When you search for messages on a particular queue, the search result list includes all messages on the associated exception queue as well, regardless of the user queue from which they originated. Consequently, if you create more than one user queue in the same queue table, the search result list may display exception messages that originated from other queues than the queue you selected.

- To review the event data for a message as an XML document, choose the message details icon in the View XML column.

Note: The message details icon is disabled if the event data for a message is empty.

- To add the information from this page to your support cart, click the Add to Support Cart button.

SYSTEM.ECXMSG

This page lets you review messages on queues with a payload type of SYSTEM.ECXMSG, including the standard Oracle XML Gateway ECX_INBOUND and ECX_OUTBOUND queues.

Enter filter criteria to locate the messages you want to review and click the Go button.

You can filter by the following message properties:

- Transaction type
- Document number
- Party site ID
- Correlation ID used to associate a message with other related messages
- Enqueue date either within the last seven days or prior to the last seven days
- Dequeue date either within the last seven days, prior to the last seven days, or on any date
- Status

Oracle Workflow Manager displays the messages on the queue for the selected agent that match your filter criteria. For each message, the list displays the message type, message standard, transaction type and subtype, document number, party ID, party site ID, party type, protocol type, protocol address, first, second, third, fourth, and fifth attributes, and the message status.

- To review the XML document for a message, choose the message details icon in the View XML column.

Note: The message details icon is disabled if the XML document for a message is empty.

- To add the information from this page to your support cart, click the Add to Support Cart button.

SYSTEM.ECX_INENGOBJ

This page lets you review messages on queues with a payload type of SYSTEM.ECX_INENGOBJ, including the standard Oracle XML Gateway ECX_IN_OAG_Q queue.

Enter filter criteria to locate the messages you want to review and click the Go button. You can filter by the following message properties:

- Message ID
- Correlation ID used to associate a message with other related messages
- Enqueue date either within the last seven days or prior to the last seven days
- Dequeue date either within the last seven days, prior to the last seven days, or on

any date

- Status

Oracle Workflow Manager displays the messages on the queue for the selected agent that match your filter criteria. For each message, the list displays the message ID, debug mode, and the message status.

To add the information from this page to your support cart, click the Add to Support Cart button.

Queue Propagation

You should schedule propagation for your local outbound agents to send event messages to their destinations. You can schedule Oracle Advanced Queuing (AQ) propagation for agents that use the SQLNET protocol by the following methods:

- Use the Distributed Database Management feature to manage AQ through Oracle Enterprise Manager. See: Oracle Enterprise Manager Support, *Oracle Streams Advanced Queuing User's Guide and Reference*.
- Run the DBMS_AQADM.Schedule_Propagation API in SQL*Plus. See: DBMS_AQADM, *Oracle Database PL/SQL Packages and Types Reference*.

If you want to use the standard WF_OUT and WF_JMS_OUT agents or custom agents for event message propagation, ensure that you schedule propagation for those agents. You do not need to schedule propagation for the WF_CONTROL or WF_NOTIFICATION_OUT agents, however, because the middle tier processes that use WF_CONTROL dequeue messages directly from its queue, and a notification mailer sends messages placed on the WF_NOTIFICATION_OUT queue.

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Related Links > Configuration > Queue Propagation

Queue Propagation

Use the Queue Propagation page to review the database initialization parameters required for queue propagation, as well as the existing propagation schedules for Business Event System agents in your instance of Oracle Workflow.

Database Initialization Parameters for Queue Propagation

For each parameter, this list shows the parameter name, actual parameter value, recommended value, and description. If the actual value does not match the recommended value, the recommended value is marked with a warning indicator icon.

The JOB_QUEUE_PROCESSES parameter defines the number of job queue processes for your instance. Oracle Workflow requires job queue processes to handle propagation of Business Event System event messages by AQ queues. The recommended number of

processes for Oracle Workflow is ten or more.

Note: In Oracle Database 10g and higher, you do not need to set the AQ_TM_PROCESSES parameter.

Queue Schedules

For each propagation schedule, the list displays the outbound queue, destination database link, job queue process executing the schedule, whether the schedule is enabled or disabled, and the error date and error message of the last unsuccessful execution. Click any column heading to sort the list by that column.

If no process is allocated to execute the schedule, you may need to increase the JOB_QUEUE_PROCESSES database initialization parameter to ensure that processes are available for propagation.

To view details for a propagation schedule, either click the queue link in the Queue column, or select the schedule and click the View Details button.

Queue Propagation Details

The Queue Propagation Details page displays the following details for a propagation schedule:

Navigation: Applications Dashboard > (pull-down menu) Workflow Manager > (B) Go > Related Links > Configuration > Queue Propagation > (B) View Details

- Destination - The destination database link.
- Process Name - The name of the job queue process executing this schedule.
- Enabled - Y if this schedule is enabled or N if the schedule is disabled. The schedule will not be executed if it is disabled.
- Last Error Date - The date of the last unsuccessful execution.
- Last Error Time - The time of the last unsuccessful execution.
- Last Error Message - The error message of the last unsuccessful execution.
- Schema - The schema that owns the queue.
- Session ID - The session ID (SID, SERIAL#) of the job executing this schedule; NULL if not currently executing.
- Propagation Window - The duration in seconds of the propagation window.
- Maximum Bytes - The maximum number of bytes propagated during a propagation

window.

- Failures - The number of times that execution of the schedule failed. If the number of failures reaches 16, the schedule will be disabled.
- Latency - The latency time in seconds that specifies how long to wait, after all messages have been propagated, before rechecking the queue for new messages to the destination. The latency represents the maximum wait time during the propagation window for a message to be propagated after it is enqueued.
- Next Run Date - The date at which the next propagation window of this schedule will be started.
- Next Run Time - The time at which the next propagation window of this schedule will be started, in HH:MI:SS format.
- Current Start Date - The date at which the current propagation window of this schedule was started.
- Current Start Time - The time at which the current propagation window of this schedule was started, in HH:MI:SS format.
- Instance - The cluster database instance number executing the schedule.
- Start Date - The date when propagation should be started, in the default date format.
- Start Time - The time when propagation should be started, in HH:MI:SS format.
- Last Run Date - The date of the last successful execution.
- Last Run Time - The time of the last successful execution, in HH:MI:SS format.
- Total Time - The total time, in seconds, spent by the system in executing this schedule.
- Total Number - The total number of messages propagated in this schedule.
- Total Bytes - The total number of bytes propagated in this schedule .
- Maximum Number - The maximum number of messages propagated during a propagation window.
- Average Number - The average number of messages propagated during a propagation window.
- Average Size - The average size of a propagated message, in bytes.

- Average Time - The average time, in seconds, to propagate a message.

License Manager

License Manager

License Manager is a utility that registers additional products, country-specific functionalities, and languages for your Oracle E-Business Suite system. Once you have contacted your Oracle sales representative, or set up your new license agreements online through the Oracle Store, you are ready to register your new products, country-specific functionalities or languages using License Manager. License Manager does not set up license agreements or determine pricing, but the registration procedure makes new components accessible to all Oracle E-Business Suite utilities.

Related to licensing is the subject of localizations, of which there are several types. This area is covered briefly at the end of the chapter.

License Manager also provides a set of reports that allows you to determine the products, country-specific functionalities and languages that are registered on your Oracle E-Business Suite system.

The main License Manager page contains three main licensing links and five report links. The licensing pages are:

- Products
- Country-specific Functionalities
- Languages

The five report links that provide licensing details about your Oracle E-Business Suite system are:

- Licensed Products
- Shared Products
- Country-specific Functionalities

- Languages
- Summary

See also: Applications Usage, page 14-20.

License Section

This section contains links to license products, country-specific functionalities, and languages.

License Products

Clicking the Products link in the License section of the License Manager main page opens the License Products page. This page displays two options and a link to show more options.

- License E-Business Suite: Select this option to register the predefined E-Business Suite of products.
- License Component Application: Select this option to register products by component applications.
- Show More Options: Select this link to show a third product licensing option, License Applications Product.
- License Applications Products: This option becomes visible when Show More Options is selected. Select this option to register Oracle E-Business Suite products individually.

Select the desired option and click Continue.

License E-Business Suite

Selecting License E-Business Suite in the Product Licensing page opens the License E-Business Suite page. This page displays all products that will be registered when you choose to register the "E-Business Suite". Once the E-Business Suite is registered, individual products within the suite cannot be unregistered. This page displays all products that can be registered and contains three columns of information:

- Select: There is a check mark for each product that will be registered.
- Focus: Select the circle icon next to a component application to see just the products in the component application.
- Name: This is the name of the component applications or the products within a component application. Click the (+) or (-) icon to hide or show the individual products within a component application.

Click Next to move on to the License E-Business Suite Add-ons page.

License E-Business Suite Add-ons

The License E-Business Suite Add-ons page lists the products that are not included in the standard "E-Business Suite" list of products. Select the E-Business Suite Add-on products on this page. This page contains three columns of information:

- **Select:** Check the check box for the products that you want to register.
- **Focus:** Select the circle icon next to a component application to see just the products in the component application.
- **Name:** This is the name of the component applications or the products within a component application. Click the (+) or (-) icon to hide or show the individual products within a component application.

To register all E-Business Suite add-ons, click the Select All link. To deselect all selected E-Business Suite add-ons, click the Select None link. Once an E-Business Suite add-on is registered, it has a check box that is disabled and cannot be unregistered.

Click Next to advance to the License E-Business Suite Review page.

License E-Business Suite Review

Navigation: Site Map (Administration) > License Manager > Products > License E-Business Suite

The License E-Business Suite Review page lists the products that you selected to register in the License E-Business Suite and License E-Business Suite Add-ons pages. This page contains two columns of information:

- **Focus:** Select the circle icon next to a component application to see just the products in the component application.
- **Name:** This is the name of the product to register. Click the triangle icon to hide or show the individual products within a component application.

Click Submit to register the products.

License Component Application

Selecting License Component Application in the Product Licensing page opens the License Component Application page. This page displays all component applications that can be registered and contains three columns of information:

- **Select:** Check the check box for the component applications that you want to register.
- **Focus:** Select the circle icon next to a component application to see just the products in the component application.
- **Name:** This is the name of the component applications or the products within a

component application. Click the (+) or (-) icon to hide or show the individual products within a component application.

To register all component applications, click the Select All link. To deselect all selected component applications, click the Select None link. Once a component application is registered, the individual products within the component application have check boxes that are greyed out and cannot be unregistered.

Click Next to advance to the License Component Application Review page.

License Component Application Review

The License Component Application Review page lists the products that you selected to register in the License Component Application page. This page contains two columns of information:

- **Focus:** Select the circle icon next to a component application to see just the products in the component application.
- **Name:** This is the name of the component application to register. Click the blue triangle to hide or show the individual products within a component application.

Click Submit to register the products.

License Applications Products

Selecting License Applications Product in the Product Licensing page opens the License Applications Products page. This page displays all products in the Oracle E-Business Suite system and allows you to register them individually.

To register all products, click the Select All link. To deselect all selected products, click the Select None link. Once a product is registered, it has a check box that is grayed out and cannot be unregistered.

Click the check box of the products that you want to register and click Next. This takes you to the License Applications Products Review page.

License Applications Product Review

The License Applications Products Review page lists the products that you selected to register in the License Applications Products page. This page contains two columns of information:

- **Product Name:** This is the name of the product to register.
- **Product Abbreviation:** This is the short name of the product to register, for example, 'AS'.

Click Submit to register the products.

License Country-specific Functionalities

Selecting Country-specific Functionalities in the License section of the License Manager main page produces the License Country-specific Functionalities page. This page

displays all country-specific functionalities in the Oracle E-Business Suite system and allows you to register them. This page contains three columns of information:

- **Select:** Check the check box for the country-specific functionality that you want to register. The already registered country-specific functionalities have check boxes that are greyed out. Once a country-specific functionality is registered, it cannot be unregistered.
- **Country Name:** This is the name of the country-specific functionality.
- **Country Short Name:** This is the short name of the country-specific functionality to register.

Once you select the country-specific functionalities that you want to register, click Next. This takes you to the License Country-specific Functionalities Review page.

License Country-specific Functionalities Review

The License Country-specific Functionalities Review page lists the country-specific functionalities that you selected to register in the License Country-specific Functionalities page. This page contains two columns of information:

- **Country Name:** This is the country name of the country-specific functionality to register.
- **Country Short Name:** This is the short name of the country-specific functionality to register, such as CO or JP.

Click Submit to register the country-specific functionality.

License Languages

Selecting Languages in the License section of the License Manager main page opens the License Languages page. This page displays all languages available for the Oracle E-Business Suite system and allows you to register them. This page contains these columns of information:

- **Select:** Check the check box for the languages that you want to register. The already registered languages have check boxes that are disabled.
- **Language Name:** This is the name of the language.
- **Language Code:** This is the language code, such as US or ESA.

Click the check box of the languages that you want to register and click Next. This takes you to the Base Language page.

Base Language

The Base Language page shows the current base language and list of languages that you

can select as a base language for your Oracle E-Business Suite system.

The Current Base Language section contains one row and two columns of information:

- Name: This is the name of the current base language.
- Language Code: This is the base language code.

The Select New Base Language section contains a row for each registered language and three columns of information:

- Select: Select the language that you want to set as the base language.
- Language Name: This is the name of the language.
- Language Code: This is the language code.

Click Next to continue to the License Languages Review page.

License Languages Review

The License Languages Review page lists the languages that you selected to register in the License Languages page and the base language that you selected in the Base Language page. There are two sections in this page.

The Selected Languages section contains a row for each language you want to register and these columns of information:

- Name: This is the name of the language to register.
- Language Code: This is the language code of the language to register, such as CA or ESA.

The Base Language section contains one row and two columns of information:

- Name: This is the name of the selected base language.
- Language Code: This is the base language code, for example, US.

Click the Submit button to register the languages and set the base language and territory.

Reports Section

This section contains links to reports.

Licensed Products Report

Clicking the Licensed Products link in the Reports section of the License Manager main page produces the Licensed Products report. The report has two sections. The first section, Summary, shows the Status information. Status is the number of products

installed and the number of products shared. Clicking on one of these status groups refreshes the second section of this report, List of Products according to the status selected.

Depending upon which group (Licensed or Shared) you clicked in the Summary section, the List of Products changes to show all licensed products or all shared products in the system. The List of Products section has four columns:

- **Select:** This option button determines which product's patch summary information is presented in the Patch Summary page.
- **Product Abbreviation:** This is the product short name, for example, FND or GL.
- **Product Name:** This is the name of the fully licensed product.
- **Status:** This is the license status of the product.

A filter at the top of the List of Products section allows you to narrow the contents of the report. You can filter by Product Abbreviation, Product Name, or (license) Status. For Status, you can choose from Licensed, Shared, or Not Licensed.

From this report you can access the Patch Information page for a specific product by selecting the product and clicking the Patch Information button, or by clicking the Product Name.

Shared Products Report

Clicking the Shared Products link in the Reports section of the License Manager main page produces the Shared Products report. The report has two sections. The first section, Summary, shows the Status information. Status is the number of products installed and the number of products shared. Clicking on one of these status groups refreshes the second section of this report, List of Products according to the status selected.

The List of Products section has four columns:

- **Select:** This option button determines which product's patch summary information is presented in the Patch Summary page.
- **Product Abbreviation:** The product short name, for example, FND or GL.
- **Product Name:** The name of the fully licensed product.
- **Status:** The license status of the product.

A filter at the top of the List of Products section allows you to narrow the contents of the report. You can filter by Product Abbreviation, Product Name, or (license) Status. For Status, you can choose from Licensed, Shared, or Not Licensed.

From this report you can access the Patch Information page for a specific product by selecting the product and clicking the Patch Information button, or by clicking the

Product Name.

Country-specific Functionalities Report

Clicking the Country-specific Functionalities link in the Reports section of the License Manager main page produces the Country-specific Functionalities report. This report displays all registered country-specific functionalities in the Oracle E-Business Suite system and contains two columns of information:

- Country Name: This is the country name of the country-specific functionality.
- Country Short Name: This is the country-specific functionality short name, such as CO or JP.

Clicking **OK** on the report returns you to the main License Manager page.

Clicking **Edit** takes you to the License Country-specific Functionalities page.

Languages Report

Clicking the Languages link in the Reports section of the License Manager main page produces the Languages report. This report displays the current database character set, the base language, and all registered languages.

The Licensed Languages section contains a row for each registered language and two columns of information:

- Language Name: This is the name of the registered language.
- Language Code: This is the short name of the registered language, such as CA or ESA.

The Base Language section contains one row and two columns of information:

- Language Name: This is the name of the base language.
- Short Name: This is the base language short name, for example, US.

Clicking **OK** on the report returns you to the main License Manager page.

Clicking **Edit** takes you to the License Languages page.

License Summary Report

Clicking the Summary link in the Reports section of the License Manager main page produces the License Summary report. This report displays a summary of all registered products, country-specific functionalities, languages, and base language. There are five sections in this report.

The Licensed Products section contains a row for each fully licensed product registered in the system and two columns of information:

- **Product Name:** This is the name of the registered product.
- **Product Abbreviation:** This is the product short name, for example, FND or GL.

The Shared Products section contains a row for each shared product registered in the system and two columns of information:

- **Product Name:** This is the name of the shared product.
- **Product Abbreviation:** This is the product short name, for example, AD or OE.

The Country-specific Functionalities section contains a row for each registered country-specific functionality and two columns of information:

- **Country Name:** This is the country name of the country-specific functionality.
- **Country Short Name:** This is the country-specific functionality short name, for example, CO or JP.

The Licensed Languages section contains a row for each registered language and two columns of information:

- **Language Name:** This is the name of the registered language.
- **Language Code:** This is the code of the registered language, for example, CA or ESA.

The Base Language section contains one row and two columns of information:

- **Language Name:** This is the name of the base language.
- **Short Name:** This is the base language short name, for example, US.

Localizations

Related to licensing, country-specific functionalities known as *localizations* provide the required business processes to meet the statutory, legal, and cultural practices of a given locality.

There are three types of localization:

- **Product Localizations** - Delivered as part of the standard product by Oracle E-Business Suite. Development
- **Add-on Localizations** - Delivered by Regional Field Centers (add-on localization teams) via My Oracle Support.
- **Partner Localizations** - Delivered by partners including ISVs and system integrators.

These three types of localizations are activated differently in Oracle E-Business Suite. Only product localizations can be activated via Rapid Install or License Manager. Add-on localizations and partner localizations are installed via special procedures created by regional teams and partners.

For more information about localizations, refer to My Oracle Support Knowledge Document 973912.1, *Oracle E-Business Suite Globalization Center*.

Functional Administrator and Functional Developer Tasks

Overview of Functional Administrator and Functional Developer Responsibilities

Oracle E-Business Suite ships two responsibilities that provide access to a subset of system administrator tasks. These tasks are primarily those setup tasks using Oracle E-Business Suite HTML-based pages.

Functional Administrator Responsibility

From the Functional Administrator responsibility you can create and/or manage the following features.

From the Security tab:

- Grants
- Permissions and Permission Sets

For more information on using grants and permissions, see: *Overview of Oracle E-Business Suite Security, Oracle E-Business Suite Security Guide*.

From the Core Services tab:

- Lookups
- Messages
- Profiles and Profile Categories
- Functions
- Menus

- Caching Framework

For information on Lookups, see: Application Utilities Lookups and Oracle Application Object Library Lookups, *Oracle E-Business Suite Developer's Guide*.

For information on Messages, see: Overview of Message Dictionary, *Oracle E-Business Suite Developer's Guide*.

For information on Profiles, see Overview of Setting User Profiles, *Oracle E-Business Setup Guide*.

For information on functions and menus, see: Overview of Oracle E-Business Suite Security, *Oracle E-Business Suite Security Guide*.

For information on the Caching Framework, see the section Caching Framework, page 17-5 and the *Oracle E-Business Suite Java Caching Framework Developer's Guide*, My Oracle Support Knowledge Document 1108093.1.

From the Personalization tab:

- Application Catalog
- Import/Export

For more information on Personalization, see the section Oracle Application Personalization Framework, page 17-3.

For more information on File Manager, see the section Generic File Manager Access Utility (FNDGFU), *Oracle E-Business Suite Setup Guide*.

For more information on Portletization, see the section Portlet Generator, *Oracle E-Business Suite Setup Guide*.

Functional Developer Responsibility

From the Functional Developer responsibility you can create and/or manage the following features.

From the Security tab:

- Objects
- Permissions and Permission Sets

For more information on objects and permissions, see: Overview of Oracle E-Business Suite Security, *Oracle E-Business Suite Security Guide* and Overview of Data Security, *Oracle E-Business Suite Security Guide*.

From the Core Services tab:

- Lookups
- Messages

- Profiles
- Functions
- Menus
- Cache Components

For more information on using the Lookups and Messages windows, refer to the online help as well as the *Oracle E-Business Suite Developer's Guide*. For more information on Profiles, see Overview of Setting User Profiles, *Oracle E-Business Suite Setup Guide*. For more information on functions and menus, see: Overview of Oracle E-Business Suite Security, *Oracle E-Business Suite Security Guide*. For more information on the Cache Components, see the *Oracle E-Business Suite Java Caching Framework Developer's Guide*, My Oracle Support Knowledge Document 1108093.1.

Oracle Application Personalization Framework

Personalization allows you to declaratively tailor the UI look-and-feel, layout or visibility of Oracle Application Framework-based (HTML-based) pages to suit business needs or user preferences.

Durability of Oracle Application Framework personalizations is largely attributed to the declarative architecture and the object-oriented approach underlying the implementation of the page. Declarative UI component definitions are stored in the form of metadata in a database repository. Personalizations are translated into offsets from the base metadata definition and stored separately. At runtime, the applicable personalizations metadata is uploaded from the repository and layered over the base metadata definition to produce the net effect. Product upgrades and patches affect only the base metadata definition, so customer personalizations continue to function properly as applicable.

For more information on personalization in Oracle Application Framework, see: Personalization, *Oracle Application Framework Personalization Guide*.

Oracle Application Framework comes with an administration user interface for personalizations which is available under the Functional Administrator responsibility. This interface contains the following two pages that can be used to personalize the pages of Oracle Application Framework-based applications at various personalization levels without modifying any code:

- Application Catalog
- Import/Export

The Application Catalog page is useful for managing several personalizations across pages and applications, especially where the administrator does not have a responsibility that can access the page directly.

You can change the layout of a page by adding rows and columns to the customizable

regions. You can also change the layout direction and order of contents inside these regions. You can update different elements and you can also add, create, or remove the contents from different regions.

Note: To rearrange contents across different regions, you must first remove them from their current location and then add them inside the new destination region.

To activate, inactive, or delete specific personalizations, or manage the translation of the personalizations made for the page in question, navigate to the *Manage Personalization Levels* page.

Depending on the type of page you selected to personalize, (configurable or non-configurable), you are automatically directed to one of the following two personalization launch pages:

Page Layout Personalization: (Configurable page). This launch page provides a boxed preview of the flexible layout structure within your page and displays controls that take you to different pages or flows where you specify and apply your actual personalizations.

Important: Page Layout Personalization: (Configurable page). This launch page provides a boxed preview of the flexible layout structure within your page and displays controls that take you to different pages or flows where you specify and apply your actual personalizations.

- Mandatory user-entered parameters
- Flow/business-logic
- Limited access to specific users
- Multi-organization access control

These parameters might not be available and the page might fail with unexpected errors. You should instead access the Personalization UI for your configurable page using the global *Personalize Page* link on the page itself, when the **Personalize Self-service Defn** profile option is enabled.

Page Hierarchy Personalization: (Non-configurable page). This launch page displays the entire structure of the selected page in a hierarchy table (HGrid), rather than as a visual boxed layout.

The Import/Export page allows you to both export meta-data to XML files, and import XML files into a MDS repository.

Both administration-level and user-level personalizations may be extracted from one database and loaded into another. This allows you the freedom to create and test personalizations in a test database before deploying the personalizations to a

production instance.

Use the **FND:Personalization Document Root Path** (FND_PERZ_DOC_ROOT_PATH) profile option to define the root path of the current deployed environment where personalizations are exported to and imported from. We recommend you set this profile to the \$APPL_TOP staging area and at the site level.

```
$APPL_TOP/<CompanyIdentifier>/<CustProductShortName>/<ProductVersion>/mds/webui
```

See *Deploying Personalizations*, *Oracle Application Framework Personalization Guide* for more information.

Caching Framework

Caching provides a powerful, flexible, easy-to-use mechanism for storing database results and other Java objects in memory for repeated usage. This mechanism minimizes expensive object initializations and database round trips, thereby improving application performance.

Application data is cached using component caches. Each component cache is identified by a name. The objects contained in a component cache are generally of the same type and share the same caching attributes. Each component cache has an associated cache loader class. The loader class has the logic for loading the cached object in case of a cache miss. When an object is requested from a component cache, if the object is found, it is returned from the cache. Otherwise, the loader is used to load the object place it in the cache.

For additional information on the Caching Framework, see the *Oracle E-Business Suite Java Caching Framework Developer's Guide*, My Oracle Support Knowledge Document 1108093.1.

Caching Framework comes with an administration user interface, which is available under the Functional Administrator responsibility. This interface contains the following three pages that can be used to implement tuning of the memory management policies and perform administrative operations:

- Overview page
- Tuning page
- Global Configuration page

As a general rule, cache administration should not be required unless there are some performance problems.

The Caching Framework Overview page provides a Cache Usage Summary, listing the following:

- Total Cache Components - All registered cache components in the system. This includes the cache components that have statistics enabled.

- **Global Idle Time** - A global setting for the elapsed time since an object was accessed last. This value applies to cache components that rely on the default Idle Time and will not override the Idle Time setting of the individual cache components.
- **Cache Components with Statistics Enabled** - All registered cache components in the system that have statistics enabled.

On the Tuning page, you can search for cache components and then measure the Caching Framework performance by enabling statistics for frequently-used components. The statistics provided include hits, misses, the hit/misses ratio and invalidation count for each cache component. You can also clear collected statistics, and clear the cache.

Tip: A cache 'miss' is when a requested object from a cache component is not found in the cache. To reduce the 'misses' value for a particular cache component, update the **Time Out Type** and **Time Out After** values of the cache component definition. An object is marked 'invalid' when the object has been *idle* beyond the idle timeout period or the object was updated, making the copy in the cache invalid. When an object is 'invalid', any subsequent `get()` operations on the object gets a new copy of the object from the database.

- **Time Out Type:** Choose either *Idle Time* (recommended) or *Time to Live*. Both values refer to the duration after which the object is marked invalid.
 - *Idle Time:* Starts from the last time the object was requested from the cache. Choose this value when the primary consideration is the memory. This option prevents infrequently used objects from being cleaned up from the cache.
 - *Time to Live:* Starts from the time the object is loaded into the cache. Choose this value when the primary consideration is data consistency. This option guarantees that the values are refreshed after the specified time interval regardless of the usage.
- **Time Out After:** This refers to the Time Out Type. We recommend choosing *Global Idle Time*.
 - *Global Idle Time:* The component cache gets a timeout value that is equal to the global idle time specified. The current default is 15 minutes. This value can be changed on the Global Configuration page.

Note: Changes to the cache components definition's **Time Out Type** and **Time Out After** values will not be put into effect until after the middle tier is bounced.

Through the Global Cache Configuration page you can update the cache statistics and

cache policy for all the cache components.

- **Cache Statistics:** You can choose to enable statistics for all the cache components. However, doing so may affect the performance of the system. You can also clear statistics for all the cache components.
- **Cache Policy:** You can set the Global Idle Time profile option, which refers to the duration after which any object is marked invalid. You can also clear all cache components, which removes all the cache components from the middle tier. Changing the cache policy can affect performance.

Note: Enabling or disabling the statistics collection of the cache components only affects the current Java Virtual Machine (JVM). To enable/disable statistics collection in other JVMs, bounce those JVMs. The same is true for changes to the Global Idle Time profile option.

Using Alerts

Overview of Oracle Alert

Oracle Alert is your complete exception control solution.

Oracle Alert gives you an immediate view of the critical activity in your database. It helps you keep on top of important or unusual business events you need to know about, as they happen. Oracle Alert gives you real-time measurements of staff and organization performance, so you can zero in on potential trouble spots immediately. You can automate routine transactions with Oracle Alert, saving your valuable time for more essential tasks. And, Oracle Alert does all this online, so you do not have to contend with a pile of paperwork.

Oracle Alert gives you the flexibility you need to monitor your business information the way you want.

For more information on Oracle Alert, see the *Oracle Alert User's Guide*.

Basic Business Needs

Oracle Alert meets the following basic business needs:

- Informs you of exception conditions as they occur
- Lets you specify the exception conditions you want to know about, as often as you want to know about them
- Informs you of exception conditions by sending alert messages through a single application -- your electronic mail
- Takes actions you specify, based upon your response to an alert message
- Automatically performs routine database tasks, according to a schedule you define
- Integrates fully with your electronic mail system

Oracle Alert Runtime Features

If you do not have a licensed copy of the full Oracle Alert product, you may still derive benefit from major Oracle Alert features by using the predefined alerts that are packaged with your Oracle E-Business Suite product.

All Oracle E-Business Suite products are packaged with a runtime version of Oracle Alert. Although all the Oracle Alert windows are available in this runtime version, not all the features in those windows are enabled. With the runtime version of Oracle Alert, you can run only the predefined alerts that are packaged with your product; you cannot create new alerts.

Alert Definitions

Alert

A mechanism that checks your database for a specific exception condition. An alert is characterized by the *SQL SELECT statement* it contains. A SQL SELECT statement tells your application what database exception to identify, as well as what output to produce for that exception.

For example, you can define an alert to flag purchase orders exceeding \$10,000, and have that alert output the name of the individual who requested the purchase order, as well as the name of the individual's manager. All predefined alerts are listed in the Alerts window of Oracle Alert.

Event Alert

An event alert monitors the occurrence of a specific exception or change in your database. An exception in your database results if you add or update information using your Oracle E-Business Suite windows. The event alert monitors the database for exceptions based on its SQL SELECT statement.

Periodic Alert

A periodic alert periodically reports key information according to a schedule that you define. Rather than notify you of immediate exceptions in the database like an event alert, a periodic alert scans for specific database information specified by its SQL SELECT statement at scheduled intervals.

Alert Action

An alert action is an action you want your alert to perform. An alert action can be dependent on the output from the alert. An alert action can fall under one of three categories:

- Detail action—an action that represents one exception found in the database

- Summary action-an action that represents multiple exceptions found in the database
- No exception action-an action that represents no exceptions found in the database

An action can include sending an electronic mail message to a mail ID, running an Oracle E-Business Suite program, running a program or script from your operating system, or running a SQL script to modify information in your database.

You can have more than one action for an alert and an action can incorporate the output of the alert. For example, you may want a particular alert to send a message to a manager, as well as run an Oracle E-Business Suite program when an exception occurs.

Action Sets

An action set is a sequence of alert actions that are enabled for a particular alert. Each action that you include in an action set can be assigned a sequence number so that you can specify the order in which the actions are performed. Some predefined alerts may also have more than one action set. You can also assign a sequence number to each action set to specify the order in which each action set is performed.

Predefined Alerts

There are two types of predefined alerts:

- **Event alerts**-for example, the Receiving Notification alert for Oracle Purchasing notifies the requestor with a mail message when an item is received and entered in the Receipts window.
- **Periodic alerts**-for example, the Forecast Over-Consumption alert for Oracle Material Planning checks every day for over-consumption of the forecast and sends you a mail message if the current forecast quantity listed in the Forecast Entries window goes below zero.

Tip: See your product's reference guide for a list of the predefined alerts that are packaged with your Oracle E-Business Suite product.

Using Predefined Alerts

All predefined alerts are initially disabled. You must enable the alerts you want to use. Select the Oracle Alert Manager responsibility when you start Oracle E-Business Suite to view or use a predefined alert. The Alert Manager responsibility gives you access to the Oracle Alert menu.

Navigate to the Alerts window to enable or edit predefined alerts. To display the predefined alert(s) for your Oracle E-Business Suite product, execute a query with your Oracle E-Business Suite product name in the Application field.

The Name field displays the name of the predefined alert. The Type field indicates if the alert is an event or a periodic alert.

You can enable an alert to run by checking the Enabled check box. You can also enter an End Date to specify the date until you want this alert run.

Choose the Alert Details button to open the Alert Details window. Choose the Alert Installations tabbed region to display the available Installations.

Enter the Oracle ID of the application installation you want your alert to run against. You can select only the Oracle IDs that are associated with the application that owns your alert. You can disable an Oracle ID for the alert temporarily by unchecking the Enabled check box.

Choose the Actions button to open the Actions window. Oracle Alert automatically displays the actions that are defined for the alert.

In the Actions window, if the Action Type is Detail, choose the Action Details button to display details for that action.

The alert action sends an alert action message to the mail ID listed in the To field of the Message Detail zone. If the mail ID is in the format *&NAME*, where *Name* is an output defined by your alert, you need not modify this field. If, however, the mail ID in the To field is not in the above format or if there is no value entered in the field, you must enter the mail ID(s) of the person(s) you wish to receive the alert action message. After modifying the contents of this window, save your work.

Navigate to the Oracle Alert Options window. Use this window to define the options Oracle Alert uses when checking your alerts.

In the Alerts window, choose the Actions Sets button to navigate to the Action Sets window. Oracle Alert automatically displays the action sets defined for the alert.

Check the Enabled check box for each action set you wish to use. You may also enter an End Date field to specify the date until you want this alert action set to be enabled.

In addition, in the Action Set Members block, check the Enabled check box for each action set member you want to use in that action set.

You may also enter an End Date to specify the date until you want this alert action set member to be enabled. When you finish, save your work.

Your predefined alert is now ready to use.

Customizing Predefined Alerts

You can customize predefined alerts in the following ways to suit your business needs:

Electronic Mail Integration

Oracle Alert leverages the Workflow Notification Mailer to send alert e-mail messages to your users. Ensure that you set up mail servers and configure the Workflow Notification Mailer to send e-mail messages according to your alert requirements. See:

Standard Alert Message Text

You can customize the message header and footer text that appears in all your alert message actions. Navigate to the Message Elements tabbed region of the Oracle Alert Options window, and four message elements appear automatically. Each element represents a specific type of message text that appears in all your alert mail messages.

In the runtime version of Oracle Alert, you need to edit only the Message Action Header and Message Action Footer elements. Simply customize the text that appears to alter the text at the beginning and end of every alert message. You may also leave the text blank if you do not want to display any standard text in your alert messages. Save your work when you are done making changes in this window.

Alert Frequency

You can schedule the frequency you wish to run each predefined periodic alert. You may want to check some alerts every day, some only once a month, still others only when you explicitly request them. You have the flexibility to monitor critical exceptions every day, or even multiple times during a 24-hour period. And, you can set less significant exceptions to a more infrequent schedule; for example, a monthly schedule.

To change the frequency of a predefined alert, navigate to the Alerts window. Perform a query to display the predefined periodic alert you wish to modify, then alter the Frequency of the periodic alert.

Alert History

Oracle Alert can keep a history of exceptions and actions for a particular alert. Use the Alerts window to alter the number of days of history you wish to keep for an alert. Simply change the Keep N Days field to the number of days of history you wish to keep.

Suppressing Duplicates

If you do not want Oracle Alert to send repeated messages for the same alert exception, you can choose to suppress duplicate messages. If Oracle Alert finds a duplicate exception condition for the alert, it simply does not execute the action set members for that alert again.

Use the Suppress Duplicates check box in the Action Sets block of the Alerts window to specify this option. The default for the Suppress Duplicates check box is unchecked. If you check the Suppress Duplicates check box, you must also make sure you keep history for the alert at least one day longer than the number of days between alert checks. Oracle Alert uses the history information to determine if an exception is a duplicate.

Message Actions

If a predefined alert involves a message action, you can customize certain aspects of that message action. Navigate to the Actions block in the Alerts window by choosing the Actions button. In this block, move your cursor to the row representing the message action you want to customize, then choose the Action Details button to open the Action Detail window for that message action. You can modify the following features of the message action:

- Recipient list—you can add or delete mail IDs in the List, To, Cc, Bcc, or Print For User fields. You should not modify any mail IDs listed with the format *&Name*, as they represent mail ID's defined by the alert output.
- Printer—you can modify the name of the printer to which you want Oracle Alert to direct the message.
- Text—you can modify the boilerplate text that you want your alert message to send. Do not edit any of the alert outputs (in the format *&Name*) used in the body of the text. For summary messages, edit only the opening and closing text within the summary message. Save your work when you finish making modifications.

Summary Threshold

Predefined alerts use one of three action types: detail action, summary action, and no exception action. A no exception action is straightforward in that Oracle Alert performs the defined action when no exceptions are found for the alert.

But how does Oracle Alert know when to perform a detail or a summary action? Oracle Alert can perform a detail action for every exception it finds, regardless of the number of exceptions, or Oracle Alert can perform a summary action for a unique set of exceptions. For example, you can receive individual mail messages for each exception found by an alert, or you can receive a single mail message summarizing all the exceptions found by the alert.

In the Members tabbed region of the Action Sets block of the Alerts window, you can set a Summary Threshold to specify how many exceptions Oracle Alert can find before it should change the action from a detail action to a summary action.

Oracle Alert Precoded Alerts

Your Oracle Alert installation contains custom alerts that are designed to help you manage your database and the data you generate when you use Oracle Alert. Oracle Alert provides eight alerts that systematically monitor your system for potential tablespace, disk space, and allocation problems, making your Database Administrators more efficient, and increasing database performance.

Occasionally, you will want to purge your database of obsolete concurrent requests, alert checks, and action set checks. Oracle Alert provides two alerts that let you

periodically remove old files, freeing up valuable tablespace and increasing database performance.

This section gives you an overview of these alerts, and suggestions on how to use them to enhance your system performance.

Terms

Before reading this discussion of precoded alerts, you may want to familiarize yourself with the following Glossary terms:

- Periodic Alert
- Exception
- Action
- Detail Action
- Summary Action
- No Exception Action
- Input

Oracle Alert DBA Alerts

Oracle Alert DBA alerts help you manage your database by notifying you regularly of:

- Tables and indexes unable to allocate another extent
- Users who are nearing their tablespace quota
- Tablespaces without adequate free space
- Tables and indexes that are too large or are fragmented
- Tables and indexes that are near their maximum extents

Customizable Alert Frequencies

Oracle Alert DBA alerts are periodic alerts, so you determine how often they check your database. Set them to run daily, weekly, or monthly, according to your database needs.

Summary and No Exception Messages

If Oracle Alert finds the database exceptions specified in a DBA alert, it sends you a message summarizing all exceptions found. If Oracle Alert finds no exceptions, it sends you a message reporting that no exceptions were found. Oracle Alert keeps you notified

of the status of your database, even if it is unchanging.

Customizable Alert Inputs

Inputs let you customize your DBA alerts. You can specify the ORACLE username, table, or index you want your alerts to target, and you can specify the threshold number of extents, maximum extents, or blocks Oracle Alert should look for. You can also define your input values at the action set level, so you can create multiple action sets that target different usernames, tables, and indexes. You can create as many action sets as you need.

Support for Multiple Database Instances

The Applications DBA application owns the Oracle Alert DBA alerts. This lets Oracle Alert perform the DBA alerts for every database instance you create, even those that reside outside Oracle Alert's database.

Applications DBA Alerts Descriptions

The following descriptions list the customizable frequency and inputs of each DBA alert.

Tables Unable to Allocate Another Extent

This alert looks for tables where the next extent is larger than the largest free extent.

Frequency	Every N Calendar Days
Inputs	Table Name, ORACLE Username

Indexes Unable to Allocate Another Extent

This alert looks for indexes where the next extent is larger than the largest free extent.

Frequency	Every N Calendar Days
Inputs	Index Name, ORACLE Username

Users Near Their Tablespace Quota

This alert detects users that are near their tablespace quota.

Frequency	Every N Calendar Days
Inputs	ORACLE Username
	Tablespace Name
	Check minimum percent free space remaining

Check maximum percent space use
Minimum total free space remaining (in bytes)
Maximum percent space used

Tablespaces Without Adequate Free Space

This alert looks for tablespaces without a specified minimum amount of free space.

Frequency Every N Calendar Days

Inputs Tablespace Name

Check total free space remaining
Check maximum size of free extents available
Maximum size of free extents available (in bytes)
Minimum total free space remaining (in bytes)

Indexes Too Large or Fragmented

This alert detects indexes that exceed a specified number of blocks or extents.

Frequency Every N Calendar Days

Inputs Index Name

ORACLE Username

Check maximum number of blocks
Check maximum number of extents
Maximum number of blocks
Maximum number of extents

Tables Too Large or Fragmented

This alert detects tables that exceed a specified number of blocks or extents.

Frequency Every N Calendar Days

Inputs Table Name

ORACLE Username

Check maximum number of blocks
Check maximum number of extents
Maximum number of blocks

Maximum number of extents

Tables Near Maximum Extents

This alert searches for tables and indexes that are within a specified number of extents of their maximum extents.

Frequency	Every N Calendar Days
Inputs	Table Name
	ORACLE Username
	Minimum number of extents remaining

Indexes Near Maximum Extents

This alert searches for tables and indexes that are within a specified number of extents of their maximum extents.

Frequency	Every N Calendar Days
Inputs	Index Name
	ORACLE Username
	Minimum number of extents remaining

Oracle Alert Purging Alerts

Two of the Oracle Alert precoded alerts are designed to help you manage the data you generate when you use Oracle Alert. While using Oracle Alert you should be able to:

- Automatically delete concurrent requests older than a specified number of days
- Automatically clean out alert checks and action set checks that are older than a specified number of days

Customizable Alert Frequencies

You determine the schedule for running your purge alerts. On the schedule you define, Oracle Alert submits the purge alerts to the Concurrent Manager, and deletes all old concurrent requests.

Customizable Alert Inputs

Inputs let you customize your alerts. You specify which application and which concurrent program you want your purge alerts to target, and you decide when your data becomes unnecessary or "old." You define your input values at the action set level, so you can create multiple action sets that target different applications and different

concurrent programs. You can create as many action sets as you need, so you can keep your system free from unnecessary files.

Oracle Alert Purging Alerts Descriptions

The following descriptions list the customizable frequency and inputs of each purging alert.

Purge Alert and Action Set Checks

This alert looks for alert and action set checks older than the number of days you specify, and runs a SQL statement script that deletes them.

Alert Type	Periodic
Periodicity	Every N Calendar Days
Inputs	Application Name, Number of days since alert check

Note: Oracle Alert will not delete alert checks and/or action set checks for a response processing alert that has open responses.

Purge Concurrent Requests

This alert looks for concurrent requests and their log and out files that are older than the number of days you specify, and runs a concurrent program that deletes them. If you enter a concurrent program name input, you should use the program name (located in the column USER_CONCURRENT_PROGRAM_NAME in the table FND_CONCURRENT_REQUESTS), and not the optional description that may accompany the concurrent program name in the Requests window.

Alert Type	Periodic
Periodicity	Every N Calendar Days
Inputs	Application Name Concurrent Program Name Number of days since concurrent request was submitted to the Concurrent Manager
Operating System Program	Deletes log file, out file, and corresponding record of each concurrent request
Arguments	Concurrent request ID

Managing Query Optimization Statistics

Oracle E-Business Suite and Query Optimization

Oracle E-Business Suite Release 12 uses cost-based optimization in order to choose the most efficient execution plan for SQL statements. Using this approach, the optimizer determines the most optimal execution plan by costing available access paths and factoring information based on statistics for the schema objects accessed by the SQL statement.

Oracle E-Business Suite requires several database initialization parameters to be set correctly in order to ensure optimal performance. Refer to My Oracle Support Knowledge Document 396009.1, and ensure that you have configured the parameters according to this note.

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. To this end, database statistics should be refreshed periodically. However, that does not necessarily imply that you should gather statistics frequently. Systems that are close to going live typically experience inserts of a large amount of data, as data from legacy systems is migrated. In that scenario, the statistics would probably need to be refreshed quite frequently (for instance, after each major load), as large loads could change the data distribution significantly. Once the system reaches steady state, the frequency of statistics collection at the schema/database level should be reduced to something like once a month. However, statistics on some volatile tables can be gathered as frequently as required.

Oracle E-Business Suite provides a set of procedures in the FND_STATS package to facilitate collection of these statistics. FND_STATS uses the DBMS_STATS package to gather statistics.

In gathering statistics, please do the following:

- Use only FND_STATS or the Gather Schema and Gather Table Statistics Concurrent Programs.

- Do NOT use the `analyze` or `dbms_stats` commands directly. The direct use of these commands is not supported and results in suboptimal plans.
- Use the `GATHER_AUTO` option to gather incrementally.
- Analyze all schemas at 10%, then specific objects at x%. Due to data skew, some products/tables benefit from higher sampling %.
- Disable the Oracle Database 10g/11g automatic job to gather statistics. Do not gather statistics excessively on entire schemas or the entire database, such as nightly or weekly.
- Note that `FND_STATS` now supports Database 11g Extended Statistics.

Oracle recommends using the `NOWORKLOAD` stats mode. Using `WORKLOAD` stats might introduce problems in execution plan instability

`FND_STATS.GATHER_SCHEMA_STATS` uses a parameter called `OPTIONS`. If set to `GATHER AUTO`, this option allows `FND_STATS` to determine automatically the tables for which statistics should be gathered, based on the change threshold. The Modifications Threshold can be adjusted by the user by passing a value for `modpercent`, which by default is equal to 10. `GATHER AUTO` uses a database feature called Table Monitoring, which needs to be enabled for all the tables. A procedure called `ENABLE_SCHEMA_MONITORING` has been provided to enable monitoring on all tables for a given schema or all Applications schemas.

Oracle E-Business Suite on Oracle Database 11g uses its improved `DBMS_STATS.AUTO_SAMPLE_SIZE` feature by default in the `FND_STATS` package for better `AUTO` sampling statistics gathering.

Gathering Statistics for the CBO

Oracle E-Business Suite provides concurrent programs that use the package `FND_STATS` to gather statistics for your applications database objects. For information on `DBMS_STATS`, see Oracle Supplied PL/SQL Packages Manual.

The following concurrent programs are available for collecting and maintaining statistics:

- Gather Table Statistics
- Backup Table Statistics
- Restore Table Statistics
- Gather Schema Statistics
- Purge `FND_STATS` History Records

Gather Table Statistics

The Gather Table Statistics program gathers the table statistics for the specified table. This program can optionally backup the existing statistics in the FND_STATTAB table before gathering new statistics. If the value of backup_flag is BACKUP, then FND_STATS exports the old statistics using dbms_stats.export_table_stats before gathering the new statistics. The exported data is stored in FND_STATTAB. If the value of backup_flag is anything other than BACKUP then the old table statistics are not saved. This program also gathers index statistics for the table by default. For a detailed description of the procedure used by this concurrent program, see: GATHER_TABLE_STATS Procedure.

Parameters

Owner Name	The owner of the table.
Table Name	The name of the table.
Estimate Percent	The sampling percentage. If left blank, a default value of 10 is used. The valid range is from 0 to 100.
Degree	The degree of parallelism to be used for gathering statistics. If a Degree is not provided, it defaults to the minimum of parallel_max_servers and cpu_count.
Partition Name	The name of the partition.
Backup Flag	The backup flag indicates whether to backup statistics. Set this flag to "BACKUP" to back up your statistics.
Granularity	The granularity of statistics to collect (only relevant for tables that are partitioned). Valid values are: <ul style="list-style-type: none">• DEFAULT - Gather global and partition-level statistics.• SUBPARTITION - Gather subpartition-level statistics.• PARTITION - Gather partition-level statistics.• GLOBAL - Gather global statistics.• ALL - Gather all (subpartition, partition, and global) statistics.
History Mode	This parameter controls the amount of history records that are created. Valid modes are LASTRUN, FULL and NONE. The default is LASTRUN. For an explanation of the different modes, please refer to the

GATHER_TABLE_STATS Procedure

Invalidate Dependent Cursors This flag indicates whether cursors dependent on the table being analyzed should be invalidated or not. This parameter is ignored if you are running a database prior to Oracle 9i Release 2 (9.2.x).

Backup Table Statistics

This concurrent program backs up the current statistics of the given table into the FND_STATTAB table. This program also backs up the related index and column statistics by default.

An identifier, commonly referred to as STATID, can be associated with the backup up statistics. This STATID allows you to restore a particular version of the statistics using the Restore Table Statistics concurrent program. Statistics for the same object can be backed up with different STATIDs. You can even backup different versions of the statistics for the same object by assigning different STATIDs.

For a detailed description of the procedure used by this concurrent program, see: BACKUP_TABLE_STATS Procedure.

Parameters

Schema Name	The name of the schema. The value ALL means all Oracle E-Business Suite schemas.
Table Name	The name of the table.
Statistics ID	An optional identifier to associate with these statistics within FND_STATTAB. The default STATID is BACKUP.
Partition Name	Name of the table partition. If the table is partitioned and if the partition name is NULL, then global and partition table statistics are exported.

Restore Table Statistics

This concurrent program allows you to restore the previously backed up table statistics for a given statistics identifier, commonly referred to as the STATID.

All index and column statistics associated with the specified table are restored as well.

For a detailed description of the procedure used by this concurrent program, see: RESTORE_TABLE_STATS Procedure

Parameters

Schema Name	The name of the schema. The value ALL means all Oracle
--------------------	--

E-Business Suite schemas.

Table Name	The name of the table.
Statistics ID	An optional identifier to associate with these statistics within FND_STATTAB. The default STATID is BACKUP.
Partition Name	Name of the table partition. If the table is partitioned and if the partition name is NULL, then global and partition table statistics are imported.

Gather Schema Statistics

This concurrent program gathers the specified schema level statistics.

Before gathering the statistics, this program can also create a backup of the current statistics, depending on the value of the Backup Flag. If for some reason, the earlier statistics need to be restored, that can be done using the Restore Schema Statistics concurrent program. The STATID used for this backup is NULL.

This program also creates histograms on the columns seeded in the FND_HISTOGRAM_COLS table.

For a detailed description of the procedure used by this concurrent program, see information on the GATHER_SCHEMA_STATS procedure.

Parameters

Schema Name	Schema for which statistics are to be gathered. Specify ALL for all Oracle E-Business Suite schemas (all schemas that have an entry in the FND_PRODUCT_INSTALLATIONS table).
Percent	The sampling percentage. If left blank, the default value of 10 is used. The valid range is from 0 to 100.
Degree	The degree of parallelism to be used for gathering statistics. If a Degree is not provided, it defaults to the minimum of parallel_max_servers and cpu_count.
Backup Flag	The backup flag indicates whether to backup statistics. Set this flag to BACKUP if you wish to back up the current statistics into the FND_STATTAB table. If NOBACKUP is used, then the GATHER_SCHEMA_STATS procedure will not backup the current statistics. This way the GATHER_SCHEMA_STATS procedure will run faster.
Restart Request ID	In the case where the Gather Schema Statistics run fails due to whatever reasons, the concurrent request can be

re-submitted and it will pick up where the failed run left off, if you provide the concurrent request_id of the failed run.

History Mode

This parameter controls the amount of history records that are created. The history records, stored in FND_STATS_HIST can be queried to find out when stats were gathered on a particular object and the amount of time it took to gather statistics on that object.

- Last Run - History records for each object are maintained only for the last gather statistics run. Each subsequent run will overwrite the previous history record for the object. This is the default behavior.
- Full - This mode does not overwrite any history information. History records are created for each run and are identified by the Request ID. If a Request ID is not provided, one is generated automatically. If this mode is used, the "Purge FND_STATS History Records " concurrent program should be run periodically to purge the FND_STATS_HIST table.
- None - This mode does not generate any history information. If this mode is used, the run cannot be restarted.

Gather Options

This parameter specifies how objects are selected for statistics gathering.

- GATHER : All tables and indexes of the schema **schemaname** are selected for stats gathering. This is the default.
- GATHER AUTO : Tables of the schema schemaname for which the percentage of modifications has exceeded modpercent are selected for statistics gathering. Indexes of these tables are selected by default. Table monitoring needs to be enabled before using this option.
- GATHER EMPTY : Statistics are gathered only for tables and indexes that are missing statistics.
- LIST AUTO : This option does not gather statistics. It only provides a listing of all the tables that will be selected for statistic gathering, if the GATHER AUTO

option is used.

- **LIST EMPTY** : This option does not gather statistics. It only provides a listing of all the tables that will be selected for statistics gathering, if the **GATHER EMPTY** option is used.

Modifications Threshold

Applicable only to **GATHER AUTO** and **LIST AUTO** options. This parameter specifies the percentage of modifications (with respect to the total rows) that have to take place on a table before it can be picked up for **AUTO** statistics gathering.

Invalidate Dependent Cursors

This flag indicates whether cursors dependent on the table being analyzed should be invalidated or not. By default, dependent cursors are invalidated. This parameter is ignored if you are running a database prior to Oracle 9i Release 2 (9.2.x).

Gather Column Statistics

This concurrent program should be used for gathering the Column Statistics, i.e. creating a histogram on a given column.

The procedure takes a backup into the **FND_STATTAB** table before gathering the statistics.

For a detailed description of the procedure used by this concurrent program, see: **GATHER_COLUMN_STATS** Procedure

Parameters

Table Owner

The owner of the table.

Table Name

The name of the table.

Column Name

The name of the column.

Estimate Percent

The sampling percentage. If left blank, a default value of 10 is used. The valid range is from 0 to 100.

Parallel Degree

The degree of parallelism to be used for gathering statistics. If a Degree is not provided, it defaults to the minimum of **parallel_max_servers** and **cpu_count**.

Bucket Size

The number histogram buckets.

Backup Flag

The backup flag indicates whether to backup statistics. Set

this flag to BACKUP if you wish to back up the current column statistics into the FND_STATTAB table. If left blank, it defaults to NOBACKUP.

Gather All Column Statistics

This concurrent program is obsolete.

Purge FND_STATS History Records

This program can be run to purge the history records from the FND_STATS_HIST table. This program should be scheduled to run periodically if statistics are being gathered with History Mode as FULL. You do not need to run this program if you gather statistics with History Mode as NONE or the default – LASTRUN.

Parameters

Purge Mode	The Purge Mode can take one of the two values: DATE or REQUEST. If the mode chosen is DATE, history records are purged based on the date range, otherwise, if it is REQUEST, records are purged based on the Request ID.
From Value	Start Date or Request ID
To Value	End Date or Request ID.

FND_STATS Package

The FND_STATS package provides procedures for gathering statistics for Oracle E-Business Suite database objects. It also provides procedures for backing up the current statistics into the table - FND_STATTAB, and restoring them back if desired. This package also allows users to specify the degree of parallelism. That helps speed up statistics gathering for large objects. FND_STATS can also maintain a history of its actions in a table called FND_STATS_HIST. The data in this table is used to provide restart ability, and can also be queried to find out the time taken to gather statistics on each object.

FND_STATS relies on the Oracle-supplied package DBMS_STATS to perform the actual statistics gathering. For more information on DBMS_STATS, refer to the Oracle database Tuning and Supplied Packages Reference manuals.

CREATE_STAT_TABLE Procedure

This procedure creates the table that is required for backing up the statistics.

There are two versions of this procedure. The first one does not need any arguments and creates the table with the default name - FND_STATTAB in the schema

corresponding to the FND product. The second version allows you to provide the schema name, table name and the tablespace for the statistics table.

Syntax

```
FND_STATS.CREATE_STAT_TABLE ;

FND_STATS.CREATE_STAT_TABLE (
  schemaname IN VARCHAR2,
  tablename  IN VARCHAR2,
  tblspcname IN VARCHAR2);
```

Parameters

schemaname	Name of the schema.
tablename	Name of the table.
tblspcname	Tablespace in which to create the statistics tables. If none is specified, then the tables are created in the user's default tablespace.

BACKUP_TABLE_STATS

This procedure backs up the statistics for the given table in the FND_STATTAB table. Setting cascade to TRUE results in all index and column statistics associated with the specified table to be stored as well. An identifier, commonly referred to as STATID, can be associated with the backup up statistics. This STATID allows you to restore a particular version of the statistics using the RESTORE_TABLE_STATS procedure.

Syntax

```
FND_STATS.BACKUP_TABLE_STATS (
  schemaname VARCHAR2,
  tablename  VARCHAR2,
  statid     VARCHAR2 DEFAULT 'BACKUP',
  partname   VARCHAR2 DEFAULT NULL,
  cascade    BOOLEAN  DEFAULT TRUE);
```

Parameters

schemaname	Name of the schema.
tablename	Name of the table.
statid	Optional identifier to associate with these statistics within FND_STATTAB.
partname	Name of the table partition. If the table is partitioned and if partname is NULL, then global and partition table statistics

are exported.

cascade

If TRUE, then column and index statistics for this table are also exported.

BACKUP_SCHEMA_STATS Procedure

This procedure can be used to backup the statistics for an entire schema. The statistics are backed up into the FND_STATTAB table. A different version can be stored by specifying a different statid. An identifier, commonly referred to as STATID, can be associated with the backup up statistics. This STATID allows you to restore a particular version of the statistics using the RESTORE_SCHEMA_STATS procedure.

Syntax

```
FND_STATS.BACKUP_SCHEMA_STATS (  
    schemaname VARCHAR2,  
    statid      VARCHAR2 DEFAULT NULL);
```

Parameters

schemaname

Name of the schema. ALL means all Oracle E-Business Suite schemas.

statid

Optional identifier to associate with these statistics within FND_STATTAB.

RESTORE_SCHEMA_STATS Procedure

This procedure restores statistics for the given schema, that were previously backed up in the FND_STATTAB table, into the dictionary. Statid can be provided to distinguish between different sets of statistics for the same object.

Syntax

```
FND_STATS.RESTORE_SCHEMA_STATS (  
    schemaname VARCHAR2,  
    statid VARCHAR2 DEFAULT NULL  
);
```

Parameters

schemaname

Name of the schema. ALL means all Oracle E-Business Suite schemas.

statid

Optional identifier to associate with these statistics within FND_STATTAB.

RESTORE_TABLE_STATS Procedure

This procedure restores statistics for the given table from the FND_STATTAB table for the given statid (optional) and transfers them back to the dictionary. Setting cascade to TRUE results in all index and column statistics associated with the specified table being imported also.

Syntax

```
FND_STATS.RESTORE_TABLE_STATS (  
    ownname  VARCHAR2,  
    tabname  VARCHAR2,  
    statid   VARCHAR2 DEFAULT NULL,  
    partname VARCHAR2 DEFAULT NULL,  
    cascade  BOOLEAN  DEFAULT TRUE,  
);
```

Parameters

ownname	Name of the schema.
tabname	Name of the table.
statid	Optional identifier to associate with these statistics within FND_STATTAB.
partname	Name of the table partition. If the table is partitioned and if partname is NULL, then global and partition table statistics are exported.
cascade	If TRUE, then column and index statistics for this table are also exported.

RESTORE_COLUMN_STATS Procedure

This procedure restores statistics for the given column from the FND_STATTAB table for the given statid (optional) and transfers them back to the dictionary. There are two versions of this procedure. One first one requires the table owner, table name and column name to be supplied. The second version restores the statistics for all the columns seeded in the FND_HISTOGRAM_COLS table.

Syntax

```
FND_STATS.RESTORE_COLUMN_STATS (  
  ownname  VARCHAR2,  
  tabname  VARCHAR2,  
  colname  VARCHAR2,  
  partname VARCHAR2 DEFAULT NULL,  
  statid   VARCHAR2 DEFAULT NULL  
);  
  
FND_STATS.RESTORE_COLUMN_STATS (  
  statid   VARCHAR2 DEFAULT NULL  
);
```

Parameters

ownname	Name of the schema.
tabname	Name of the table.
colname	Name of the column. Optional identifier to associate with these statistics within FND_STATTAB.
partname	Name of the table partition. If the table is partitioned and if partname is NULL, then global and partition table statistics are exported.
statid	Optional identifier to associate with these statistics within FND_STATTAB.

ENABLE_SCHEMA_MONITORING Procedure

This procedure should be used for enabling the Monitoring option for all tables in the specified schema. Monitoring option should be enabled before using the GATHER AUTO or LIST AUTO option of GATHER_SCHEMA_STATS. If the value of the schemaname argument is ALL, then the Monitoring option is enabled for all tables that belong to all schemas registered in Oracle E-Business Suite.

Syntax

```
FND_STATS.ENABLE_SCHEMA_MONITORING (  
  schemaname VARCHAR2 DEFAULT 'ALL');
```

Parameters

schemaname	Name of the schema for which Monitoring should be enabled.
-------------------	--

DISABLE_SCHEMA_MONITORING Procedure

This procedure should be used for disabling the Monitoring option for all tables in the specified schema. If the value of the schemaname argument is ALL, then the Monitoring option is disabled for all tables that belong to all schemas registered in Oracle E-Business Suite.

Syntax

```
FND_STATS.DISABLE_SCHEMA_MONITORING (  
    schemaname VARCHAR2 DEFAULT 'ALL');
```

Parameters

schemaname	Name of the schema for which Monitoring should be disabled.
-------------------	---

GATHER_SCHEMA_STATS Procedure

This procedure gathers statistics for all objects in a schema. Statistics are gathered with the granularity of DEFAULT. This procedure is also available through the concurrent program "Gather Schema Statistics." If this procedure fails at any time during operation, supplying the request ID for the request that failed can restart it. The request ID can be captured when the program is started from concurrent manager or can be queried from the FND_STATS_HIST table.

GATHER_SCHEMA_STATS cannot be executed directly in sqlplus because of an OUT parameter. The procedure GATHER_SCHEMA_STATISTICS has been provided for gathering schema statistics from the sqlplus prompt.

Syntax

```
FND_STATS.GATHER_SCHEMA_STATS (  
  schemaname          VARCHAR2,  
  estimate_percent    NUMBER DEFAULT NULL,  
  degree              NUMBER DEFAULT NULL,  
  internal_flag       NUMBER DEFAULT NULL,  
  Errors_OUT Error_Out,  
  request_id          NUMBER default null,  
  hmode               VARCHAR2 default 'LASTRUN',  
  options in          VARCHAR2 default 'GATHER',  
  modpercent          NUMBER default 10,  
  invalidate          VARCHAR2 default 'Y'  
);  
  
FND_STATS.GATHER_SCHEMA_STATISTICS (  
  schemaname          VARCHAR2,  
  estimate_percent    NUMBER DEFAULT NULL,  
  degree              NUMBER DEFAULT NULL,  
  internal_flag       NUMBER DEFAULT NULL,  
  request_id          NUMBER DEFAULT NULL,  
  hmode               VARCHAR2 DEFAULT 'LASTRUN',  
  options in          VARCHAR2 DEFAULT 'GATHER',  
  modpercent          NUMBER DEFAULT 10,  
  invalidate          VARCHAR2 DEFAULT 'Y'  
);
```

Parameters

schemaname	Schema to analyze. ALL means all Oracle E-Business Suite schemas.
estimate_percent	The sampling percentage. If a value is not provided, the default value of 10 is used. The valid range is from 0 to 100.
degree	The degree of parallelism to be used for gathering statistics. If a degree is not provided, it defaults to the minimum of parallel_max_servers and cpu_count.
internal_flag	The backup flag indicates whether to backup statistics. Set this flag to BACKUP if you wish to back up the current statistics into the FND_STATTAB table. If NOBACKUP is used, then the GATHER_SCHEMA_STATS procedure will not backup the current statistics. This way the GATHER_SCHEMA_STATS procedure will run faster.
errors	User defined Type for holding the Error messages .
Request_id	A request_id can be provided to identify the history records for a given statistics gathering run. This parameter is also used for providing restart ability. In case, a statistics gathering run fails due to whatever reasons, subsequent

submission can pick up where the failed run left off, if you provide the request_id of the failed run.

Hmode

This parameter controls the amount of history records that are created. The history records, stored in FND_STATS_HIST can be queried to find out when statistics were gathered on a particular object and the amount of time it took to gather statistics on that object.

LASTRUN - History records for each schema are maintained only for the last gather statistics run. Each subsequent run will overwrite the previous history record for the index. This is the default behavior.

FULL - This mode does not overwrite any history information. History records are created for each run and are identified by the Request ID. If a Request ID is not provided, one is generated automatically. If this mode is used, the "Purge FND_STATS History Records" concurrent program should be run periodically to purge the FND_STATS_HIST table.

NONE - This mode does not generate any history information. If this mode is used, the run cannot be restarted.

Options

This parameter specifies how objects are selected for statistics gathering.

GATHER - All tables and indexes of the schema <schemaname> are selected for stats gathering. This is the default.

GATHER AUTO - Tables of the schema schemaname for which the percentage of modifications has exceeded modpercent are selected for statistics gathering. Indexes of these tables are selected by default. Table monitoring needs to be enabled before using this option.

GATHER EMPTY - Statistics are gathered only for tables and indexes that are missing statistics.

LIST AUTO - This option does not gather statistics. It only provides a listing of all the tables that will be selected for statistic gathering, if the GATHER AUTO option is used.

LIST EMPTY - This option does not gather statistics. It only provides a listing of all the tables that will be selected for statistics gathering, if the GATHER EMPTY option is used.

Modpercent

Applicable only to GATHER AUTO and LIST AUTO

options. This parameter specifies the percentage of modifications (with respect to the total rows) that have to take place on a table before it can be picked up for AUTO statistics gathering.

Invalidate This flag indicates whether cursors dependent on the table being analyzed should be invalidated. By default, dependent cursors are invalidated. This parameter is ignored if you are running a database prior to Oracle 9i Release 2 (9.2.x).

Exceptions

ORA-20000: Schema does not exist or insufficient privileges.
ORA-20001: Bad input value.

GATHER_INDEX_STATS Procedure

This procedure gathers statistics for the specified index.

Syntax

```
FND_STATS.GATHER_INDEX_STATS (  
    ownname    VARCHAR2,  
    indname    VARCHAR2,  
    percent    NUMBER DEFAULT NULL,  
    partname   VARCHAR2 DEFAULT NULL,  
    backup_flag VARCHAR2 DEFAULT NULL,  
    hmode     VARCHAR2 DEFAULT 'LASTRUN',  
    invalidate VARCHAR2 DEFAULT 'Y'  
);
```

Parameters

ownname	Schema of index to analyze.
indname	Name of index.
percent	The sampling percentage. If left blank, the default value of 10 is used. The valid range is from 0 to 100.
partname	Partition name.
backup_flag	The backup flag indicates whether to backup statistics. Set this flag to BACKUP if you wish to back up the current column statistics into the FND_STATTAB table. If left blank, it defaults to NOBACKUP.
Hmode	This parameter controls the amount of history records that

are created.

LASTRUN - History records for each index are maintained only for the last gather statistics run. Each subsequent run will overwrite the previous history record for the index. This is the default behavior.

FULL - This mode does not overwrite any history information. History records are created for each run and are identified by the Request ID. If a Request ID is not provided, one is generated automatically. If this mode is used, the "Purge FND_STATS History Records" concurrent program should be run periodically to purge the FND_STATS_HIST table.

NONE - This mode does not generate any history information. If this mode is used, the run cannot be restarted.

Invalidate

This flag indicates whether cursors dependent on the index being analyzed should be invalidated. By default, dependent cursors are invalidated.

GATHER_TABLE_STATS Procedure

This procedure gathers table, column and index statistics. It attempts to parallelize as much of the work as possible. This operation does not parallelize if the user does not have select privilege on the table being analyzed.

Syntax

```
FND_STATS.GATHER_TABLE_STATS (  
    ownname    VARCHAR2,  
    tabname    VARCHAR2,  
    percent    NUMBER DEFAULT NULL,  
    degree     NUMBER DEFAULT NULL,  
    partname   VARCHAR2 DEFAULT NULL,  
    backup_flag VARCHAR2 DEFAULT NULL,  
    cascade    BOOLEAN DEFAULT TRUE,  
    granularity VARCHAR2 DEFAULT 'DEFAULT',  
    hmode      VARCHAR2 DEFAULT 'LASTRUN',  
    invalidate VARCHAR2 DEFAULT 'Y'  
);
```

Parameters

ownname	Owner of the table.
tabname	Name of the table.
percent	The sampling percentage. If left blank, the default value of

10 is used. The valid range is from 0 to 100.

degree	The degree of parallelism to be used for gathering statistics. If a degree is not provided, it defaults to the minimum of <code>parallel_max_servers</code> and <code>cpu_count</code> .
partname	Name of the partition.
backup_flag	The backup flag indicates whether to backup statistics. Set this flag to <code>BACKUP</code> if you wish to back up the current table statistics into the <code>FND_STATTAB</code> table. If left blank, it defaults to <code>NOBACKUP</code> .
cascade	When set to <code>TRUE</code> index statistics are gathered in addition to gathering statistics for the specified table. Index statistics gathering is not parallelized. Using this option is equivalent to running the <code>GATHER_INDEX_STATS</code> procedure on each of the table's indexes
granularity	The granularity of statistics to collect (only relevant for tables that are partitioned). Valid values are: <code>DEFAULT</code> - Gather global and partition-level statistics. <code>SUBPARTITION</code> - Gather subpartition-level statistics. <code>PARTITION</code> - Gather partition-level statistics. <code>GLOBAL</code> - Gather global statistics. <code>ALL</code> - Gather all (subpartition, partition, and global) statistics.
Hmode	This parameter controls the amount of history records that are created. <code>LASTRUN</code> - History records for each index are maintained only for the last gather statistics run. Each subsequent run will overwrite the previous history record for the index. This is the default behavior. <code>FULL</code> - This mode does not overwrite any history information. History records are created for each run and are identified by the Request ID. If a Request ID is not provided, one is generated automatically. If this mode is used, the "Purge FND_STATS History Records" concurrent program should be run periodically to purge the <code>FND_STATS_HIST</code> table. <code>NONE</code> - This mode does not generate any history information. If this mode is used, the run cannot be

restarted.

Invalidate

This flag indicates whether cursors dependent on the index being analyzed should be invalidated. By default, dependent cursors are invalidated.

GATHER_COLUMN_STATS Procedure

This procedure should be used for gathering the Column Statistics, that is, creating a histogram on a given column.

There are two versions of this procedure. The first one gathers statistics on all columns seeded in the FND_HISTOGRAM_COLS for the given appl_id. If NULL, all seeded histograms are created. The other version gathers column statistics for the specified column.

Syntax

```
FND_STATS.GATHER_COLUMN_STATS (
    appl_id      NUMBER DEFAULT NULL,
    percent      NUMBER DEFAULT NULL,
    degree       NUMBER DEFAULT NULL,
    backup_flag  VARCHAR2 DEFAULT NULL,
    Errors       OUT Error_Out,
    hmode        VARCHAR2 DEFAULT 'LASTRUN',
    invalidate   VARCHAR2 DEFAULT 'Y'
);

FND_STATS.GATHER_COLUMN_STATS (
    ownname      VARCHAR2,
    tablename    VARCHAR2,
    colname      VARCHAR2,
    percent      NUMBER DEFAULT NULL,
    degree       NUMBER DEFAULT NULL,
    hsize        NUMBER DEFAULT 254,
    backup_flag  VARCHAR2 DEFAULT NULL,
    partname     VARCHAR2 DEFAULT NULL,
    hmode        VARCHAR2 DEFAULT 'LASTRUN',
    invalidate   VARCHAR2 DEFAULT 'Y'
);
```

Parameters

appl_id	Application ID.
ownname	Owner of the table.
colname	Column name.
tablename	Table name.
partname	Name of the partition.

percent	The sampling percentage. If left blank, the default value of 10 is used. The valid range is from 0 to 100.
degree	The degree of parallelism to be used for gathering statistics. If a degree is not provided, it defaults to the minimum of parallel_max_servers and cpu_count.
hsize	Number of buckets in the histogram.
backup_flag	The backup flag indicates whether to backup statistics. Set this flag to BACKUP if you wish to back up the current column statistics into the FND_STATTAB table. If left blank, it defaults to NOBACKUP.
errors	User defined Type for holding the Error messages.
hmode	<p>This parameter controls the amount of history records that are created.</p> <p>LASTRUN - History records for each index are maintained only for the last gather statistics run. Each subsequent run will overwrite the previous history record for the index. This is the default behavior.</p> <p>FULL - This mode does not overwrite any history information. History records are created for each run and are identified by the Request ID. If a Request ID is not provided, one is generated automatically. If this mode is used, the "Purge FND_STATS History Records" concurrent program should be run periodically to purge the FND_STATS_HIST table.</p> <p>NONE - This mode does not generate any history information. If this mode is used, the run cannot be restarted.</p>
Invalidate	This flag indicates whether cursors dependent on the index being analyzed should be invalidated. By default, dependent cursors are invalidated.

GATHER_ALL_COLUMN_STATS Procedure

This procedure gathers column statistics, i.e. creates histograms on all columns that are seeded in the FND_HISTOGRAM_COLS, belonging to the specified schema .

Syntax

```
FND_STATS.GATHER_ALL_COLUMN_STATS (  
  ownname      VARCHAR2 ,  
  percent      NUMBER DEFAULT NULL,  
  degree       NUMBER DEFAULT NULL,  
  hmode        VARCHAR2 DEFAULT 'LASTRUN',  
  invalidate   VARCHAR2 DEFAULT 'Y'  
);
```

Parameters

ownname	Schema for which seeded histograms have to be created. ALL means all Applications schemas.
percent	The sampling percentage. If left blank, the default value of 10 is used. The valid range is from 0 to 100.
degree	The degree of parallelism to be used for gathering statistics. If a degree is not provided, it defaults to the minimum of parallel_max_servers and cpu_count.
Hmode	This parameter controls the amount of history records that are created. LASTRUN - History records for each index are maintained only for the last gather statistics run. Each subsequent run will overwrite the previous history record for the index. This is the default behavior FULL - This mode does not overwrite any history information. History records are created for each run and are identified by the Request ID. If a Request ID is not provided, one is generated automatically. If this mode is used, the "Purge FND_STATS History Records" concurrent program should be run periodically to purge the FND_STATS_HIST table. NONE - This mode does not generate any history information. If this mode is used, the run cannot be restarted.
Invalidate	This flag indicates whether cursors dependent on the index being analyzed should be invalidated. By default, dependent cursors are invalidated.

ANALYZE_ALL_COLUMNS Procedure

This procedure is obsolete.

LOAD_XCLUD_STATS Procedure

This procedure is obsolete.

PURGE_STAT_HISTORY Procedure

This procedure should be used for purging the unwanted history records from the `fnd_stats_hist` table. There are two versions of this procedure. The first one takes in a range of request ids and deletes all history records that fall within that range. The second version takes a range of dates as arguments and all the history records falling in-between that range are deleted. The delete takes place as an autonomous transaction.

Syntax

```
FND_STATS.PURGE_STAT_HIST (
    From_req_id NUMBER,
    To_req_id   NUMBER);

FND_STATS.PURGE_STAT_HIST(
    Purge_from_date VARCHAR2,
    Purge_to_date   VARCHAR2);
```

Parameters

<code>from_req_id</code>	Start Request ID.
<code>to_req_id</code>	End Request ID.
<code>purge_from_date</code>	Start Purge Date.
<code>purge_to_date</code>	End Purge Date.

CHECK_HISTOGRAM_COLS Procedure

For a given list of comma-separated tables, this procedure checks the data in all the leading columns of all the non-unique indexes of those tables and determines if histograms need to be created for those columns. The algorithm for this procedure is:

```
select
decode(floor(sum(tot)/(max(cnt)*FACTOR)),0,'YES','NO') HIST
from (select count(col) cnt , count(*) tot
from tab sample (PERCENT)
where col is not null
group by col);
```

The decode statement determines whether a single value occupies 1/FACTOR or more of the sample PERCENT.

If `sum(cnt)` is very small (a small non-null sample), the results may be inaccurate. A `count(*)` of at least 3000 is recommended. The procedure is run from a SQL prompt after setting the server output on.

Syntax

```
FND_STATS.CHECK_HISTOGRAM_COLS (  
  tablelist VARCHAR2,  
  factor    NUMBER DEFAULT 75,  
  percent   NUMBER DEFAULT 10,  
  degree    NUMBER DEFAULT NULL);
```

Parameters

tablelist	A comma separated list of tables. It should be of the form schema.tablename. A wildcard in the tablename is also allowed. For example, tablelist=>'oe.so%head% , pa.pa_exp% , ar.ra_customers'. The owner part is mandatory.
factor	The factor for calculating the histograms.
percent	Sample percent.
degree	Degree of parallelization.

VERIFY_STATS Procedure

For a given list of comma-separated tables, or for a given schema name, this procedure reports the statistics in the data dictionary tables for the tables, indexes, and histograms.

Syntax

```
FND_STATS.VERIFY_STATS (  
  schemaname  VARCHAR2 DEFAULT NULL,  
  tablelist   VARCHAR2 DEFAULT NULL,  
  days_old    NUMBER DEFAULT NULL,  
  column_stat BOOLEAN DEFAULT FALSE);
```

Parameters

schemaname	The name of a schema. If schemaname is NULL (which is the default), then the procedure reports on the given list of tables.
tablelist	A comma-separated list of tables. If the tablename is not of the form <schema>.<tablename> then the schema is the value of the schemaname parameter. If the tablelist is NULL (the default), then the procedure reports on all the tables for the specified schemaname.
days_old	The procedure only reports those tables whose statistics are older than the days_old number of days. The default is

NULL, which means the procedure will report on all the tables.

column_stat

If TRUE, the procedure reports column statistics for the export_table_stats table. The default is FALSE.

Using and Administering Oracle E-Business Suite Secure Enterprise Search

Overview of Oracle E-Business Suite Secure Enterprise Search

Oracle E-Business Suite Secure Enterprise Search is a centralized, secure search vehicle with consistent user interfaces throughout the Oracle E-Business Suite. By leveraging Oracle Secure Enterprise Search (SES), Oracle E-Business Suite Secure Enterprise Search enables a powerful keyword search on applications content in a faster, user-friendly way without compromising on the security and context sensitive information.

Before users can search on applications content, searchable objects must be set up first, constructed with secure context, and indexed into a full text search engine by Oracle SES in order to be ready for query. To accomplish this goal, Oracle E-Business Suite Secure Enterprise Search uses a flexible mechanism to help analyze these searchable objects, group related objects into categories, and build security rules around them for easier, secure search and fast result display.

In fact, searchable objects are business objects that are made available for text search. For example, a purchase order is a searchable object that can be defined as a set of searchable properties or business attributes along with its relationship to other searchable objects. This abstraction allows searchable objects to be bound to different context at run time and grouped into searchable categories.

Searchable objects are created with searchable attributes. These attributes allow the objects to be indexed, applied with security rules, and displayed with structured search results. Before users query, a search administrator grants appropriate data access privileges to users to secure application sensitive data from unauthorized access before deploying these objects to an Oracle SES instance.

At crawl time, the Oracle SES search engine starts a crawling job for a specific business object type. Based on a object type, searchable business objects or attributes get retrieved, indexed, and stored in the Oracle SES index store.

At query time, when a user performs a search through the centralized user interface, he

or she is actually searching against a preindexed store which contains numerous objects or metadata that has been preprocessed with indexes at crawl time. The search engine queries the results enforced by security rules and constructs the hits returning as search results displayed to the user.

Key features of Oracle E-Business Suite Secure Enterprise Search include:

- A centralized global search capability provides user rich experience of searching text across the entire Oracle E-Business Suite.
- It allows you to search structured and unstructured data like attachment through business categories and further narrow down your search results by business entities or attributes.
- Security context defined at multiple levels controls the accessibility of application sensitive data to only authorized users.
- It leverages extensive search capabilities provided by Oracle SES to search application content in a secure and user-friendly way.
- It provides pluggable search region capability which allows a specific searchable object to be embedded in a page for context sensitive search.
- A search administrator can proactively control and manage crawling schedules and statuses using the administrative pages.
- It provides multiple language support allowing application users to perform internationalized searches.

To have a better understanding of Oracle E-Business Suite Secure Enterprise Search, the following topics are discussed in this section:

- Terms and Definitions, page 20-2
- Architecture Overview, page 20-4
- Design Time, page 20-6
- Crawl Time, page 20-7
- Query Time, page 20-13

Oracle E-Business Suite Secure Enterprise Search Related Terms and Definitions

To better understand and administer Oracle E-Business Suite Secure Enterprise Search, this section provides relevant terminologies and their definitions used in Oracle E-Business Suite Secure Enterprise Search.

Searchable Objects

Searchable objects are business objects that are made available for text search; they are used in an abstract way for exposing business data to search engines. For example, a purchase order as a searchable object would be defined as a set of searchable properties and its relationship to other searchable objects.

Search Category

Related searchable objects can be grouped into a search category and it is also called a searchable group.

Oracle E-Business Suite Secure Enterprise Search leverages Role-Based Access Control (RBAC) model to associate searchable groups with permission sets and grant the group access privileges to authorized users.

Search Context

The binding information could be specific to a search engine. In order to make the search service open, Oracle E-Business Suite Secure Enterprise Search needs to abstract out the search engine internals and makes search engine a service that can be replaced by one another at the deploy time.

Search context is an application within which search services will be provided for searchable objects.

Search Engine

Search engine is an application or service that encapsulates the need of text search on a resource. It uses a number of well-defined sub modules to perform the necessary tasks. For Oracle E-Business Suite Secure Enterprise Search, Oracle SES is the search engine that makes search service feasible.

Crawler

Crawlers are software agents used by a search engine to retrieve content for a given data source.

Indexer

An indexer is a software module that is used by a search engine to create an index from each crawled document.

Once indexes are created for a particular data source, they are available for search through a set of Web Service APIs (Searcher interface).

Searcher

A searcher is a software module that allows external users to query into pre crawled and indexed stores. It is responsible for matching keywords and predicates to

documents, and then return them to the user.

Security Plug-in

To help protect unauthorized access to application information, security plug-in is used to enforce search security at the object level. Security plug-in is a Java class that implements the security methods to generate the access control list (ACL) for a document and to fetch Security Keys for a user.

An ACL is a list of permissions attached to an object specifying who or what is allowed to access the object and what operations are allowed to be performed. Oracle SES authorization plug-in works on the basis of the ACL-based security model and Security Keys for a document to authorize users or revoke the access to a search result.

User Authorization Cache (UAC)

This Oracle SES feature provides a framework allowing the Security Keys for a particular user, a specific data source, or a search object in Oracle E-Business Suite can be cached in Oracle SES.

By leveraging this framework from Oracle SES, when a user performs a search, the UAC is first looked up for the availability of the Security Keys for that user. If the keys are not found, then the Security Keys will be fetched synchronously during the query.

Query Rewrite

Query rewrite is a feature offered through a plug-in component that can rewrite the query to reflect current user context such as security before the query is sent to a search engine.

Data Security

Data security is a generic authorization model used by many applications within the Oracle E-Business Suite. It controls what users can see on application data through security grants.

Function Security

Function security is the basic access control in Oracle E-Business Suite. It restricts user access to individual menus and menu options within the system regardless of which application data in the row.

Oracle E-Business Suite Secure Enterprise Search uses the function security feature to guard the application content access through the menus and responsibilities assigned to each application user.

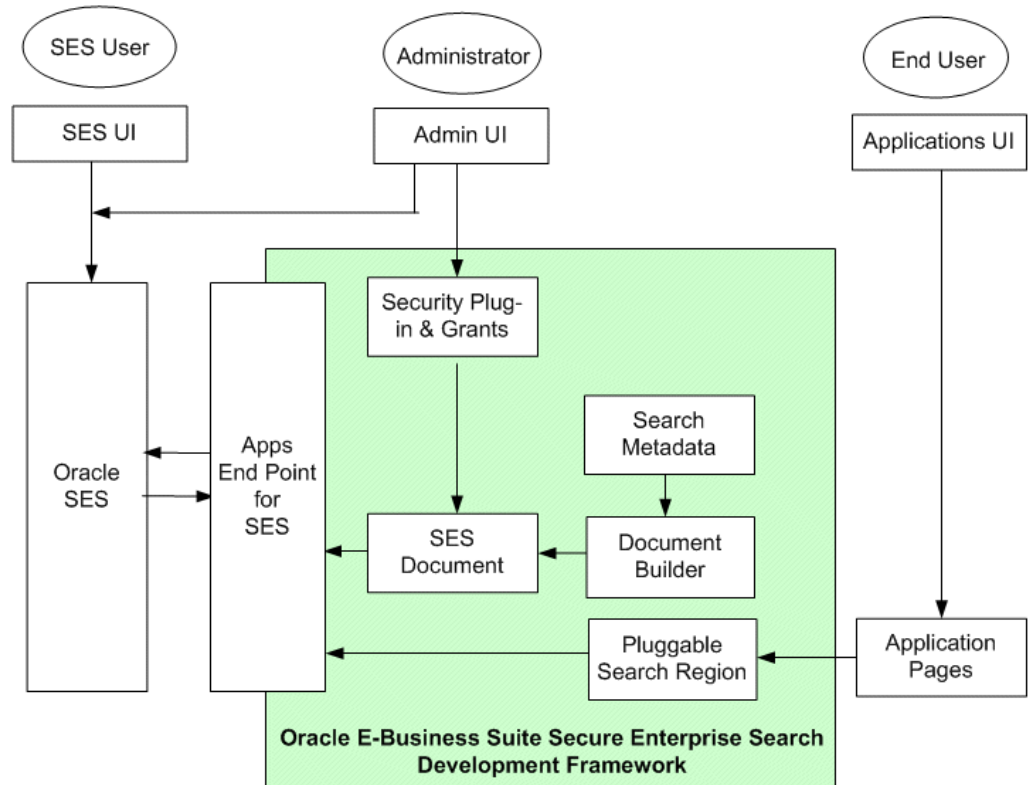
Architecture Overview

Oracle E-Business Suite Secure Enterprise Search development framework establishes

searchable objects from search metadata. This search metadata is then used during crawl time to conduct searches and store data, and used during query time to qualify results.

The following architecture diagram illustrates how metadata is used in defining searchable objects, and the interaction between Oracle E-Business Suite Secure Enterprise Search and Oracle SES:

Architecture Diagram



Business objects with searchable attributes become searchable metadata. Oracle E-Business Suite Secure Enterprise Search utilizes SES Document Builder to construct this searchable metadata or object which may contain complex business structure into a flattened searchable document. This document is also known as SES Document.

A search administrator creates security grants through roles or responsibilities, and necessary security plug-ins to secure searchable objects.

The search administrator or system administrator configures the necessary Oracle SES proxy parameters and setup tasks both in Oracle E-Business Suite Secure Enterprise Search and Oracle SES. This enables Oracle SES to crawl Oracle E-Business Suite (EBS), and Oracle E-Business Suite to query Oracle SES.

When an application user performs a search through application interfaces, a query is executed by invoking a search against a preindexed search store in SES.

Design Time


During the design phase, searchable objects with searchable attributes are created in Search Modeler and loaded to Oracle E-Business Suite as search metadata. These attributes allow searchable objects to be indexed, applied with security rules, and displayed with structured search results.

The search administrator grants appropriate data access privileges to users through roles or responsibilities to secure application sensitive data from unauthorized access before or after deploying objects to Oracle SES; the Oracle SES administrator then manages crawling schedules so that deployed data sources can be crawled for a specific object type and indexed.

Note: Once searchable objects are deployed, crawling schedules are automatically created along with data sources in Oracle SES. For more information on how to manage crawling schedules, see *Administering Crawls in Oracle SES*, page 20-72.

The following diagram illustrates the interaction flow during the design time:

Design Time Process Diagram



1. The search administrator creates security grants to users through roles or responsibilities before or after deploying searchable objects to Oracle SES as data sources.
2. The Oracle SES administrator manages crawling schedules that contain data sources.
3. The crawler manager picks the data source for crawl based on the schedules.

For more information about how to create searchable objects in Search Modeler, see *Creating Searchable Objects, Oracle E-Business Suite Search Modeler User's Guide* available from My Oracle Support Knowledge Document 781366.1, Search Modeler 1.1 for Oracle E-Business Suite Readme.

Crawl Time

To produce satisfying search results in a timely fashion, crawling and indexing are essential tasks to a successful search. At crawl time, crawling is done by several distributed crawlers. Oracle SES crawler is a Java process activated by a set schedule. When activated, the crawler spawns a configurable number of processor threads that

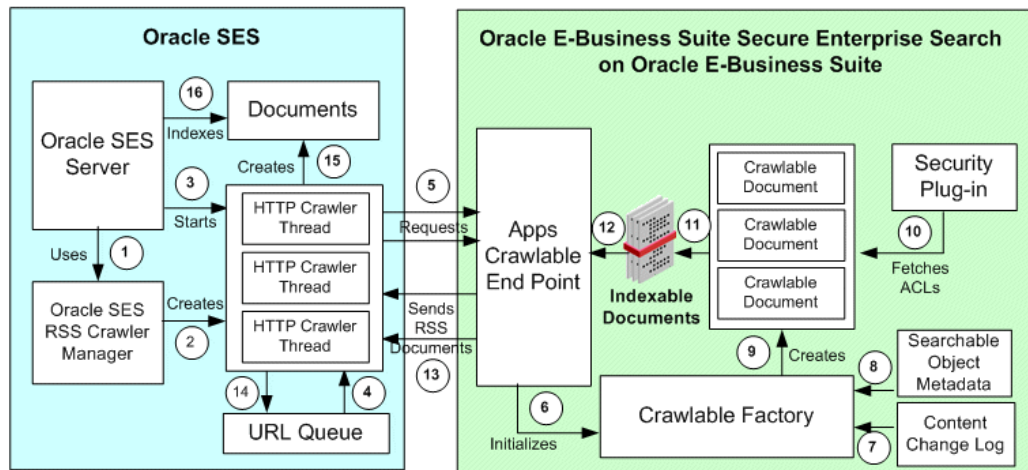
fetch information from various sources and index the documents. This index is used for searching sources.

Some crawlers are designed to crawl Oracle E-Business Suite users and provide user documents to Oracle SES. Oracle SES in turn invokes its authorization plug-in to generate document security access keys for each user crawled pertaining to the Oracle E-Business Suite source type and caches these keys for the authorized users and specific searchable objects or data sources. When a user performs a search, these previously cached security access keys will be used which provides a quick search result with security enforced.

Crawl Time for Indexable Documents

After searchable objects are deployed to Oracle SES as data sources contained in crawling schedules, Oracle SES starts crawling jobs in the Oracle E-Business Suite. A "crawlable" Oracle E-Business Suite means a secure end point that has been made crawlable to Oracle SES. This allows application data to be crawled and indexed into an Oracle SES store. The following diagram illustrates the interaction flow of Oracle SES crawler tasks:

Crawl Time Interaction Diagram for Indexable Documents



1. Oracle SES initializes RSS Crawler Manager.
2. Oracle SES Crawler Manager spawns and initializes a preconfigured number of crawler threads.
3. Oracle SES Starts the crawlers.

Note: The crawler maps links and analyzes relationships. Whenever the crawler encounters embedded non-HTML, or non-textual documents during the crawling, it automatically

detects the document type and filters and indexes the document.

4. Crawler threads pick up crawlable URLs from the URL Queue. URL Queue is populated using controlFeed mechanism as described in step 5.
5. Crawler threads contact Oracle E-Business Suite Crawling End Point, which is a servlet registered in `oafm` container.

The requests come as post requests with URL parameters as in `http(s)://<ebs apache host>:<web host>/webservices/AppSearch/[ConfigFeed | ControlFeed | DataFeed]/Search Object Name?user<ebs user having FND_SEARCH_CRAWLER resp>&password=<password>`.

Note: ConfigFeed and ControlFeed are crawling mechanisms to generate crawlable URLs in multiple batches of preconfigured sizes, so that crawling can proceed in parallel. These are used to generate the initial "URL Queue" in Oracle SES.

DataFeed is the actual crawling request, which has been illustrated in the diagram.

6. Once the Oracle E-Business Suite Crawlable End Point receives the crawling requests, it initializes the Crawlable Factory whose purpose is to fetch the content from Oracle E-Business Suite database.

Please note that Crawlable Factory is also responsible for splitting the original application content large data set into smaller work units through AD Parallel Update package, and then crawling the units in parallel by using the multi-thread crawling mechanism provided by Oracle SES.

Note: The Crawlable Factory is the place where an initial crawl taken place. The initial crawl refers to the first time a searchable object is crawled.

7. Content change log provides application changes that are indexed to the Crawlable Factory.
8. Search metadata is loaded to the Crawlable Factory.
9. Crawlable Factory creates crawlable documents, which conform to some schema provided by the indexing vendor.
10. While creating indexable documents, the Access Control List (ACL) is fetched for each document using the search plug-in associated with the searchable object definition. The `getAcl()` and `getSecureAttrAcl()` methods of the search

plug-in are invoked to generate the ACLs.

For more information about security plug-in, see *Search Security Plug-ins*, page 20-43.

11. Documents are ready to be consumed by a search/indexing engine.
12. Oracle E-Business Suite crawler threads pick the documents.
13. The indexable documents which are in the form of a RSS feed are passed to Oracle SES through the Oracle E-Business Suite End Point URL in response to the crawling request mentioned in step 5.

These documents conform to Oracle SES crawlable schema and should have following information:

- Metadata
 - Content to be indexed
 - Dependent document URLs (such as actionable links, attachments, or related documents or links)
14. On retrieving the document, Oracle SES indexing engine analyzes the RSS feed received from Oracle E-Business Suite and places the neighboring URLs into the URL Queue.

Typically the neighboring URLs in Oracle E-Business Suite are the attachment fetch URLs.
 15. Oracle SES indexing engine transforms the documents in Oracle SES readable format by extracting keywords.
 16. Finally indexing process indexes the documents.

Indexed documents are stored in the precrawled index store in Oracle SES.

Crawl Time for User Authorization Cache Source

To reduce the search response time of synchronously fetching Security Keys for an authorized user during user query, cached Security Keys for a particular user and a searchable object or data source are precrawled and stored in Oracle SES and then used directly at query time. This solution by using cached access keys to authorize a document access privilege for a user at query time is leveraged from Oracle SES User Authorization Cache (UAC) feature. For more information about this feature, see *User Authorization Cache*, page 20-63.

Note: To protect sensitive application data from unauthorized access,

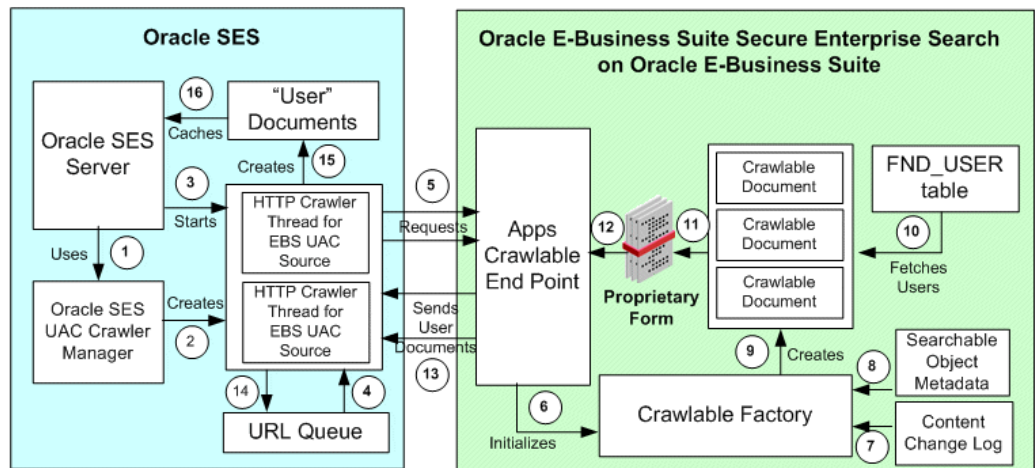
Access Control List (ACL) and Security Keys are generated through a security plug-in that is attached to a searchable object to enforce security at the object level.

An ACL is a list of permissions attached to an object specifying who or what is allowed to access the object and what operations are allowed to be performed. Security Keys are generated for a user to match the prebuilt ACLs to access specific documents based on user privileges.

For more information about security plug-in, see Search Security Plug-ins, page 20-43.

The following diagram illustrates the interaction flow of Oracle SES crawler tasks for Oracle E-Business Suite "User Authorization Cache" (UAC) source:

Crawl Time Interaction Diagram for User Authorization Cache Source



This interaction diagram for UAC source is similar to the crawling tasks for indexable documents. The major differences between these two are highlighted as follows:

- UAC Crawler Manager is initialized to preconfigure the crawler threads specifically for Oracle E-Business Suite UAC source type (step 2).
- While creating documents for User Crawl, crawlable instances use FND_USER, which is the source of Oracle E-Business Suite users. This holds true for Single Sign-On integrated Oracle E-Business Suite instances as well (step 10).
- Oracle E-Business Suite End Point URL responds to the crawling request mentioned in step 5 by sending the user documents, which are XML-based documents conforming to a mutually agreed upon schema.

Oracle E-Business Suite calculates the total number of users satisfying user list criteria mentioned in Oracle SES. If the number is less than the profile value defined

in the *FND: Search Crawl Batch Size* profile option, it directly gives user list to Oracle SES. Otherwise, it gives a 'User Control Feed' which in turn is used again by Oracle SES to create final user list feeds. (step 13).

- On retrieving the document, Oracle SES indexing engine analyzes the XML feed received from Oracle E-Business Suite and places the neighboring URLs into the URL Queue (step 14).
- Finally Oracle SES engine caches the Oracle E-Business Suite users. The precrawled UAC sources are stored in Oracle SES (step 16).

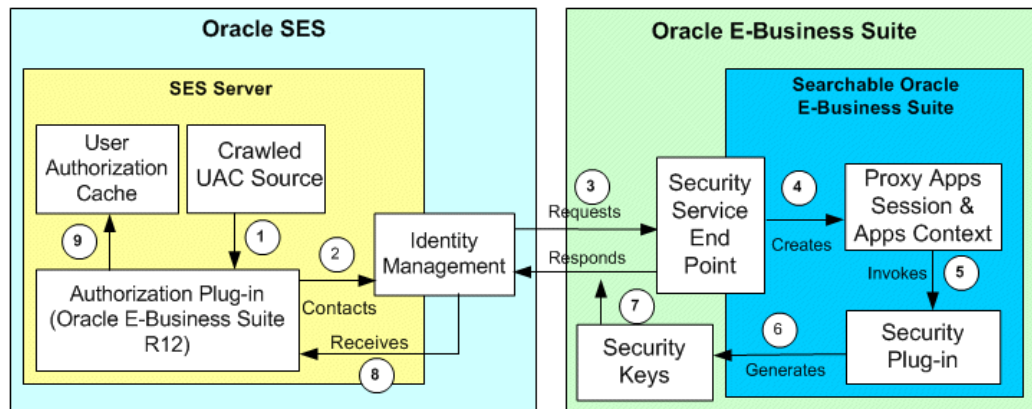
This crawling process generates a list of Oracle E-Business Suite users for whom the Security Keys need to be cached in Oracle SES for the predefined "User Authorization Cache" source type.

Fetching Security Keys Offline

In order to provide quick search results back to a user and eliminate possible time-outs of fetching Security Keys simultaneously during query due to complex application logic of deriving the keys, by leveraging the User Authorization Cache framework from Oracle SES, the user Security Keys can be generated as an offline process.

The following diagram illustrates the interaction flow for fetching Security Keys offline:

Interaction Diagram for Fetching Security Keys Offline



1. The crawled UAC source retrieves users and associated data sources from Oracle E-Business Suite.

Based on each user and data source, the authorization plug-in is invoked

2. Authorization plug-in contacts Identity Management to initiate the Security Key fetch process.
3. Identity Management plug-in sends a request to the Oracle E-Business Suite Security Service End Point to fetch the Security Keys for an Oracle E-Business Suite

user and a data source (i.e. search object in Oracle E-Business Suite).

The request is in the form: `http://<ebs server>:<port>/AppSearch/SecurityService?user=<proxy user>&password=<proxy password>`.

An XML message containing the user for whom the Security Keys are requested and the search object name is posted.

4. Security establishes the proxy session and applications context. The credential is verified for the proxy user name and password, which Oracle SES posts with the request.

The session is trusted or updated on behalf of the actual search user for whom the Security Keys have been requested.

5. The search plug-in is invoked by the Security Service End Point to generate the Security Keys. It is executed in the same proxy session.
6. The `getSecurityKeys()` and `getSecureAttrKeys()` methods are executed to generate the Security Keys for the proxy context.

Since the context is always incomplete, security plug-ins have to be aware of such scenarios.

7. Security Service End Point responds to the request mentioned in step 3 by sending the Security Keys for the requested user.
8. The authorization plug-in receives the Security Keys.
9. Security Keys are cached for a given user and a specific search object or data source.

Query Time

When an application user performs a search from the centralized search user interfaces, the user actually queries from a preindexed store in Oracle SES.

It is important to note that searchable group security rule and search plug-in security are enforced for a user query. For the searchable group security, not every searchable group can be seen or displayed to a user. Only those who have the group access privileges can find the group names displayed from the list of values for search selection. For search plug-in security, it can be used at crawl time and query time to fetch ACLs and generate Security Keys to protect unauthorized access to application data.

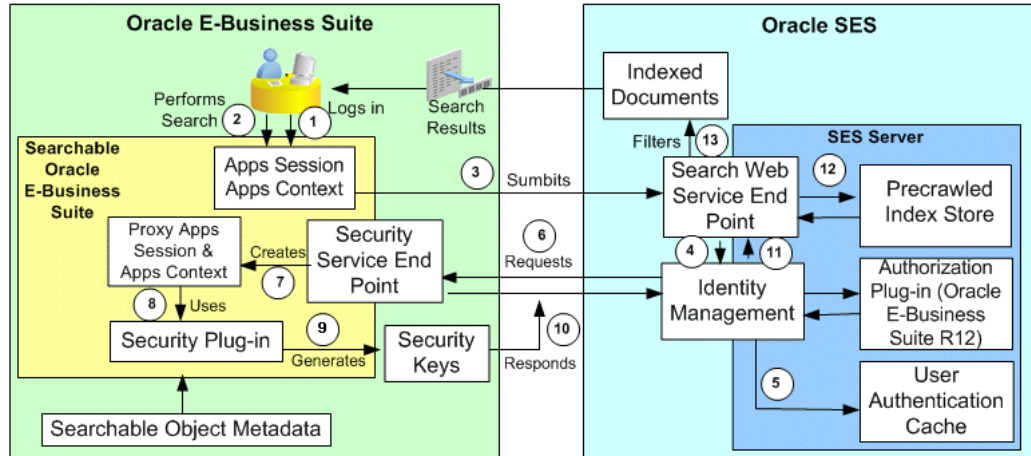
For information on how to secure searchable objects through searchable group security and security plug-in, see *Securing Searchable Objects*, page 20-38. For information on how to perform a search, see *Oracle E-Business Suite User's Guide*.

Query from Oracle E-Business Suite

You can perform a search from the Oracle E-Business Suite centralized search user interfaces.

The following diagram illustrates the query time interaction flow when performing a search through Oracle E-Business Suite:

Query Time Interaction Flow from Oracle E-Business Suite



1. A user logs on to Oracle E-Business Suite. A proxy session is created along with the initialization of applications context for the user.

The applications context may be incomplete at this stage depending on whether the user has selected a responsibility or not.

2. The user accesses the Oracle E-Business Suite Secure Enterprise Search toolbar and submits a search query within the same session and context.

3. The query is submitted to the Oracle SES client APIs, which are hosted within Oracle E-Business Suite. The Oracle SES client APIs in turn make Web service calls to the Search Web Service End Point published by Oracle SES server.

The Web service call includes the search keywords, filters if any, and user information amongst the most important parameters.

4. Once the search service is invoked, Oracle SES contacts the Identity Manager.

Identity Management is set up as part of the configuration for the integration between Oracle E-Business Suite and Oracle SES. Oracle SES has specific identity manager for Oracle E-Business Suite Release 12. This Identity Manager configuration needs Oracle E-Business Suite Security Service End Point and a proxy application user name and password to establish a proxy session.

For setup configuration for Oracle SES integration, see Setting Up Oracle E-Business

Suite Secure Enterprise Search for Oracle SES Integration, page 20-31.

5. If a User Crawler initiates at the crawl time, the Security Keys for a user, data source, or searchable object can be retrieved offline and cached in Oracle SES.

Oracle SES first looks up the Security Keys for the object and logged-in user in User Authentication Cache (UAC).

- If a match is found and the cache is usable, proceed to Step 12.
- If there is no match found, proceed to the next Step 6.

For more information on this feature, see User Authorization Cache, page 20-63.

6. Identity Manager requests Security Key information for the search user from the Security Service End Point. The Security Service End Point is registered as a servlet in `oafm` container.
7. Once the Security Service End Point receives a request for Security Keys, it initializes a proxy session. The proxy username/password credential is verified for the request. The session is then trusted or updated on behalf of the actual search user for whom the Security Keys have been requested.

Note: The proxy applications context may be incomplete since the responsibility information may or may not be there. Therefore, a special plug-in mechanism is provided to create the complete context information. For more information about the plug-in mechanism, see Understanding Security Logic and General Plug-in Mechanism, page 20-47.

8. The search plug-in is invoked in the same proxy session by the Security Service End Point to generate the Security Keys.
9. The `getSecurityKeys()` and `getSecureAttrKeys()` methods of the search plug-in are executed to generate the Security Keys for the proxy context.
Since the context may be incomplete, security plug-ins have to be aware of such scenarios.
10. The Security Service End Point responds to the request mentioned in Step 5 by sending the Security Keys for the search user. The request-response happens over HTTP protocol.

Oracle SES ensures that it does not wait indefinitely for the response to complete by setting a time-out on the request.

Note: The time-out value is configurable. This is done to ensure

responsiveness of the overall search solution.

11. The search service receives the authorization keys from Identity Management.
12. Search service retrieves indexed documents from the index store as per the search criteria given.
13. Indexed documents are filtered by Oracle SES after applying the Security Keys/Authorization Keys. This way, only the authorized documents are retrieved for the search user. The filtered indexed documents are returned to the query user for viewing and further action.

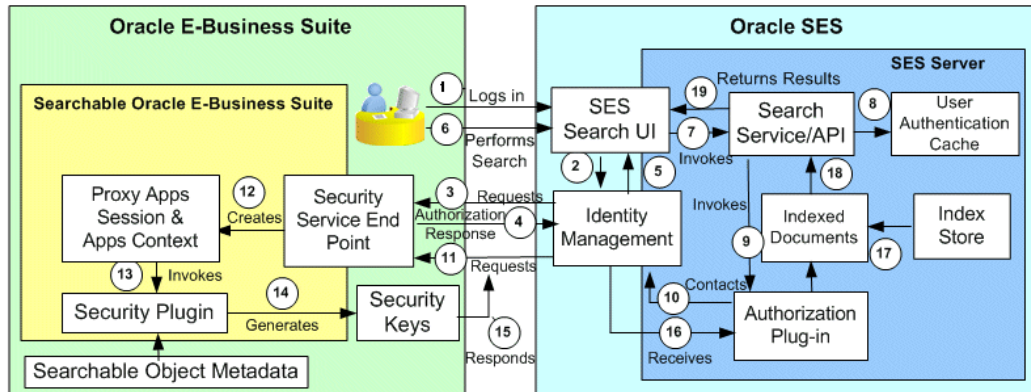
Query from Oracle SES

If an Oracle E-Business Suite user tries to log on and performs a search through the Oracle SES search user interface instead, the user's login credentials need to be authenticated first. At this stage, user login validation does not require any search plug-in.

Once the login is authenticated, the user can perform a search in the Oracle SES search UI with similar query time architecture as query in the Oracle E-Business Suite.

The following diagram illustrates the query time interaction flow when performing a search through Oracle SES:

Query Time Interaction Flow from Oracle SES



1. An Oracle E-Business Suite user attempts to sign in to Oracle SES Search UI using Oracle E-Business Suite username and password.
2. Oracle SES contacts Identity Manager to verify the user credentials.
3. Identity Manager sends an authentication request to the Security Service End Point of Oracle E-Business Suite. The request is sent over HTTP protocol.

It is typically of the form `http(s)://<ebs server>:<port>/webservices/AppSearch/SecurityService`. An XML message containing the exact authentication service requested is posted.

4. The Security Service End Point validates the login credentials and responds to the request by sending another XML message.
5. Identity Manager parses the response message and sends the success or failure response to the Oracle SES Search UI.
6. After successful login, the user submits a search query. It may consist of keywords, filters, and other search criteria.
7. Upon receiving the search request, Oracle SES invokes an appropriate search service or API.
8. If a User Crawler initiates at the crawl time, the Security Keys for a user, data source, or searchable object can be retrieved offline and cached in Oracle SES. Oracle SES first looks up the Security Keys for the object and logged-in user in User Authentication Cache (UAC).
 - If a match is found and the cache is usable, proceed to Step 17.
 - If there is no match found, proceed to the next Step 9.

For more information on this feature, see *User Authorization Cache*, page 20-63.

9. The search service in turn invokes the authorization plug-in to get the Security Keys for the current search user.
10. Authorization plug-in contacts the Identity Manager to fetch the Security Keys.
11. Identity Manager requests Security Key information for the search user from the Security Service End Point. This is done over HTTP. The request message contains the proxy username and password, which is stored in Oracle SES as part of the configuration.
12. Security establishes the proxy session and applications context. The proxy username/password credential is verified for the request from Oracle SES. The session is then trusted or updated on behalf of the actual search user for whom the Security Keys have been requested.

Note: Please note that the proxy applications context is always incomplete since the responsibility information may not be there while logging into Oracle SES. This is the major difference between searching from Oracle E-Business Suite and searching from Oracle

SES Search user interface.

13. The search plug-in is invoked in the same proxy session by the Security Service End Point to generate the Security Keys.

14. The `getSecurityKeys()` and `getSecureAttrKeys()` methods of the search plug-in are executed to generate the Security Keys for the proxy context.

Since the context is always incomplete, security plug-ins have to be aware of such scenarios.

15. The Security Service End Point responds to the request mentioned in Step 5 by sending the Security Keys for the search user. The request-response happens over HTTP protocol.

Oracle SES ensures that it does not wait indefinitely for the response to complete by setting a time-out on the request.

Note: The time-out value is configurable. This is done to ensure responsiveness of the overall search solution.

16. The authorization plug-in receives the Security Keys for the query user.

17. The Oracle SES search service or API retrieves indexed documents from index store, matching the search keywords and filters.

18. Indexed documents are filtered by the Security Keys retrieved for the query user.

19. Filtered search results are returned back to the query user.

Using Oracle E-Business Suite Secure Enterprise Search

By leveraging Oracle Secure Enterprise Search (SES), Oracle E-Business Suite Secure Enterprise Search provides a centralized search capability allowing you to use text search on applications content without compromising on the security and context sensitive information.

Through a secure text search, not only can you enjoy the immediate search results with accuracy, but you can also easily refine the search by narrowing down the results or tracing back to original results in an efficient way.

Note: Oracle E-Business Suite Secure Enterprise Search does not replace all the Search pages, instead it is deployed where it makes the most sense in searching for applications content. See Querying and Viewing

Data, Oracle E-Business Suite User's Guide for other search pages.

The screenshot displays the Oracle E-Business Suite Enterprise Search interface. At the top, the Oracle logo and 'E-Business Suite' are visible, along with navigation links like Navigator, Favorites, Diagnostics, Home, Logout, Preferences, Help, and Personalize Page. The search bar contains the text 'oracle' and a 'Go' button. Below the search bar, the 'Apps Search Results page' is shown. The interface is divided into several regions: 'Expand By' (Keyword: oracle, Group: Object Type Interfaces, Object: Object Type Interfaces, Filter: IrepName), 'Filter Region' (ContractParty: ECE_POO, Search With Filter button), 'Narrow By' (Object Type Interfaces, IrepName, DisplayName, ShortDescription, ProductName, Status, InterfaceType), and 'Contracts PO Header Search' results. The results are sorted by Relevance and show three items: Purchase Order Header View (40%), Purchase Order Line View (40%), and Purchase Order Shipment View (40%). The page number '1' is displayed at the bottom of the results section.

From the Oracle E-Business Suite Secure Enterprise Search region embedded in your application, you can:

1. Perform a search in the **Search region**.
2. View your **search results** in classic or tabular format.
3. Narrow down your search results using the **Narrow By region**.
4. Roll back to your previous search results using the **Expand By region**.
5. Enter filters in the **Filter region** to narrow down the search results on a searchable object only.

Using the Search Region

The Search region appears on top of the Oracle E-Business Suite Home page or your product home page.

To perform a search using the Search region

1. Select a searchable business category from the LOV.

To secure the application data to only those who have the access privileges, Oracle E-Business Suite Secure Enterprise Search only displays the searchable categories in the LOV that you have been authorized for access.

For example, if you are authorized to access 'Customer Relation' business category, you should find 'Customer Relation' from the LOV.

Note: If "All" business category is selected, then you will not find search results for objects that require security based on full security context privileges such as "Vacancies" object in iRecruitment. This is because "All" category search is designed to give results for objects with normal security (or not based on full security context) only. Instead, you will need to search on the group (such as "Vacancies") or narrow down to the object to get results if you have the access privileges.

If the 'Customer Relation' business category contains Contracts and Customer Account business objects, after executing a search on 'Customer Relation' category, you will find these two business objects displayed in a tree structure in the Narrow By region. See: Narrowing Down the Search Results Using Narrow By, page 20-23.

2. Enter a keyword search or enter partial values containing wildcard characters (regular expressions wildcard '*', not '%') on the selected category.

Note: Because date based columns are precrawled and indexed in the form of YYYY-MM-SS, if you want to base search on a date value, use the same format 'YYYY-MM-SS' as the query string. For example, use '2007-04-25' to search for records dated on 25-Apr-2007.

The search text entered here is also recorded in the Expand By region. See: Tracing Back the Results Using Expand By, page 20-24.

3. Select your Search Results Display Preference value from the following drop-down selection options.

- Standard (default): By default, search results will be displayed and rendered in Oracle E-Business Suite.

See: Viewing the Search Results in Oracle E-Business Suite, page 20-21

- Advanced: If 'Advanced' is selected, without logging out from Oracle E-Business Suite, the results will be displayed and rendered directly in the Oracle SES user interface if Single Sign On has been configured and the search results page in Oracle SES has been customized with your desired result

display.

See: Viewing the Search Results in Oracle SES, page 20-23

Click **Go** to execute your search and have the results displayed based on your selected display value.

Viewing the Search Results

Oracle E-Business Suite Secure Enterprise Search provides a flexible search result display mechanism allowing you to view the results either within the Oracle E-Business Suite or directly from the Oracle SES user interface based on your selected search results display value.

Viewing the Search Results in Oracle E-Business Suite

By default, 'Standard' Oracle E-Business Suite is the default value in the Search Results Display Preference field. This lets you view the search results in the Results region within the Oracle E-Business Suite.

The search result details include the last crawled time for your selected search object or group, current page number and results displayed, total number of estimated hits, time taken by search engine in performing the search, and the search results for each element.

Note: The total number of estimated hits is just a rough estimate. To have a more accurate estimate, Oracle SES allows administrators to set a parameter in global setting to report an exact count of search results. However, this option impacts query response time. See *Oracle Secure Enterprise Search Administrator's Guide* for details.

Each search result page contains a list of result elements that match the search criteria. The result can be displayed in the following different formats:

- **Classic Style:** All result elements are displayed in the classic style if the search is on a searchable category.
- **Tabular Style:** If a search is on a searchable object selected in the Narrow By region, the Result Style drop-down field appears letting you change the display from the default classic style to tabular format. The tabular style displays the search results in a table or spreadsheet like format, while the classic style displays the results in an Internet like format.

To easily navigate to each result element contained in the page and make the search result more user-friendly, the following information can be shown as part of the result display for each element:

- **Relevance Ranking:** A percentage value (or called 'Score') is displayed in a

parenthesis as part of the result element title link indicating the relevance of the result element is to your search criteria. For example, a result element with a score value of 78 percentage indicating a high relevance to the search criteria compared to an element with a score value of 40 percentage. A result element with the greater percentage value will be shown higher in the result list.

Click the title link of a result element lets you drill down the element details. For example, click on a job title 'Accounting Manager' (with a score or relevance percentage value) link to view the Accounting Manager job details.

- Score: It describes the source object name to which searchable object that each result element belongs.

This feature is particularly useful if a searchable group has multiple searchable objects or in case of "All" category search. This helps you easily identify the source of hit.

If the result is displayed in the classic style, you might find the following links available (if these attributes are identified and associated with the object you search on):

- Actionable Links: This lets you visit external URLs.
- Related Search: If a relation exists between the object searched on and other objects. This lets you perform a search on related object.

The related search content is secured by authorization. Only users with appropriate access privileges can execute search on the related objects. Otherwise, an unauthorized access message appears indicating that you do not have the access privileges.

For example, 'Customer Invoices' might be displayed as a related object while searching for 'Customer Relation'. Clicking on the 'Customer Invoices' object link will execute another search on the related object only if you have the access privileges. Otherwise, an unauthorized access message appears.

Note: For better query time performance, the search result will display related objects along with other search hits even though you do not have the access privileges to some of the objects.

Sorting Search Results

To have the search results displayed in your desired order, Oracle E-Business Suite Secure Enterprise Search provides the sorting option allowing you to sort the entire set of search results based on your selected value from the Sort By drop-down list.

The drop-down list will be available in both classic and tabular display format. The sorting criteria in the drop-down list will contain all the displayed attributes on which the sorting can be performed. Initially, the result will be sorted by Relevance in descending order. You can optionally change the sorting order by clicking on the Up Arrow icon next to the Sort By drop-down list to sort the result in ascending order.

Exporting Tabular-based Search Results

If a search result is displayed in the tabular format, the **Export** button is available in the search results region. Clicking the **Export** button lets you export the tabular-based search result content directly from the search results region to Microsoft Excel. You can save the exported data to a designated directory to use later.

Note: The program that is launched while saving exported data is dependent on your computer setting. To launch a .csv file in Microsoft Excel, the .csv file type needs to be mapped to open with Microsoft Excel. If it is not mapped, Microsoft Windows will ask you to choose a program with which to open the file.

This export feature is only available for the search results displayed in the tabular format, but not for the classic display.

The exported data can be served as an offline report which is particularly useful in several business functions.

For example, you can search inventory items by part number, specification, lot number, location or by stock levels or any other indicator as needed. After the search, you can export the results to Microsoft Excel and use the data for reconciliation with physical inventory. Additionally, you can search for any outstanding invoices for a particular customer by giving a date range (age bracket). Export the search results displayed in the tabular format and then perform offline activities such as reconciliation of accounts or following up with the customers on the invoices.

Viewing the Search Results in Oracle SES

Instead of rendering the search results region within Oracle E-Business Suite, you can view the search results directly in Oracle SES user interface if 'Advanced' is selected as the Search Results Display Preference value.

Oracle SES UI provides rich functionality in searching and grouping as well as the flexibility to customize the search results display for various business needs. This feature allows you to directly invoke and view the search results in Oracle SES UI without logging out from Oracle E-Business Suite if the Single Sign On has been configured and you have the search results page in Oracle SES UI customized with your desired result display.

If you want to refine your search, click the **Advanced Search** link to open the Advanced Search page where you can enter more search criteria or attributes for your search.

How to perform search in Oracle SES user interface and customize the search results, see *Oracle Secure Enterprise Search Administrator's Guide*.

Narrowing Down the Search Results Using Narrow By

Use the Narrow By region to narrow down the search results by selecting the following search components from a tree:

- **Searchable Group**

A business category is comprised of searchable groups and objects. A searchable group is preceded with the "check" icon when displayed in the tree. By selecting a searchable group link from the tree, all other searchable groups or objects contained in the selected group are displayed as a child node. If you select a searchable group from the child node, the selected group becomes a parent node in the tree, and the selected group name is added to the Expand By region allowing you to trace back the original search result. If a selected child node is a searchable object, then this object becomes a parent node and all the fields underneath become a child node.

Note that if a search is on a searchable group, the result is displayed in classic format only, and the Result Style field is not visible.

- **Searchable Object**

A searchable object node is preceded with the "object" icon in the tree. When refining your search on a searchable object, the selected object name becomes a parent node in the tree, and all the fields underneath become a child node. These field-based child nodes allow you to further narrow down the search results by applying filters. See: Using Filters on Searchable Objects, page 20-24.

Tracing Back the Results Using Expand By

Use the Expand By region to roll back to the original search result after you modify the search in the Narrow By region. The Expand By region displays all the search criteria you used for a search including the text entered in the Search region, group names, object names, or any filters. Clicking on any links in the Expand By region will bring the result back to the point from where you came. For example, clicking on the 'Oracle' filter link removes the filter from a search and expands the search result without being restricted by the filter, 'Oracle'. jkl

Using Filters on Searchable Objects

Use the Filter region to narrow down your search on a selected object. When you select a searchable object in the Narrow By region, the Filter region appears allowing you to apply filters on the selected object. By entering filter text and clicking **Search With Filter**, a search with filter criteria is executed.

Note: To search on an exact phrase through filters, use quotation marks (") around the phrase. For example, enter "contract imported" as the filter text to find the exact phrase from the search result. Without the quotation marks, you will find either 'contract' or 'imported' from the result.

Please note that filters do not consider words "to", "the" and "other". If any of those words are entered as filter text, you will find that search

results not containing above values are also displayed. This is because those words are considered as stopwords in Oracle SES. Hence, they are removed from queries.

See: Narrowing Down the Search Results Using Narrow By, page 20-23.

Performing Administrative and Setup Tasks

Since all searchable objects are precrawled and indexed in the Oracle SES index store before a search invokes, a search administrator or system administrator must perform administrative setup tasks. These tasks include enabling searches in the E-Business Suite, creating a search administrator who is responsible for setting up, and configuring Oracle E-Business Suite Secure Enterprise Search for Oracle SES integration.

This section includes the following topics:

- Creating a Search Administrator, page 20-25
- Setting Up Oracle E-Business Suite Secure Enterprise Search, page 20-26
- Setting Up Oracle E-Business Suite Secure Enterprise Search for Oracle SES Integration, page 20-31

Creating a Search Administrator

To have Oracle E-Business Suite Secure Enterprise Search work properly, a search administrator must be set up first in order to configure and maintain administrative tasks before users can perform searches on applications data.

Note: It is important to know that a search administrator is not only responsible for configuring and setting up essential tasks, but also responsible for managing crawling schedules and administering crawls in Oracle SES which are typically not performed by a system administrator. It is highly recommended that you create a new user (such as `sesadmin`) for that role, instead of using an existing system administrator user (`sysadmin`) assigned with necessary responsibilities. For more information, see *Oracle E-Business Suite Secure Enterprise Search Best Practices, Release 12*, My Oracle Support Knowledge Document 744820.1.

Use the following steps to create a search administrator:

1. Create a user (such as `sesadmin`) who will be the search administrator.
2. Assign the following responsibilities to the user `sesadmin`:

- Application Search Administrator responsibility (FND_SEARCH_ADMIN)
- FND Search Crawler responsibility (FND_SEARCH_CRAWLER)

Once a search administrator is created, the same user name and password information will be entered in the Application Search Administration page as part of the setup parameters for Oracle SES Integration, as well as entered in the Oracle SES administrative pages to validate and authenticate users for secured searches on Oracle E-Business Suite or add secure federated searches.

For detailed configuration and setup steps, see *Configuring Search Proxy Parameters for Oracle SES*, page 20-34 and *Performing Setup Steps in Oracle SES*, page 20-36.

Setting Up Oracle E-Business Suite Secure Enterprise Search

Setup Overview

Oracle E-Business Suite Secure Enterprise Search is comprised of database, middle-tier, and UI components. It also relies on external dependencies to have the function work properly. Before setting up Oracle E-Business Suite Secure Enterprise Search and performing administrative tasks, a search administrator or system administrator must first understand the product dependencies and the integration between Oracle SES.

Product Dependencies

Oracle E-Business Suite Secure Enterprise Search has dependencies on the following products in order to have its features work properly:

- **Oracle Secure Enterprise Search (SES)**

Oracle E-Business Suite Secure Enterprise Search relies on an external search engine to provide text search capability, and this search function is provided by Oracle SES.

Note: Oracle SES 11g client libraries are leveraged in this release.

- **AD Parallel**

Oracle E-Business Suite Secure Enterprise Search utilizes the AD Parallel Update package to facilitate the crawlable factory in expediting the performance of an initial crawl which usually involves a large set of data.

- **OA Framework**

Oracle E-Business Suite Secure Enterprise Search depends on OA Framework to enable the reusable search region displayed in the OA Home page.

Note: Since Oracle E-Business Suite Secure Enterprise Search can

embed reusable search regions either in the OA Home Page or as a plug-in, the OA Home Page is dependent on Oracle E-Business Suite Secure Enterprise Search as well.

- **Function Security**

Oracle E-Business Suite Secure Enterprise Search uses function security to guard the application content access through the menus and responsibilities assigned to an application user.

- **Data Security**

Oracle E-Business Suite Secure Enterprise Search uses data security to control what users can see on the application data through security grants. Only authorized users can view searchable objects.

Note: The security access to searchable objects are also implemented through the Role-Based Access Control Security for Oracle E-Business Suite Secure Enterprise Search, page 20-39.

- **Web Service Technology Stack**

Oracle E-Business Suite Secure Enterprise Search forms query and performs a search against Oracle SES through Web Service Technology Stack.

Setup Tasks

Since all searchable objects are precrawled and indexed in the Oracle SES index store before a search is invoked, a search administrator or system administrator needs to perform general setup tasks in Oracle E-Business Suite and search-related setup tasks both in Oracle E-Business Suite Secure Enterprise Search and Oracle SES administrative pages.

This section contains the following topics:

- Enabling Searches in Oracle E-Business Suite, page 20-28
 - Setting Language Preferences, page 20-28
 - Setting Profile Options, page 20-28
 - Assigning FND Search Crawler Responsibility to an FND User, page 20-30
 - Performing Personalization Setup Steps, page 20-30
- Setting Up Oracle E-Business Suite Secure Enterprise Search for Oracle SES Integration, page 20-31

- Configuring Search Parameters for Oracle SES, page 20-34
- Performing Setups in Oracle SES, page 20-36

Enabling Searches in Oracle E-Business Suite

The setup steps in Oracle E-Business Suite include the following tasks:

- Setting language preferences
- Setting profile options
- Assigning the FND Search Crawler responsibility to an FND user
- Performing personalization setup steps for displaying the Enterprise Search region

Setting Language Preferences

To have the search and result displayed in your preferred language, a search administrator must set a default language in the General Preferences page if it is not English.

For information on how to set language preferences, refer to Set Preferences section, Getting Started with Oracle E-Business Suite chapter, *Oracle E-Business Suite User's Guide* for details.

Setting Profile Options

Oracle E-Business Suite Secure Enterprise Search uses profile options to define necessary setup parameters so that searches can be enabled in the Oracle E-Business Suite. These profiles determine the following features:

- The availability of the Oracle E-Business Suite Secure Enterprise Search feature
- The valid URL for an external Oracle SES instance access
- The version of Oracle SES to be used for integrating with Oracle E-Business Suite Secure Enterprise Search
- The timeout value in seconds for an FND user logging into Oracle SES for query

Note: A valid FND user indicates that the user's application login information such as user name and password must be stored in the FND_USER table.

- The timeout value in seconds for an administrator logging into Oracle SES for administrative tasks
- The site-wide batch size used by AD Parallel for crawling

The following table lists the profile options used in Oracle E-Business Suite Secure

Enterprise Search:

Profile Option	Description	Required	Default Value
FND: Search Enabling Flag	Use this site level profile option to control whether Oracle SES integration is enabled for the site. Oracle E-Business Suite Secure Enterprise Search must have it set to Yes indicating this feature is enabled.	Yes	N
FND: Search Engine URL	Use this site level profile option to specify a valid URL with the format <code>http://<hostname>:<portnumber></code> for an external Oracle SES instance to which query will be made against. This profile value must be provided if the site is Oracle SES enabled.	Yes	N/A
FND: Search SES Version	Use this site level profile option to specify a valid version of Oracle SES for integrating with Oracle E-Business Suite Secure Enterprise Search. The profile value should have minimum two characters, and the first two characters should be digits (such as 11.1.2.2, 11g).	Yes	No default value for this option.

Profile Option	Description	Required	Default Value
FND: Search Session Timeout Value for Query	Use this site level profile option to control the timeout value in seconds for an FND user logging into Oracle SES. The session expires if this amount of time passes since the last activity by the user.	Yes	1200
FND: Search Session Timeout Value for Admin Tasks	Use this site level profile option to control the timeout value in seconds for an administrator logging into Oracle SES. The session expires if this amount of time passes since the last activity by the administrator.	Yes	1200
FND: Search Crawl Batch Size	This profile allows application administrators to set the site-wide batch size used by AD Parallel.	Yes	1000

Assigning the FND Search Crawler Responsibility to an FND User

The FND Search Crawler responsibility must be assigned to an FND user and its user name and password must also be provided in the search administrative page before synchronizing applications metadata with Oracle SES. If the user information changes, you must update it and synchronize the data again.

Performing Personalization Setup Steps for Displaying the Enterprise Search Region

To ensure that the Enterprise Search region appears on top of the Oracle E-Business Suite Home page, you need to perform the following personalization steps:

1. Log in to Oracle E-Business Suite with the system administrator's user name and password.
2. Select the Functional Administrator responsibility from the Navigator menu.

3. Select the Personalization tab and the Application Catalog subtab to open the Application Catalog page.
4. In the Search region, enter `/oracle/apps/fnd/search/webui` in the Document Path field as the search criteria and click **Go**.


All document names that match the search criteria should be displayed in the search result table.
5. Click the **Personalize Page** icon for the `/oracle/apps/fnd/search/webui/AppsSearchRG` document name listed in the result table to open the Choose Personalization Context page.
6. Enter 'Applications Home Page' (OAHOMEPAGE) in the Function field and click **Apply** to open the Personalize Region page for your document.
7. In the Personalization Structure region, select the **Personalize** icon for the Row Layout field to open the Personalize Row Layout page.
8. Set the rendered property to 'true' at the 'Function: Applications Home Page' level.
9. Navigate back to the Home page after applying the personalization change.
10. Review if the Enterprise Search region is getting rendered on the Oracle E-Business Suite Home page.
11. Perform a search by entering search criteria and click **Go** to verify if the search result is displayed in the search result page.

Setting Up Oracle E-Business Suite Secure Enterprise Search for Oracle SES Integration

To have seamless integration with Oracle SES, after enabling searches in Oracle E-Business Suite, the system administrator or search administrator must perform configuration tasks in Oracle E-Business Suite Secure Enterprise Search and Oracle SES.

The following diagram illustrates the high-level integration flow for Oracle E-Business Suite and Oracle SES integration:

Oracle SES and Oracle E-Business Suite Integration Workflow



After the installation of Oracle SES and Oracle E-Business Suite, search related metadata and business objects must be created and made it available in Oracle SES; metadata and relevant security rules are implemented and employed by Oracle E-Business Suite. To make the search available in Oracle E-Business Suite, necessary setup tasks must be performed in Oracle E-Business Suite, Oracle E-Business Suite Secure Enterprise Search, and Oracle SES.

For example, these tasks include setting language preferences and profile options for the applications to enable Oracle E-Business Suite Secure Enterprise Search, configuring search proxy parameters to facilitate remote access to Oracle SES, and performing administrative setup steps in Oracle SES for integration.

With appropriate setup and configuration between Oracle SES and Oracle E-Business Suite Secure Enterprise Search, searchable objects can be successfully deployed to Oracle SES instance, and initial and incremental crawls can be launched in Oracle SES. Users can perform queries from Oracle E-Business Suite against a precrawled index store in Oracle SES.

Important: For troubleshooting frequently encountered issues during installation and setups, see *Oracle E-Business Suite Secure Enterprise Search Troubleshooting Guidelines, Release 12*, My Oracle Support Knowledge Document 726239.1 for details.

This section covers the following setup tasks in Oracle E-Business Suite Secure Enterprise Search and Oracle SES:

1. Installing Oracle E-Business Suite Secure Enterprise Search, page 20-33

This section provides installation information so that Oracle E-Business Suite Secure Enterprise Search can integrate with Oracle SES.

2. Configuring Search Proxy Parameters for Oracle SES, page 20-34

This step includes setting administrative proxy and query proxy parameters, such as user name, password, and timeout value in seconds for an administrator and a valid FND user to access a remote Oracle SES instance.

3. Performing Setup Steps in Oracle SES, page 20-36

To ensure the seamless integration between Oracle E-Business Suite Secure Enterprise Search and Oracle SES, the system administrator or search administrator must perform additional setup steps in Oracle SES. These steps include setting up connections between Oracle SES and an identity management system, and adding federation entities to Oracle SES.

Installing Oracle E-Business Suite Secure Enterprise Search

Oracle E-Business Suite Secure Enterprise Search is released with Oracle E-Business Suite and Oracle Secure Enterprise Search (SES). To have it installed properly, perform the following installation steps:

1. Install or upgrade your instance to Oracle E-Business Suite Release 12.2.
2. Install Oracle SES 11.1.2.2 from Oracle Technology Network (OTN) page (<http://www.oracle.com/technetwork>), or upgrade to Oracle SES 11.1.2.2 from Oracle SES 10.1.8.4 or Oracle SES 11.1.2.0.

Refer to the *Oracle Secure Enterprise Search Installation and Upgrade Guide* 11g Release 1 (11.1.2.2) for installation details and upgrade information from Oracle SES 10.1.8.4 or Oracle SES 11.1.2.0 to Oracle SES 11.1.2.2.

Oracle SES can be integrated with Oracle E-Business Suite Release 12.2. The minimum supported version of Oracle SES in this release is SES 11.1.2.2. For more installation information, see *Installing Oracle E-Business Suite Secure Enterprise Search, Release 12*, My Oracle Support Knowledge Document 462377.1 for details.

Note: Oracle SES 11.1.2.2 uses two separate JVMs for running the crawler and search applications. The crawler is run using Sun JRE whereas the search application is run using JRockit JRE. Both JREs are available under \$ORACLE_HOME of Oracle SES installation. If Oracle E-Business Suite is on SSL enabled environment, when integrating with Oracle SES 11.1.2.2 instance, the Oracle E-Business

Suite SSL certificate has to be imported into both the Oracle SES keystores (JRE truststores) using keytool.

Remember the port and the associated password for user name 'eqsys' during the installation. This information will be used later in configuring Oracle E-Business Suite Secure Enterprise Search to enable Oracle SES integration.

Once you complete the installation for both Oracle E-Business Suite and Oracle SES, you must also perform administrative setup tasks both in Oracle SES and Oracle E-Business Suite Secure Enterprise Search to configure the system.

See:

- Configuring Search Proxy Parameters for Oracle SES, page 20-34
- Performing Setup Steps in Oracle SES, page 20-36

For more installation information, see *Installing Oracle E-Business Suite Secure Enterprise Search, Release 12*, My Oracle Support Knowledge Document 462377.1 for details.

To ensure that the Enterprise Search region appears on top of the Oracle E-Business Suite Home page, perform the personalization steps mentioned earlier. See: *Performing Personalization Setup Steps*, page 20-30.

Configuring Search Proxy Parameters in the Configuration Tab

Use the Configuration tab to set proxy parameters to enable Oracle SES instance access. This includes setting proxy parameters for an administrator and a valid FND user who has the FND Search Crawler responsibility.

Important: Changes in the proxy parameters including user name and password will require redeploying all searchable objects. If these objects have been crawled, then redeployment will not make data updates in Oracle SES. To resynchronize Oracle SES data, you must manually delete the data source of the same name in Oracle SES first, and then redeploy the objects.

Configuring Parameters to Access Oracle SES

The screenshot shows the Oracle Application Search Administrator interface. The top navigation bar includes 'ORACLE Application Search Administrator', 'Navigator', 'Favorites', 'Diagnostics', 'Home', 'Logout', 'Preferences', and 'Personalize Page'. The main content area is titled 'Application Search Administration' and has tabs for 'Schedules', 'Configuration', and 'Searchable Objects'. The 'Configuration' tab is active, showing a 'Personalize Stock Layout' section with three main parameter groups: 'Parameters', 'Admin Proxy', and 'Query Proxy'. Each group has an 'Update' checkbox and a text input field. The 'Parameters' group includes 'Personalize "SES End Point"' (with a note about URL format and an 'Update' checkbox), 'SES End Point URL' (input: http://rws60066rems.us.oracle.com:7777), 'SES Version' (input: 10.1.8.4), 'Admin Proxy' (with a note about valid admin user and password and an 'Update' checkbox), 'User Name' (input: eqsys), 'Password' (input: *****, masked), and 'Time Out' (input: 1200, unit: Secs). The 'Query Proxy' group includes 'Personalize "Query Proxy"' (with a note about FND username and password and an 'Update' checkbox), 'User Name' (input: sesadmin), 'Password' (input: *****, masked), and 'Time Out' (input: 1200, unit: Secs). On the right side, there is a 'Tasks' section with 'Personalize "Tasks"', 'Optimize Indexes' (with a note about index optimization and an 'Optimize Index' button), and 'Additional Tasks' (with a note about navigation and a 'SES Admin Login' link).

Use the following steps to configure search proxy parameters for an administrator and a valid FND user:

1. Log on to Oracle E-Business Suite with the Application Search Administrator responsibility and select the Application Search Administration link from the Navigator window.
2. From the Application Search Administration window, select the Configuration tab.
3. Specify the following information in the SES End Point region:
 - SES End Point URL: Enter an URL address with the format `http://<hostname>:<portnumber>`, such as `http://us.example.com:1234` in this field. This is an external Oracle SES instance to which query will be made against.

If you have the FND: Search Engine URL profile value defined, then you should see the URL value populated automatically.

To update this field, select the **Update** check box in the SES End Point region to enter new URL address. Click **Update** at the bottom of the page to save your change.

- SES Version: Enter an appropriate Oracle SES version that your system will be integrated with. It should have minimum two characters, and the first two

characters should be digits, such as 11.1.2.2, or 11g.

Values entered here will be stored in the 'FND: Search SES Version' profile option.

4. Specify the administrative proxy parameters including User Name, Password, and Time Out values in the Admin Proxy region.

The Time Out value field can be populated automatically if you set the 'FND: Search Session Timeout Value for Admin Tasks' profile value.

Note: To integrate Oracle E-Business Suite Secure Enterprise Search with Oracle SES, you need to set the Admin Proxy section as follows:

- User Name: `eqsys`
- Password: Use the same password for `eqsys` user name when you installed the Oracle SES.
- Time Out: 1200 secs

To update these fields, select the **Update** check box in the Admin Proxy region to make the changes. Click **Update** at the bottom of the page to save your change.

5. Specify the query proxy parameters including User Name, Password, and Time Out values for a valid FND user with the FND Search Crawler responsibility. This query user name and password is usually set to the system administrator `sysadmin` or search administrator `sesadmin` who has appropriate search responsibilities.

Like the Admin Proxy region, the Time Out value field can be populated automatically if you set the 'FND: Search Session Timeout Value for Query' profile value.

To update these fields, select the **Update** check box in the Query Proxy region to make the changes. Click **Update** at the bottom of the page to save your work.

Important: Once you change the query proxy parameters, the Oracle SES instance needs to be restarted to reflect the changes.

Performing Setup Steps in Oracle SES

Oracle E-Business Suite Secure Enterprise Search integrates with Oracle Secure Enterprise Search (SES) to provide powerful text search features. It allows Oracle SES to crawl application content and return results for query.

To ensure its seamless integration with Oracle SES, the search administrator needs to perform the following administrative tasks in Oracle SES after completing necessary

setup steps in Oracle E-Business Suite Secure Enterprise Search:

1. Log on to the Oracle SES administrative user interface using `http://<hostname>:<portnumber>/search/admin`. You can also access it from the **SES Admin Login** link in the Configuration tab of the Oracle E-Business Suite Secure Enterprise Search administrative page.

2. Select the Global Settings tab from the Secure Enterprise Search page to configure the following settings:

- Select **Identity Management Setup** from the System section to set up connections between Oracle SES and an identity management system to validate and authenticate users for secured searches.

Select

`oracle.search.plugin.security.identity.ebs.EBSIdentityPlug
inMgr` from the Available Identity Plug-in region and click **Activate**.

In the Activate Identity Plug-in page, enter the following parameter values to define the selected Identity Plug-in settings for all authentication and validation activity in Oracle SES:

- HTTP endpoint for authentication: Enter an end point URL for Oracle E-Business Release 12 authentication, such as `http://us.example.com:port/webservices/AppSearch/SecurityService`.
- User ID: Enter the search administrator's user name that you created earlier. See: Creating a Search Administrator, page 20-25.
- Password: Enter the search administrator's login password.

Click **Finish** to return to the Global Settings page.

- Select **Federation Trusted Entities** from the Search section to add federation entities. Oracle SES uses these entities to provide secure federated searches.

In the Federation Trusted Entities page, enter the following information:

- Entity Name: Enter the search administrator's user name that you created earlier. See: Creating a Search Administrator, page 20-25.
- Entity Password: Enter the search administrator's login password.
- Select the **Use Identity Plug-in for authentication** check box to authenticate through the active identity plug-in.

Click **Add** to return to the Global Settings page.

- Select the **Crawler Configuration** link from the Sources section to ensure the crawler logging setting is appropriate.

In the Crawler Logging region, ensure the crawler log file directory path including log file name defined in the Crawler Log File Directory field is less than the supported length of 100 characters for Oracle SES 11.1.2.2 integration.

Click **Apply**.

3. Restart both Oracle SES and Oracle E-Business Suite instances.

Important: You must make sure that the `search.properties` file in Oracle SES server is also properly configured. Use the following steps to set the time in milliseconds to wait for security filter refresh task to finish during query processing:

1. Locate the `search.properties` file in `$ORACLE_HOME/search/webapp/config` directory.
2. Set the security filter refresh task wait time value:
`sec_filter_refresh_wait_time=20000`

For more information on Oracle SES integration setup steps, see the *Oracle Secure Enterprise Search Administrator's Guide* for details.

Securing Searchable Objects

Security is the most critical feature that is designed to guard application content from unauthorized access. To ensure that the right person has access to appropriate data at the right time, searchable objects or metadata must be enforced by security rules before they can be made available for search within the Oracle E-Business Suite.

Oracle E-Business Suite Secure Enterprise Search provides a flexible mechanism to enforce and secure searchable objects without compromising on the data integrity and content sensitivity. To effectively manage search security both at the group and object levels, and reduce the search response time, the following security mechanisms are used in enabling Oracle E-Business Suite Secure Enterprise Search:

- **Role-Based Access Control (RBAC) Security**, page 20-39

With RBAC security model, search security can be enforced at the group level through security grants. This section discusses the Role-Based Access Control (RBAC) security model and steps in creating security grants to ensure the application content sensitive data is only accessed by authorized people.

- **Search Security Plug-ins**, page 20-43

Security plug-ins provide another layer of security control at the object level. A security plug-in can be added to a searchable object at the design time; it can also be used at crawl time to fetch ACLs and generate Security Keys offline or at query time

if needed to protect unauthorized access to application data. Sample security plug-ins are also described in this section.


- **User Authorization Cache**, page 20-63

User Authorization Cache (UAC) feature allows previously cached Security Keys for a specific user and a searchable object or data source to be used at query time. This greatly reduces the search response time and provides a quick result with security enforced. When a user performs a search, the cached information is first examined to look up the query user or an object. If the match is found, the associated keys will be compared with the pre-built ACLs to authorize or revoke the document access privileges. In case there is no match found, then the Security Keys are fetched synchronously.

Role-Based Access Control (RBAC) Security

Oracle E-Business Suite Secure Enterprise Search uses Role-Based Access Control (RBAC) security to secure searchable objects through roles, and user access to applications data is determined by the roles granted to the user. This approach builds upon Data Security and Function Security, but it goes beyond both of them.

Role-Based Access Control Security



Function security is the base layer of access control in Oracle E-Business Suite. It restricts user access to individual menus and menu options within the system, but it does not restrict access to the data contained within those menus. Data security provides access control on the application data, and the actions a user can perform on the data. With data security, users can be restricted by security rules to access or view only certain types of data on the screen once they have selected a menu while an administrator can have more data access to the same page.

With RBAC, access control is defined through roles, and a role can be configured to consolidate the responsibilities, permissions, function and data security policies that

users require to perform a specific task. This solution simplifies mass updates of user permissions because changes can be done through roles which inherit the new sets of permissions automatically. Based on the job functions, each role can be assigned a specific permission or permission set if needed. For example, a sales organization may include 'Sales Representative', 'Sales Manager', and 'Sales Support' roles. The 'Sales Manager' role would include a permission set allowing the manager role to perform a job function for both the representative and support roles.

By leveraging the concept of permission sets, Oracle E-Business Suite Secure Enterprise Search allows related searchable objects to be grouped and sequenced to construct searchable groups; these searchable groups are associated with a function role and then they are assigned to users through security grants. When a user logs on to the E-Business Suite and performs a search, Oracle E-Business Suite Secure Enterprise Search filters the secured searchable objects based on the grant and displays the list to the user who has authorized privileges.

For example, Oracle E-Business Suite Secure Enterprise Search uses search function security to provide a permission on searchable objects and then grant to different roles. When a searchable object 'purchase order' is defined, in order for a user to search on this purchase order object, she or he must have been assigned to a role that holds a grant of purchase order access permission. Once the user logs on the applications, she or he should be able to see the purchase order displayed from the searchable object list for search selection.

For more information on Data Security, Function Security, and RBAC security models, see *Oracle E-Business Suite Security Guide* for details.

Creating Security Grants

To secure application data access to a user with right privileges, the system administrator or search administrator needs to administer the security grant which includes:

1. Create Objects, page 20-40
2. Create Permissions, page 20-41
3. Create Permission Sets, page 20-42
4. Grant Permission Sets, page 20-42

Prerequisites: Please note that the system or search administrator needs to have Functional Developer role to create objects, permissions, and permission sets, as well as the Functional Administrator role to create grants.

To create security grants:

Creating Objects

1. Log on to Oracle E-Business Suite with the Functional Developer responsibility.

2. From the Security tab, click the Objects subtab, and select **Create Object**.
3. Enter the following fields to create an object:
 - Name: Enter a display name.
 - Code: Enter a name, such as WF_SEARCHABLE_NTF.
 - Application Name: Select an application name.
 - Database Object: This must be FND_OBJECTS.
 - Description: Enter a proper description for the object you want to create.
 - Column: Select the first column name as CRAWL_CRAWLABLE. The Type field populated automatically with VARCHAR2.
4. Click **Apply**.
5. Select the Object Instance Set tab and click **Create Instance Set**.
6. Enter the following information to create an instance set:
 - Name: Enter a display name for the instance set.
 - Code: Enter a code.
 - Description: Enter a proper description for the instance set.
 - Predicate: Enter a predicate.

Creating Permissions

1. Log in with the Functional Developer responsibility. From the Security tab, select the Permissions subtab and click **Create Permission**.
2. Enter the following fields to create a permission:
 - Name: Enter a name for the permission, such as WF: Searchable Notifications.

Note: The permission name entered here will be displayed as a searchable object name in the Narrow By region which allows you to refine your search from the Search Results page. For more information on how to use Oracle E-Business Suite Secure Enterprise Search, see *Oracle E-Business Suite User's Guide*.
 - Code: Enter a standard code, such as WF_SEARCHABLE_NTF.

- Description: Enter a proper description for the permission.
- Object Name: Select the object you created in the previous steps.
- Add to Permission Set: Select a permission set for this field if you have a permission set created.

3. Click **Apply**.

Creating Permission Sets

1. Log in with the Functional Developer responsibility. From the Security tab, select the Permission Sets subtab and click **Create Permission Set**.

2. Enter the following fields to create a permission set:

- Name: Enter a name for the permission set, such as ATG Searchables.

Note: If you are authorized to have the security access to the permission set name you entered here, when you perform a search on the Oracle E-Business Suite Home Page or a product home page, you should find this permission set name displayed from the business category drop-down list for your selection. For more information on how to use Oracle E-Business Suite Secure Enterprise Search, see *Oracle E-Business Suite User's Guide*.

- Code: Enter a standard code, such as SESG_WF_NTF.

Important: Your permission set must be prefixed with SESG.

- Description: Enter a proper description for the permission set.

3. Click **Add Another Row** and enter the following information:

- Permission: Select a permission you created earlier, such as WF: Searchable Notifications.
- Add more permissions as appropriate.

4. Click **Apply**.

Granting Permission Sets

This process requires the 'Functional Administrator' role to create grants.

1. Log in with the Functional Administrator responsibility. From the Security tab,

select the Grants subtab and click **Create Grant**.

2. Enter the following fields in the Create Grant: Define Grant page:
 - Name: Enter a name for the grant, such as ATG Searchables Grant.
 - Description: Enter a proper description for the grant.
 - Enter proper information in the Effective From and Effective To fields.
3. Enter the following information in the Security Context region:
 - Grantee Type: Select a proper grantee type, such as Group of Users.
 - Grantee: Enter System Administrator.
4. Click **Next**.
5. In the Set region, select a permission set to grant, such as ATG Searchables and click **Next**.
6. Review the grant details and click **Apply**.

Search Security Plug-ins

In addition to securing your search at the group level through security grants, Oracle E-Business Suite Secure Enterprise Search uses security plug-in to strengthen security further down to the object level. Since searchable objects are the key elements in the crawling mechanism, this type of security mechanism can be easily implemented and enforced at crawl time and even can be dynamically executed during user query. Its flexible, object-based security plug-in feature can effectively guard and protect application sensitive data such as HRMS employee data, General Ledger data in a legal entity from unauthorized access or transactions across organizations if in a multiple-organization environment.

Security plug-in is a Java class that implements security methods to support custom or user-defined security rules at the object level and in turn to secure your search.

At design time, a security plug-in can be added to a searchable object during the object creation through the metadata-based Search Modeler user interface.

At crawl time, while creating indexable documents, two search methods (`getAcl()` and `getSecureAttrAcl()`) of the plug-in associated with the object definition are invoked to generate the access control list (ACL) for each document.

Note: An ACL is a list of permissions attached to an object specifying who or what is allowed to access the object and what operations are allowed to be performed.

Oracle SES authorization plug-in works on the basis of the ACL-based security model and Security Keys for a document to authorize users or revoke the access to a search result. Through the authorization plug-in implementation of Oracle E-Business Suite connector in Oracle SES, all searches within Oracle E-Business Suite can be authorized and leveraged from the SES search engine.

At query time, when a user performs a search, different sets of search methods (`getSecurityKeys()` and `getSecureAttrKeys()`) of the plug-in are executed to generate the **Security Keys** for the user in order to match the pre-built ACLs. Any matched indexed documents will then be retrieved for the user. Unmatched or unauthorized documents get dynamically filtered out in the process.

Security Keys and User Authorization Cache (UAC)

To reduce the search response time of fetching Security Keys simultaneously during user query, User Authorization Cache (UAC) framework in Oracle SES is leveraged to allow Security Keys to be generated as an offline process if a User Crawler initiates at the crawl time.

This user crawling process generates a list of Oracle E-Business Suite users for whom the Security Keys needs to be cached in Oracle SES. Security Keys are then generated against the user list by executing (`getSecurityKeys()` and `getSecureAttrKeys()`) methods of the plug-in. These generated keys for a given user and a specific searchable object or data source are cached as User Authorization Cache and will be looked up during user query to see if any match for a given source and user and whether the cache is usable.

For more information about User Authorization Cache feature and how it works, see User Authorization Cache, page 20-63.

How to add a search security plug-in to an object, see Creating Searchable Objects, *Oracle E-Business Suite Search Modeler User's Guide* available from My Oracle Support Knowledge Document 781366.1, Search Modeler 1.1 for Oracle E-Business Suite Readme.

This section includes the following topics:

- How Security Plug-in Works, page 20-45
- Understanding Security Logic and General Plug-in Mechanism, page 20-47
- ACL-based Security, page 20-49
- Query Rewrite Security, page 20-51
- Supporting Security Models with Search Plug-ins, page 20-55
- Other Considerations, page 20-61

How Security Plug-in Works

To effectively guard application content from unauthorized access and support various business requirements within Oracle E-Business Suite, security plug-in mechanism is implemented to ensure the search security and context sensitive information only accessible to appropriate users.

This section highlights and further explains the roles of security plug-in from crawl and query different perspectives. It includes the following topics:

- Crawl Time to Generate ACLs, page 20-45
- Query Time to Generate Security Keys, page 20-45

Crawl Time to Generate ACLs

Security plug-ins are used to fetch ACLs at crawl time.

When Oracle E-Business Suite Crawlable End Point receives crawl requests from Oracle SES crawler threads, the Crawlable Factory is initialized to fetch the indexable content from Oracle E-Business Suite database and create crawlable documents. While creating indexable documents, the security plug-in associated with the searchable object definition will be used through the invocation of the `getAcl()` and `getSecureAttrAcl()` methods to generate ACLs for the documents.

At this time, these indexable documents in the form of RSS feed is ready to be consumed. Oracle E-Business Suite crawler threads pick up the documents; the Crawlable End Point sends them back to the SES indexing engine as crawling responses. The SES indexing engine will then analyze the documents and transform them into indexed documents with readable format.

Query Time to Generate Security Keys

At query time, security plug-ins are used to generate Security Keys for the query user.

Query through Oracle E-Business Suite

When a user performs a search through the Oracle E-Business Suite user interface, a search session is created and the applications context is also initialized for the user. The applications context may be incomplete at this stage depending on if the user has selected a responsibility or not after logging on to the Oracle E-Business Suite.

Note: Applications context information is required for application users to perform certain business transactions or to be used in security plug-in to generate the ACLs and Security Keys for the users. It contains username, responsibility, responsibility application, and security group information.

When the query is submitted to the SES client APIs, the APIs in turn invoke the Web service calls in the Oracle SES server. To ensure the user is authorized for a search, Oracle SES first looks up the previously cached Security Keys for the object and logged-in user in User Authentication Cache (UAC). If a match is found and the cache is

usable, the associated keys will be used to compare the pre-built ACLs. Any matched indexed documents will be retrieved for the user. If no match is found, Identity Manager in Oracle SES requests Security Keys for the user through Security Service End Point. A proxy session is initialized to verify the credentials of the proxy username and password required by Oracle SES for the user. This proxy session is trusted or updated on behalf of the actual search user for whom the security keys have been requested.

Please note that the proxy applications context may be incomplete since the search can be performed either with or without the responsibility information. To generate Security Keys for the user in order to perform certain business transactions or activities that require full applications context information, you must extend the `oracle.apps.fnd.search.impl.ContextSecurable` plug-in class to create the complete context information. For more information about the plug-in mechanism, see *Understanding Security Logic and General Plug-in Mechanism*, page 20-47.

Security plug-in is also invoked by the Security Service End Point to generate the Security Keys through the execution of the `getSecurityKeys()` and `getSecureAttrKeys()` methods for the proxy content.

Once the Security Keys are generated, the Security Service End Point sends the keys back in response to the earlier request from Identity Manager. This request-response happens over HTTP protocol.

To ensure that it does not wait indefinitely for the response to complete, Oracle SES can set a timeout message on the request. The timeout value is configurable.

Query through Oracle SES

When an Oracle E-Business Suite user performs a search through the Oracle SES user interface, the security checks can be performed in the following two stages:

1. **Login Security Authentication:** This stage validates the user's login credentials through Oracle SES without security or authorization plug-ins.
2. **Search Security Authorization:** This stage begins when a user submits a search query after successful login. Search plug-in is used in the same way as described in querying through Oracle E-Business Suite that is to generate Security Keys for the query user.

Note: The major difference between searching from within Oracle E-Business Suite and from Oracle SES Search UI is that while searching from the Oracle SES Search UI, the proxy applications context is always incomplete since the responsibility information may not be there.

If certain business transactions or activities that require full applications context information, you must extend the `oracle.apps.fnd.search.impl.ContextSecurable` plug-in class to create the complete context information. For more

information about the plug-in mechanism, see Understanding Security Logic and General Plug-in Mechanism, page 20-47.

The user query is submitted to the SES client APIs which in turn invoke the Web service calls in the Oracle SES server. To ensure the user is authorized for a search, Oracle SES first looks up the previously cached Security Keys for the object and logged-in user in User Authentication Cache (UAC). If a match is found and the cache is usable, the associated keys will be used to compare the pre-built ACLs. Any matched indexed documents will be retrieved for the user. If no match is found, authorization plug-in contacts the Identity Manager to fetch the Security Keys. Identity Manager sends a request message containing the proxy username and password to Oracle E-Business Suite Security Service End Point. Security End Point establishes the proxy session and Applications Context. After the user credential (proxy username and password) is verified, the proxy session is trusted or updated on behalf of the actual search user for whom the Security Keys have been requested.

Security plug-in is invoked by the Security Service End Point to generate the Security Keys through the execution of the `getSecurityKeys()` and `getSecureAttrKeys()` methods for the proxy content.

Once the Security Keys are generated, the Security Service End Point sends the keys back in response to the earlier request from Identity Manager. Authorization plug-in receives the Security Keys for the search user.

Use Security Keys to Match the Pre-built ACLs

Oracle SES search service or APIs retrieve indexed documents from index store, matching the search keywords and filters. Indexed documents with pre-built ACLs are filtered by the Security Keys retrieved for the search user. Filtered search results are returned back to the query user. Unmatched or unauthorized documents get dynamically filtered out in the process.

Understanding Security Logic and General Plug-in Mechanism

Implementing Security Logic

Oracle E-Business Suite Secure Enterprise Search provides security through an interface. Once implemented, various methods of this interface can be called at different stages to enforce the security on the content of a searchable object. Each searchable object can have a plug-in Java class, nominated at design time through the Search Modeler user interface. If this class implements the `Securable` interface, the rules implemented by this class are enforced on the searchable object.

This `Securable` interface security plug-in Java class includes the following security methods:

- **`isAclEnabled`**

Crawlers use this method to determine if a document is ACL guarded or not. The `getAcl()` and `getSecurityKeys()` methods are called only when this method returns a value of *true*.

- **getAcl**

This method returns a list of tokens as access control during crawl time to extract each doc with ACL. This method is called for every document before they are sent to Oracle SES for indexing. The string value must be uppercase, alphanumeric, and not longer than 30 characters. This list is used as a predicate when the query is sent to Oracle SES. For example, if you return a string value of "XYZ123ABC", then only the person who holds the key "XYZ123ABC" can access this document, which is then returned by the `getSecurityKeys()` method.

- **getSecurityKeys**

This method returns an array of keys that the user holds at the time the query is made. You can get the session user from search context that is passed into this method. This method is called only once per session.

- **getSecureAttrAcl**

In some cases, ACLs may be related to an attribute of the searchable object. If this is true, then you must specify which attribute of the searchable object is a secure attribute. When it is specified and the `Securable` interface is implemented by its plug-in, this method is called for each secure attribute of each document. This method returns a list of ACL for secure attributes at crawl time. It uses the same rules as for `getAcl()`, except it is associated with a particular attribute. This method is paired with `getSecureAttrKeys()` as `getSecurityKeys()` is with `getAcl()`.

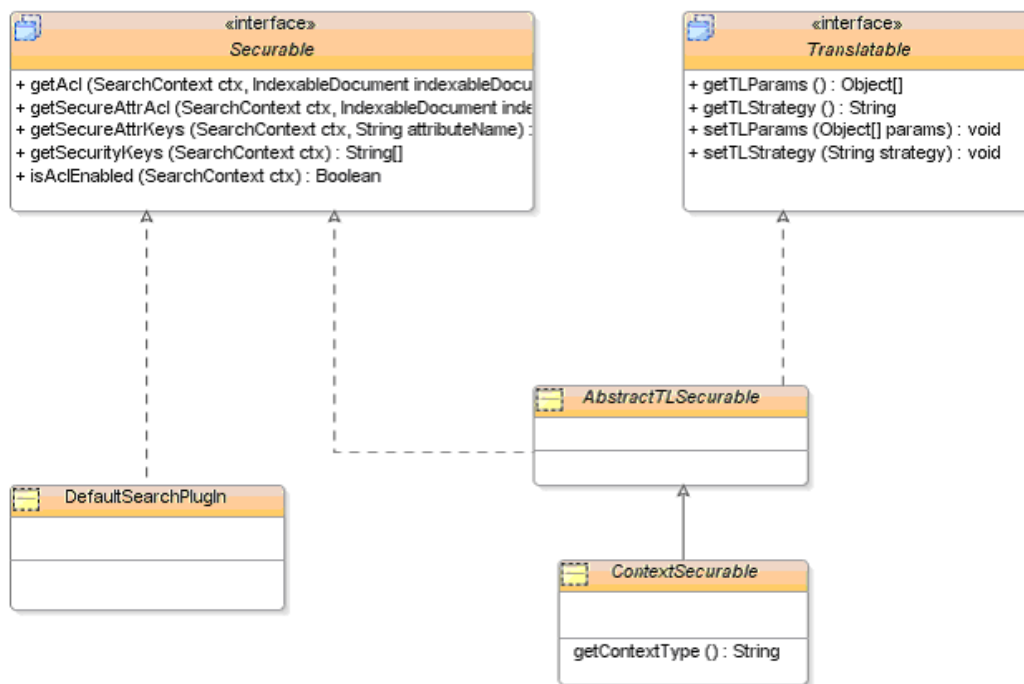
- **getSecureAttrKeys**

This method is called for each secure attribute per session. It returns a match string list based on the user's access.

General Security Plug-in Mechanism

Certain security plug-ins not only provide `Securable` feature, but also provide `translation` feature to searchable attribute. The relationship between `Securable` interface, `translation` interface, and other associated Java classes is illustrated in the following plug-in class hierarchy diagram:

General Plug-in Class Hierarchy Diagram



- `oracle.apps.fnd.search.AbstractTLSecurable`: It is an abstract Java class which supports translation of Searchable and Displayed attribute names. Any business requirement which requires the displayed attribute names to be translated should extend this class.
- `oracle.apps.fnd.search.impl.ContextSecurable`: It is an abstract Java class which can have the full applications context information in order to generate the ACLs and Security Keys. Any plug-in business logic which would require full applications context should extend this class. It is imperative that any subclass of `ContextSecurable` plug-in class has to implement translations as well.
- `oracle.apps.fnd.search.impl.DefaultSearchPlugIn`: It is a concrete Java class which secures the search results based on whether the search user has access to the target UI function or not. This Java class (and any of its subclass) does not support translation of attribute names and full applications context.

ACL-based Security

An access control list (ACL) is a list of permissions attached to an object. The list specifies who or what is allowed to access the object and what operations are allowed to be performed on the object.

In the ACL-based security, each entry in the list specifies a subject and an operation. For example, the entry (Alice, delete) on the ACL for file WXY gives Alice permission to

delete file WXY. Each crawled entry is associated with a locker, which is created in the form of an ACL. Both the locker and the content get indexed in Oracle SES.

When a user performs a search on an object in Oracle E-Business Suite, the system first checks the list for an applicable entry in order to decide whether to proceed with the query.

When crawled, the `getAcl()` method for each document is called by the crawler and it returns an ACL, which is indexed by Oracle SES along with the document.

Please note that the ACL-based security approach can also be used at query time along with `getSecureAttrKeys()` for additional security based on secure attributes. For more information, see Supporting Security Models with Search Plug-ins, page 20-55.

A Security Example with ACL-based Security

In this example, the ACL is a list of responsibilities that have access to functions in `FND_FORM_FUNCTIONS`. Resolving this relationship requires some complex logic and a number of tables including `FND_MENUS` and `FND_MENU_ENTRIES`.

Function ID	Name	Content	ACL
1	Edit	Oracle Workflow	10 20
2	Update	Oracle Test	10
3	Create	Oracle Financial	30
4	Delete	Oracle Personnel	40

At query time, a search user needs to acquire key(s) for a secure searchable object. In this example, it is a list of responsibilities assigned to the user. The `getSecurityKeys()` method returns this list when it is called. The query is rewritten with the key(s) and posted to Oracle SES. (This is the equivalent of adding security predicates to a query in SQL before hitting the database table).

The user SYSADMIN logs in and issues a query search on the content *Oracle*. Before hitting Oracle SES, the `getSecureAttrKeys()` method is called with the proper user context and returns a list of responsibilities assigned to SYSADMIN, which is "10", "20", and "30". The query is rewritten as:

```
(ACL_KEY: 10 OR 20 or 30) AND content: oracle
```

Using the above example, this query returns the following:

Function ID	Name	Content	ACL
1	Edit	Oracle Workflow	10 20

Function ID	Name	Content	ACL
2	Update	Oracle Test	10
3	Create	Oracle Financials	30

This approach incurs some cost at crawl time because it calls the `getAcl()` method for each row in `FND_FORM_FUNCTIONS`. However, this is acceptable if the underlying table is relatively small.

Tip: In this example, there are approximately 40,000 records in `FND_FORM_FUNCTIONS` and it takes about ten minutes to crawl the entire table.

Query Rewrite Security

By leveraging the searchable attribute feature, Query Rewrite provides another layer of security mechanism during user query to secure application content.

To use this Query Rewrite security, one or more searchable object attributes have to be marked as "Secured" at design time during object creation so that the `getSecurityKeys()` method of the search plug-in can be invoked for each "Secured" attribute at the time of user query.

At crawl time, since the secure attribute concept is used in this mechanism, no access control list will be generated during crawl for the documents of a searchable object.

At query time, when a user performs a search, the user acquires key(s) for a secure object to authenticate the operation based on the applications context. Oracle SES fetches the run-time keys for the user by invoking `getSecurityKeys()` method for each secure attribute, and returns a list of Security Keys for the user to access the secure object. The query is rewritten with keys and posted to Oracle SES. This query rewrite concept is similar to add security predicates to a query in SQL before hitting the database table. As a result, only proper data for the authorized user will be returned as the search result, but the entire query rewritten process is transparent to the user.

For example, a Purchase Order should be visible only to the buyer who places the order. With this design principle, `BUYER_ID` acting as an identifier should be marked as a secure attribute and it should be associated with the document of Purchase Order searchable object during the design time to secure the order content. When a user searches for "laptop docking station" related purchase orders, the query will be rewritten with the predicate, `BUYER_ID = "<buyer id of user>"`. This approach reinforces the secured object Purchase Order and only allows the person who places the order to have the relevant UI and order access privileges.

Tip: Query rewrite method should be used when the number of keys is limited in number. If the number of keys is higher, use ACL method instead.

With the query oriented approach, you nominate one or more attributes as secure attributes for each searchable object. For example, the BUYER_ID can be a secure identifier for a purchase order.

At query time, (before searching Oracle SES), a `getSecureAttrKeys()` method is called for each secure attribute such as BUYER_ID. The `getSecureAttrKeys()` method implementation should map the number of IDs that the current user has access to.

Note: The query only returns results that the user has access to.

When using this approach, there is no need to form Access Control Lists (ACLs) and security is basically enforced at query time, which means there is less risk of security data being out of date. However, this approach does not work if a user has too many keys. To rewrite a query that has thousands of keys is unrealistic. In order for this approach to perform you must provide a way to limit the number of keys returned by `getSecureAttrKeys()`.

Tip: When there is no clear way to limit keys, or it is too expensive to resolve keys at query time, you should use the crawl oriented approach.

A Security Example with ACL with Query Rewrite

The ACL-based security approach can also be used at query time along with `getSecureAttrKeys()` for additional security based on secure attributes.

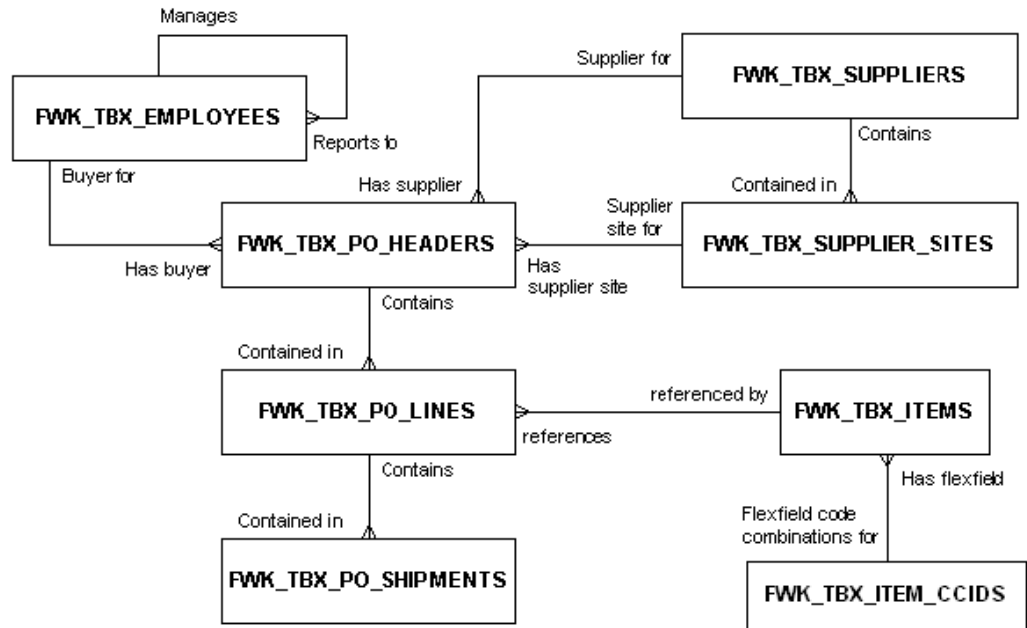
Security Rule Based on Employee Hierarchy

Take purchase orders as an example to explain the security mechanism that combines both ACL-based security and Query Rewrite security.

Purchase orders are usually owned by the user (employee) who initiates the purchase, and the employee hierarchy is usually used as the rule for its visibility to others. For example, an employee can access only his or her own purchase orders, whereas a manager can access the purchase orders he or she owns and those that are initiated by his or her reportees.

The following diagram illustrates a typical entity relationship for a purchase order.

Purchase Order Entity Relationships



The following list of attributes have been selected for displaying:

Entity	Attribute
FWK_TBX_PO_HEADERS	HEADER_ID (PK) STATUS_CODE DESCRIPTION BUYER_ID (FK)
FWK_PO_LINES	LINE_ID (PK) HEADER_ID (FK) ITEM_ID (FK) ITEM_DESCRIPTION
FWK_TBX_ITEMS	ITEM_ID (PK) ITEM_DESCRIPTION

Entity	Attribute
FWK_TBX_EMPLOYEE	EMPLOYEE_ID (PK)
	FIRST_NAME
	LAST_NAME
	EMAIL_ADDRESS
	MANAGER_ID

To use Query Rewrite security, you should mark the `BUYER_ID` attribute as a "Secured" attribute. The content is then indexed in Oracle SES during crawl time without an ACL.

The following is a sample listing of purchase orders that are initiated by different employees. Employees *adillon* and *bcarey* report to *ekane*, *ekane* reports to *rlavery*, and *rlavery* reports to *khart*.

Header ID	Status Code	Description	Content	BUYER_ID
1	Open	Dell Computer	Oracle Workflow	12 (adillon)
2	Closed	Apple, Inc	Oracle SES	13 (khart)
3	Open	Oracle	Oracle Test	14 (ekane)
4	Approved	Microsoft	Oracle Financial	15 (bcarey)
5	Approved	Oracle	Oracle	13 (khart)
6	Closed	Dell Computer	Oracle Framework	10 (rlavery)
7	Open	Oracle	Oracle Personnel	14 (ekane)

At query time, an authenticated user acquires a key, or keys, via the `getSecureAttrKeys()` method. This method passes applications context information along with the secure attribute such as `BUYER_ID`. It returns a list of keys to access purchase orders. The query is rewritten with the keys and posted to Oracle SES.

Using the above list of purchase orders, if the user *ekane* performs a keyword search on 'Oracle', she would have the keys 12, 14, and 15. The results would be:

Header ID	Status Code	Description	Content	BUYER_ID
1	Open	Dell Computer	Oracle Workflow	12 (adillon)
3	Open	Oracle	Oracle Test	14 (ekane)
4	Approved	Microsoft	Oracle Financial	15 (bcarey)
7	Open	Oracle	Oracle Personnel	14 (ekane)

If *adillon* performs the same search only one row is returned. However, this approach becomes more complicated when the head of a real department performs a search because they own the entire hierarchy and may have thousands of keys. For this case you can add logic in the `getSecureAttrKeys()` method so that only a specific number of keys or levels of hierarchy is returned.

Supporting Security Models with Search Plug-ins

To secure sensitive application data from unauthorized access and support complex security needs within Oracle E-Business Suite, Oracle E-Business Suite Secure Enterprise Search provides seeded security search plug-ins. These plug-ins are pre-built public Java classes which support well-known application security models. With the flexible plug-in security mechanism, users can search and navigate to appropriate transaction pages with security enforced to obtain needed information.

Oracle E-Business Suite Secure Enterprise Search supports the following security models with seeded search plug-ins:

- Business Group Based Search Security
- Legal Entity Based Search Security
- Organization Based Search Security
- Employee Hierarchy Based Search Security

Common Security Features of Seeded Search Plug-ins

Although these seeded search plug-ins are provided for various business reasons to secure sensitive application data, they all have the following common security features:

- **Designed Based on ACL and Query Rewrite Security Models**

All these four seeded search plug-ins are designed based on both the ACL and Query Rewrite security models. As a result, the `isAclEnabled()` method must return "true" in order for the `getAcl()` and `getSecurityKeys()` methods to be called later on during the crawl and query.

- ***Crawl Time ACL Implementation***

In order to fetch ACLs while creating indexable documents at crawl time, the `getAcl()` and `getSecureAttrAcl()` methods of the plug-in have to be implemented.

Additionally, the `getAcl()` method should extend the `DefaultSearchPlugIn` concrete Java class (`oracle.apps.fnd.search.impl.DefaultSearchPlugIn`), so that the access to specific target UI function can be securely enforced. For more information about `DefaultSearchPlugIn`, see *Understanding Security Logic and General Plug-in Mechanism*, page 20-47.

- ***Query Time Security Key Implementation***

At query time, after a query user's credential (proxy username and password) is verified and updated in response to a request of the user security validation, the `getSecurityKeys()` method should be simply invoked with super class implementation, and the `getSecureAttrKeys()` method returns the list of matched business groups that the query user has access to.

- **A Secure Attribute Needed for Query Rewrite**

A *secure attribute* is needed to implement the Query Rewrite security mechanism.

A searchable object contains a number of database tables or views; each table or view contains a number of columns that are bound to business data. While defining a searchable object through Search Modeler, you can select needed columns for each selected table name (entity) for your object. These table columns are called *attributes* that can be indexed for search. If an attribute contains certain feature that can be used to secure documents during search, then this attribute can be considered as a secure attribute. For example, the `BUSINESS_GROUP_ID` column that acts as a Business Group identifier can be a *secure attribute* for HR tables.

At query time, in order for query engine to generate Security Keys, (before searching Oracle SES) the `getSecureAttrKeys()` method is called for each *secure attribute* such as `BUSINESS_GROUP_ID`. The `getSecureAttrKeys()` method then returns a number of Business Group IDs that the query user has access to. As a result, the query to Oracle SES is rewritten.

"Secured" and "Stored" Attribute

To ensure the security check of each search plug-in can be enforced, all seeded plug-ins must contain an unique, secure table column (attribute) as a secure identifier. To differentiate the "secure" feature from other attributes, the attribute property "Secured" should be selected while defining a searchable object in Search Modeler. An attribute marked with "Secured" property indicates that this attribute can be used for securing the document using search plug-in. Additionally, select "Stored" property for the attribute. This indicates that this attribute can be stored in Oracle SES.

Note: If an attribute is marked not "stored", it cannot be displayed in the search result summary.

Other available attribute properties can be like "displayed", "title", "Is attachement", etc. For more attribute property information, see *Creating Searchable Objects, Oracle E-Business Suite Search Modeler User's Guide* available from My Oracle Support Knowledge Document 781366.1, Search Modeler 1.1 for Oracle E-Business Suite Readme.

For example, to ensure business group based security, Business Group Search Plug-in will require that the search object definition should have BUSINESS_GROUP_ID attribute, which should be marked at least "Stored" and "Secured".

For more information about ACL and Query Rewrite security models, see *ACL-based Security*, page 20-49 and *Query Rewrite Security*, page 20-51.

In addition to the common security features, each search plug-in contains various security requirements and secure attribute information. They are further explained in the following sections:

- Business Group Based Search Security, page 20-57
- Legal Entity Based Search Security, page 20-59
- Organization Based Search Security, page 20-59
- Employee Hierarchy Based Search Security, page 20-60

Business Group Based Search Security

A good example of this type of security model is Oracle E-Business Suite Core HRMS system.

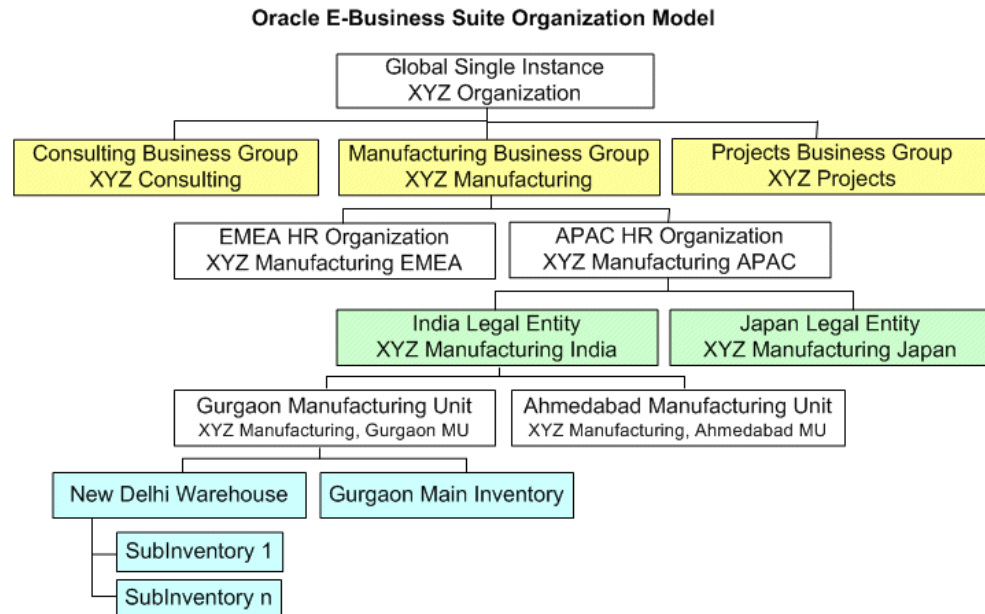
- **Security Requirements:** A Business Group is an organization that is set up and configured in Oracle Human Resources. It is the highest level classification in the organization hierarchy of the Oracle E-Business Suite, and it often relates to country-specific legislation. It may correspond to your entire enterprise or to a major grouping such as a subsidiary or operating division. If your payroll tax and employment authorities permit it, you can group employees of different registered companies together for reporting.

In this security model, data within a business group is only visible to members of that business group. Even an HRMS superuser (e.g. HR Director) can only see data pertaining to the same business group.

For example, an XYZ Organization has Consulting, Manufacturing, and Projects business groups. Consulting business group data can only be seen by the Consulting group members, not by the Manufacturing and Projects group

members.

Oracle E-Business Suite Organization Model



To secure the group data access only limited to the associated group members while retrieving search results, data is striped in the HRMS tables using `BUSINESS_GROUP_ID` identifier.

- **Business Group Search Plug-in (**
`oracle.apps.fnd.search.impl.BusinessGroupSearchPlugin`**) Java Class:** Use this seeded business group search plug-in to secure search results depending on the search user's access to one or more HRMS Business Groups. This search plug-in must extend the `DefaultSearchPlugIn` concrete Java class (`oracle.apps.fnd.search.impl.DefaultSearchPlugIn`) to ensure that the secure access to the target UI function is also validated.

- **Secure Attribute:** Since a secure attribute is needed to implement the Query Rewrite security mechanism, all HRMS tables are striped with `BUSINESS_GROUP_ID` column, which acts as a Business Group identifier.

Since a secure attribute is needed to implement the Query Rewrite security mechanism, all HRMS tables are striped with `BUSINESS_GROUP_ID` column, which acts as a Business Group identifier.

This plug-in requires that the search object definition should have `BUSINESS_GROUP_ID` attribute, which should be marked at least "Stored" and "Secured".

Legal Entity Based Search Security

- **Security Requirements:** Legal Entity in the Oracle E-Business Suite corresponds closely to "company" in the legal world. It has the right to own property, the right to trade, and the responsibility to comply with appropriate laws. You can store information about a registered company or other real world legal entity in the "legal entity".

The Ledger represents an accounting representation for an organization that is accountable in a self-contained way. A ledger owner might be a legal entity. Thus, the General Ledger (GL) application in Oracle E-Business Suite secures data at Legal Entity level.

To secure GL data during search, Legal Entity level security needs to be enforced by using the `LEDGER_ID` identifier in GL tables while retrieving search results.

- **Legal Entity Search Plug-in (`oracle.apps.fnd.search.impl.LegalEntitySearchPlugin`) Java class:** Use this legal entity search plug-in to restrict search results by the legal entities, which are accessible to the query user.

Similar to the business group search plug-in, this legal entity search plug-in must extend the `DefaultSearchPlugIn` concrete Java class (`oracle.apps.fnd.search.impl.DefaultSearchPlugIn`) to ensure that the secure access to the target General Ledger UI function is also validated.

- **Secure Attribute:** Since a secure attribute is needed to implement the Query Rewrite security mechanism, all GL tables are striped with `LEDGER_ID` column, which acts as a unique identifier as a business group identifier does to a business group security model.

This plug-in requires that the search object definition should have `LEDGER_ID` attribute, which should be marked at least "Stored" and "Secured".

Organization Based Search Security

Oracle E-Business Suite supports the concepts of multiple organizations as well as "Multiple Organizations Access Control (MOAC)" security model.

- **Security Requirements:** *Multiple Organizations* can be sets of books, business groups, legal entities, operating units, or inventory organizations. You can define multiple organizations and the relationships between them through a single installation of Oracle E-Business Suite.

Operating Units (OUs) are good examples of multiple organizations and they are often identified with security. (The "Manufacturing Units" mentioned in the Oracle E-Business Suite Organization Model diagram in the Business Group Based Search Security, page 20-57 closely resemble the multiple organizations.)

With MOAC security model, users are given access to the data they handle through "responsibilities". A responsibility is associated with a specific OU, or with several OUs. By securing application data in this way, users can access and process

transaction only for the particular operating unit or set of operating units to which they have been granted access. In other words, data pertaining to one organization is normally not visible to another organization unless the user has permission to transact across organizations.

While performing a search on such objects, the organization level data security needs to be enforced through `ORG_ID` identifier for a wide variety of business entities including Purchase Orders, Sales Orders, Payables Invoices, Receivables Invoices, etc.

- **Organization Search Plug-in (`oracle.apps.fnd.search.impl.OrganizationSearchPlugin`) Java Class:** Use this organization search plug-in to restrict search results by organizations, which are accessible to the query user.

Similar to the business group search plug-in, this legal entity search plug-in must extend the `DefaultSearchPlugIn` concrete Java class (`oracle.apps.fnd.search.impl.DefaultSearchPlugIn`) to ensure that the secure access to the target UI function associated with organizations is also validated.

- **Secure Attribute:** Since a secure attribute is needed to implement the Query Rewrite security mechanism, most Oracle Financial applications tables (e.g. AP, AR, FA etc.) are striped with `ORG_ID` column, which acts as an unique identifier as a business group identifier does to a business group security model.

This plug-in requires that the search object definition should have `ORG_ID` attribute, which should be marked at least "Stored" and "Secured".

Employee Hierarchy Based Search Security

This type of security model secures data based on employee hierarchy. Good examples can be iExpense, iProcurement, and iLearning within Oracle E-Business Suite. These application modules search on a particular employee's expenses, procurement, and training information based on employee hierarchy.

- **Security Requirements:** Based on the employee hierarchy, data for a particular employee is only visible within his reporting hierarchy. Therefore, while performing a search on such object, search results have to be secured through an unique person identifier.

- **Employee Hierarchy Search Plug-in (`oracle.apps.fnd.search.impl.EmployeeHierarchySearchPlugin`) Java Class:** Use this organization search plug-in to restrict search results by employee hierarchy in an organization, which are accessible to the query user.

Similar to other seeded search plug-ins, this legal entity search plug-in must extend the `DefaultSearchPlugIn` concrete Java class (`oracle.apps.fnd.search.impl.DefaultSearchPlugIn`) to ensure that the secure access to the target UI function associated with organizations is also

validated.

- **Secure Attribute:** Since a secure attribute is needed to implement the Query Rewrite security mechanism, most applications which require employee hierarchy, refer to `PER_ALL_PEOPLE_F` table. The `PERSON_ID` column is the unique person identifier (along with effective dates).

This plug-in requires that the search object definition should have `PERSON_ID` attribute, which should be marked at least "Stored" and "Secured".

Other Considerations

Oracle E-Business Suite Secure Enterprise Search allows various security rules to be added to secure your searchable objects and application content. However, there are some security limitations and performance need to be considered.

Limitations

If you have more than one security attribute implemented, the principle is that both security rules must be satisfied. This may prevent some use cases from working.

For example, purchase orders are allowed to be seen by buyer, approver, and accountant. However, the accountant is actually a role which can be held by different people at different times, while the buyer and approver are recognized by their employee Ids. If this case occurs, set the `emp_id` as a secure attribute. This way, when `getSecureAttrAcl` for `emp_id` is called, the `buyer_id` is returned along with a list of responsibilities that are granted to access purchase order. The logic is paired with `getSecureAttrKeys`, which basically returns the buyer's direct employee Id as well as their responsibilities.

Performance

Since a search plug-in is used both during crawl and query, it adds overhead to performance in various times of the object life cycle. This is especially true in `getAcl` and `getSecureAttrAcl` since these methods are called row by row.

Crawl Time Performance

For crawl time performance, there are two possible expensive operations when crawling a searchable object:

- Get Content

Get content implies JDBC calls to collect information to form indexable documents for Oracle SES to index. The performance of this operation depends on how well the view objects are defined, and how complex a searchable object is.

For example, if a searchable object assembles a large set of objects, such as a purchase order, it will take time to crawl because each document will need to source data from a number of tables (views).

- Get Acl

Get ACL is performed on a row-by-row basis by definition because each document

will have different ACLs. For a complex security model, `getAcl` might involve multiple database trips. This could incur a high cost and should be carefully balanced with query time performance in order to achieve overall performance.

Query Time Performance

Security also has impact on query time performance. This is due to the fact that for a securable searchable object, the query must be rewritten with access keys by calling `getSecurityKeys`. This function call usually involves database calls.

For example, during query execution, Oracle SES authorization plug-in mechanism contacts Oracle E-Business Suite Security Service End Point over HTTP protocol. The Security Service End Point is used to authenticate an Oracle E-Business Suite user and generate the Security Keys for the query user. The Security Service End Point is implemented as a servlet and registered as "AppSearch" servlet in `oafm` container. Therefore, any security service request is subject to the risk of HTTP Timeout. That is when Oracle SES authorization plug-in mechanism contacts the Security Service in Oracle E-Business Suite, the request has to be completed within a predefined amount of time.

As a guideline, the HTTP time-out value should be set to 30000 milliseconds. The time taken to execute the search plug-in is quite proportional to the overall execution time of a query. Hence for a responsive application, the order of execution has to be classified as follows:

1. Simple plug-in execution: 5000 milliseconds
2. Medium complexity plug-in execution: 10000 milliseconds
3. Complex plug-in execution: 20000 milliseconds

Please note that this has to be irrespective of the data volume, which a customer might encounter. Query time performance normally has higher priority than crawl time performance. It must be balanced on a case-by-case basis.

Note: During query, Oracle SES fetches the runtime keys for the current application user using `getSecurityKeys()` or `getSecureAttrKeys()`. Oracle SES waits for a predetermined but configurable amount of time for these methods to retrieve the results. In case of a timeout, Oracle SES assumes the security keys are null for the current user and caches them. Most of the cases, it results in getting no search hits. This is one of the foremost reasons of not getting desired search results.

Improving Query Time Performance Using Cache

Please note that query time performance can be greatly improved by using previously cached security access keys stored in Oracle SES for a particular user, data source, or object. This greatly reduces the query response time of synchronously fetching the

Security Keys for a user or gets timed out if the cache exists. For more information on how to use this feature, see User Authorization Cache, page 20-63.

User Authorization Cache

User Authorization Cache (UAC) framework provides a mechanism allowing the security access keys for a particular user, a specific data source, or a searchable object in Oracle E-Business Suite can be precrawled, cached, and stored in Oracle SES.

By leveraging the UAC feature from Oracle SES, when a user performs a search, instead of fetching the access keys synchronously for that user or object during user query, the previously cached Security Keys will be first looked up in SES for the availability of the keys for that user or object. If a match is found and the cache is usable, the associated keys will be used to compare with the pre-built ACLs. Any matched indexed documents will then be retrieved for the user. Unmatched or unauthorized documents get dynamically filtered out in the process. If there is no match found, the Security Keys will then be fetched and built security filters synchronously during the query. Any matched indexed documents based on the Security Keys and ACLs will be retrieved for the user.

Note: Although Oracle SES contains UAC feature, UAC for Oracle E-Business Suite will be fully enabled in a later release of Oracle SES. Full benefits of UAC will be visible only then.

By using the previously cached keys to authorize or revoke the document access privilege (in contrast of generating the keys real time during user query), this feature greatly reduces the search response time and in turn provides quick search results with security enforced.

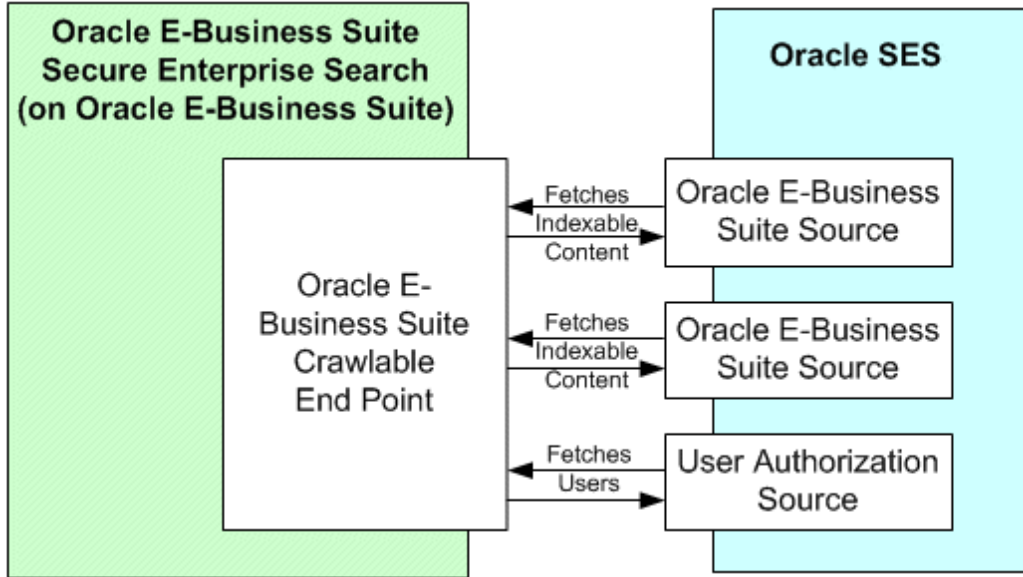
Note: Security Keys are generated through the execution of the `getSecurityKeys()` and `getSecureAttrKeys()` methods of a security plug-in. For more security plug-in information, see Search Security Plug-ins, page 20-43.

How Does User Authorization Cache Work

To enable UAC feature, a specific User Crawler should be initialized to crawl Oracle E-Business Suite users and provides User documents to Oracle SES while other crawlers are for crawling and indexing searchable documents.

The following diagram illustrates the high level crawler tasks:

Crawler Interaction Flow



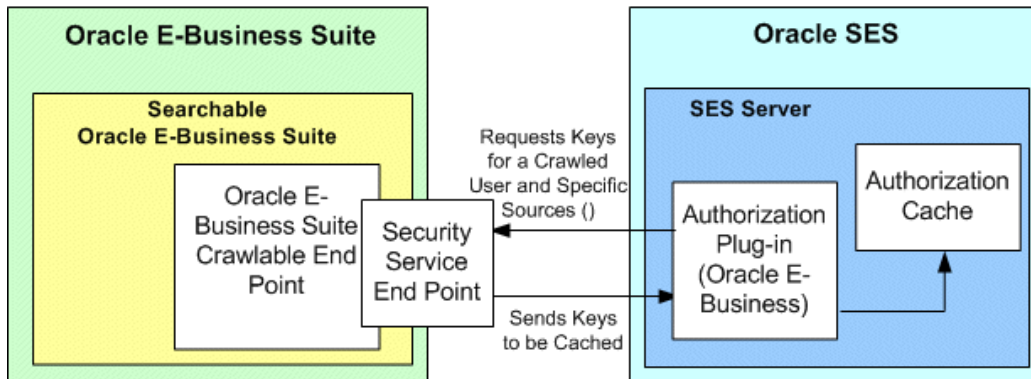
The User Crawler process generates a list of Oracle E-Business Suite users for whom the Security Keys need to be cached in Oracle SES for the predefined "User Authorization Cache" source type. See *Defining and Updating UAC Source*, page 20-65.

Fetching Security Keys Offline

In order to provide quick search results back to a user and eliminate possible time-outs of fetching Security Keys simultaneously during query due to complex application logic of deriving the keys, the user Security Keys can be generated as an offline process.

The following diagram illustrates the high level authorization cache population flow:

Authorization Cache Population High Level Flow Diagram



When a list of Oracle E-Business Suite user documents is generated by the User Crawler, Oracle SES invokes its Authorization Plug-in to contact Identity Management

to fetch the Security Keys. This is done by sending a request for the keys of a crawled user and a specific data source to Oracle E-Business Suite Security Service End Point. The Service End Point in turn invokes the search plug-in to generate the Security Keys which is executed in the same proxy session where the user credential (username and password) is verified. The `getSecurityKeys()` and `getSecureAttrKeys()` methods of a search plug-in are executed to generate the Security Keys. The Service End Point sends the generated keys for a given user and a specific searchable object or data source to be cached to the Authorization Plug-in and have them cached and stored in Oracle SES.

Using Cached Keys for Query Time

When a user performs a search, these previously cached Security Keys will be examined first in Oracle SES to see if the cache exists for a given source and user, as well as whether the cache is usable.

- If a match is found and the cache is usable, the associated keys will be used to compare with the pre-built ACLs. Any matched indexed documents will be retrieved for the user.
- If there is no match found, the Security Keys will be fetched real time during the query and built security filters. Any matched indexed documents based on the Security Keys and ACLs will be retrieved for the user. Unmatched or unauthorized documents get dynamically filtered out in the process.

Please note that the user authentication cache is populated only when the User Crawler gets executed at the crawl time. If the Security Keys are fetched synchronously during user query, the cache will not be populated.

Defining and Updating UAC Source

In order to crawl Oracle E-Business Suite users in a source system, a special "User Authorization Cache" source type should be defined.

Creating UAC Source

For Oracle E-Business Suite which has seamless integration with Oracle SES, whenever any searchable object is deployed either from Oracle E-Business Suite or from Search Modeler for the first time, one source of "User Authorization Cache" type will be transparently created. This source will have the following information:

Parameter Name	Value
Name	Oracle E-Business Suite UAC
Type	User Authorization Cache

Parameter Name	Value
User Search Query	<p>""</p> <p>Note: Query expression defines the set of users to be crawled. For example, a* means to crawl all users whose names begin with the letter a, and null value (*) means to crawl all Oracle E-Business Suite users.</p> <p>The SES Administrator can enter comma separated user names in wild card format, for example, OPERATIONS, BPALMER, SYSADMIN*. The names entered will be preserved during successive deployment of objects.</p>
Source names for which security attributes should be crawled	<p>This parameter will have comma separated values of Sources already deployed. These source names will be automatically updated on deployment of objects.</p>

The above information lets Oracle SES know about the Oracle E-Business Suite "Sources" for a specific user for which the security keys need to be fetched. In other words, the UAC source maps users with the sources whose security keys need to be cached.

Updating UAC Source

Any searchable object deployed subsequently will automatically update the "Source names for which security attributes should be crawled" parameter value to include the name of the source currently being deployed.

Manually Updating UAC Source

However, in case a "User" crawl is in progress while such deployment is attempted, the update to the UAC source might fail. Therefore, a Search Administrator might have to manually add the source name later on.

To manually update the UAC source, log on to Oracle SES and select Sources tab. For Source Type, select "User Authorization Cache" and then click the Update icon. The Update User-Defined Source page is displayed allowing you to update the source information.

Managing the UAC Crawling Schedules

A Search Administrator needs to schedule the crawling job of "E-Business Suite UAC" source at a regular interval to fulfill your business needs.

How to set the crawling frequency and manage crawling schedules, see *Administering Crawls in Oracle SES*, page 20-72 and *Managing Crawling Schedules*, page 20-79.

Administering Searchable Objects

Searchable objects are business objects that are made available for text search; they are used in an abstract way for exposing business data to search engines. For example, a purchase order as a searchable object would be defined as a set of searchable properties and its relationship to other searchable objects. Oracle E-Business Suite Secure Enterprise Search uses this abstraction concept to group objects in a logical way at runtime.

To secure all searchable objects containing sensitive application context only exposed to appropriate users before they are deployed to Oracle SES, and to effectively manage and administer data sources after the deployment, system administrator or search administrator needs to perform the following tasks:

1. *Securing Searchable Objects Using Security Grants*, page 20-67

Before deploying searchable objects to Oracle SES and making them available to users, these objects must be secured first. By leveraging the concept of the Role-Based Access Control (RBAC) security model, administrators can create security grants to ensure the application content sensitive data is only accessed by authorized people.

2. *Deploying Searchable Objects to Oracle SES*, page 20-68

Once searchable objects are ready to be deployed, the system administrator or search administrator can deploy them to the Oracle SES instance. Since not all searchable objects can be successfully deployed to Oracle SES, deployment guidelines and additional tasks are described in this section.

3. *Administering Crawls in Oracle SES*, page 20-72

Once searchable objects are deployed to the Oracle SES instance, the crawling schedules by default are automatically created in Oracle SES and visible in the Oracle E-Business Suite. The system administrator or search administrator must first manually edit the default schedules with desired crawling frequencies and start the initial crawl.

Securing Searchable Objects Using Security Grants

As soon as a searchable object is created and patched into Oracle E-Business Suite, it is crawlable in Oracle SES. To make it available for users to search without compromising the data integrity and content sensitivity, the security context must be constructed around the searchable object first. By utilizing the Role-Based Access Control (RBAC) security model, Oracle E-Business Suite Secure Enterprise Search provides a flexible solution that can easily embed application security into a full text search service, and

this solution allows only authorized users with appropriate access privileges to search on or view applications data against a preindexed Oracle SES store.

For more information about the RBAC model and how to create security grants, see Role-Based Access Control (RBAC) Security, page 20-39.

Deploying Searchable Objects to Oracle SES

Once searchable objects are ready to be deployed to the Oracle SES instance that you set up earlier in the Configuration tab, the system administrator or search administrator can deploy a single object or deploy all objects simultaneously from a search.

The deployment process can create the following items in Oracle SES:

- Create a data source for each deployed searchable object
- Create an 'Oracle E-Business Suite Release 12' data source type associated with each data source
- Create a schedule per data source
- Create all source groups

This creation includes a source group per data source, and a source group per permission set.

Important: Once searchable objects are deployed to Oracle SES, default schedules for each searchable object are created automatically in Oracle SES, but they are set to have a manual launch for the initial crawl. A system administrator or search administrator must manually edit the default schedule by setting up crawling frequency through the use of the administrative page in Oracle SES and starting the initial crawl. Otherwise, the initial crawl will not be automatically started. For more information on setting up crawling frequency and starting an initial crawl, see *Administering Crawls*, page 20-72.

Please note that this synchronization process with an Oracle SES instance can only deploy the objects that have never been deployed to Oracle SES. Once they are deployed, any future deployment will not update the Oracle SES instance unless you manually delete the data source of the same name in Oracle SES and redeploy it again. Also, if you change the proxy user name and password, business objects that have already been crawled cannot be updated or resynchronized with the Oracle SES instance.

For more details on deployment, see *Deployment Concepts and Guidelines in Oracle SES*, page 20-70.

Deploying Searchable Objects

ORACLE Application Search Administrator

Diagnostics Home Logout Preferences

Application Search Administration

Schedules Configuration **Searchable Objects**

Search

Display Name

Name

Hide More Search Options

UI Function Name

Driving Table

Source File Name

Source File Product

Previous 1-5 Next 5

Details	Name	Display Name	Description	Last Update	Last Crawled	Deploy
<input type="checkbox"/> Show	oracle.apps.fnd.rep.server.AllIREPClassesVO	AllIREPClassesVO	IREP classes	23-Jul-2007 06:09:02	25-Jul-2007 05:00:00	<input type="button" value="Deploy"/>
<input type="checkbox"/> Show	oracle.apps.fnd.rep.server.AllIREPInterfacesVO	AllIREPInterfacesVO	IREP classes	20-Jul-2007 06:23:41		<input type="button" value="Deploy"/>
<input type="checkbox"/> Show	oracle.apps.fnd.rep.server.AllIREPObjectsVO	AllIREPObjectsVO	Fnd Objects	23-Jul-2007 06:02:51	25-Jul-2007 05:00:00	<input type="button" value="Deploy"/>
<input type="checkbox"/> Show	oracle.apps.fnd.wf.worklist.server.AllNotificationsVO	AllNotificationsVO	Notifications Search	27-Jul-2007 04:24:01		<input type="button" value="Deploy"/>
<input type="checkbox"/> Show	oracle.apps.eam.asset.server.AppSearchVO	AppSearchVO	Asset Search	25-Jul-2007 11:46:45		<input type="button" value="Deploy"/>

About this Page Privacy Statement

Diagnostics Home Logout Preferences

Copyright (c) 2006, Oracle. All rights reserved.

To deploy searchable objects:

1. Log on to Oracle E-Business Suite with the Application Search Administrator responsibility and select the Application Search Administration link from the Navigator window.
2. From the Application Search Administration window, select the Searchable Objects tab.
3. Enter simple search criteria in the Search region, such as Display Name and Name fields. Click **Go** to execute the search.

Optionally, click the **Show More Search Options** link to enter more search criteria, such as UI Function Name, Driving Table, Source File Name, and Source File Product fields.

4. From the search result table, you can choose one searchable object you want to deploy to an Oracle SES instance and click the **Deploy** icon for the object.
5. Click **Deploy All** to deploy all the objects from the result table.

Note: The selection of **Deploy All** is to deploy all the objects from the search result to an Oracle SES instance, and this would mean a reload of Oracle SES references.

- Click the **Show** link for an object to view the object details. These details include the searchable object's properties information and detailed breakdown for each object member's attributes whether it is displayed, titled, indexed, stored, or secured.

Displaying Searchable Object Details

The screenshot shows the Oracle Application Search Administrator interface. At the top, there is a navigation bar with 'ORACLE Application Search Administrator' and links for 'Diagnostics', 'Home', 'Logout', and 'Preferences'. Below this, the 'Application Search Administration' section is active, with sub-tabs for 'Schedules', 'Configuration', and 'Searchable Objects'. The 'Searchable Objects' tab is selected, showing a search form with 'Display Name' and 'Name' (AppSearch) fields, and buttons for 'Go', 'Clear', and 'Deploy All'. Below the search form, a table lists search results. The first result is expanded to show detailed information:

Details	Name	Display Name	Description	Last Update	Last Crawled	Deploy
<input type="checkbox"/> Hide	oracle.apps.eam.asset.server.AppSearchVO	AppSearchVO	Asset Search	25-Jul-2007 11:46:45		<input type="checkbox"/>
Object Name	VIEWOBJECT:/oracle/apps/eam/asset/server/AppSearchVO		Source Product	Enterprise Asset Management		
SearchPlugin	oracle.apps.eam.asset.webui.AlmSearchPlugin		Source File Path	java/asset/server		
Driving Table	csi_item_instances cii		Source File Name	AppSearchVO.xml		
UI Function	EAM Self Service Asset Details Page		Source File Version	12.0		
UI Function Display Name	EAM_ASSET_DETAILS					
UI Params	instanceId= <InstanceId> & OrgId= <MaintOrgId>					

Below the details, there are sections for 'Incremental Crawl Triggers' (Change Event, Date Based Columns) and 'Searchable Object Members'. A table lists members with columns: Name, Type, Description, Is Displayed, Is Title, Indexed, Stored, and Is Secure. The table contains 13 rows of data. At the bottom, there is a 'Show' link for another object: 'oracle.apps.csi.instance.search.server.AssetAppSearchVO'.

Click the **Hide** link to close the detailed view.

Deployment Concepts and Guidelines in Oracle SES

In addition to deploying searchable objects to an Oracle SES instance, the system administrator or search administrator must be aware of the following concepts and may need to perform additional tasks if necessary:

- Data sources
- Crawling schedules
- Data source groups

Data Sources

A data source is one kind of data that you might want to search on. For example, if your data is in Web pages, then Web source is your data source. In other words, it is a particular end point where data can be retrieved. Each data source has a data type associated with it, such as Oracle E-Business Suite 12. Searchable objects contain many business attributes and these attributes can be retrieved and indexed for a given data source type during crawling.

Since the deploy process will only synchronize objects that they have never been deployed to an Oracle SES instance, if an object has been deployed, any future deployment for the same object will not update the instance unless you manually delete the data source of the same name in the instance and then deploy it again. Also, if you change the proxy user name and password, business objects that have already been crawled cannot be resynchronized with the Oracle SES instance.

Deployment Guidelines

Use the following guidelines to have searchable objects successfully deployed to the Oracle SES instance:

- If a data source of the same name does not exist, create the data source with new parameters.
- If a data source is created, you can only delete the data source if it is not crawled and create the data source with new parameters.
- If a data source exists and has been crawled already, do not delete the source or create a new one. Instead, you need to manually update the data source parameters in the Oracle SES administrative pages and redeploy it from the E-Business Suite.

To access the Oracle SES administrative user interfaces, select the Configuration tab in the Applications Search Administration page. In the Tasks region, click the **SES Admin Login** link under the Additional Tasks section to navigate to the Oracle SES administrative user interfaces. In the Oracle SES instance, use the Home tab and the Source subtab to edit or create data sources. For more information on data source creation and parameters, see the *Oracle Secure Enterprise Search Administrator's Guide* for details.

Crawling Schedules for Searchable Objects

Once searchable objects are deployed to Oracle SES, default schedules for each searchable object are generated automatically in Oracle SES, but they are set to have manual launch the initial crawl. It is very important that system administrator or search administrator must manually edit the default schedule by setting up crawling frequency through the use of the administrative page in Oracle SES and starting the initial crawl. Otherwise, the initial crawl will never be automatically started. See: *Administering Crawls in Oracle SES*, page 20-72

During the data deployment process, if the data source gets created, the existing schedule will be deleted and a new one should be created. However, for a data source that has been crawled already, its schedule will not be recreated.

Data Source Groups

A data source group is a concept used in Oracle SES to group a number of crawled indexes for an aggregated search. For each searchable object, a default data source group is created with the same name which includes only the data source for this object.

To enable Oracle SES to perform searches on groups, all the groups that have been created in the E-Business Suite application instance should have corresponding source groups created in the Oracle SES instance as well.

For example, for each permission set that starts with `SESG`, a data source group will also be created and populated with Oracle SES references. The permissions included in the permission set that is linked to a searchable object will have their data sources included in the group.

For example, a **permission set** `SESG_SEARCH_CRM` includes the following permissions:

- **ServiceRequest** permission (contains 'ServiceRequest' searchable object)
- **Customer** permission (contains 'Customer' searchable object)
- **Contract** permission (contains 'Contract' searchable object)

Oracle E-Business Suite Secure Enterprise Search uses this mechanism to allow an application user to perform text search in a searchable group and refine or narrow down the search result using the searchable objects contained in the group:

- Each individual searchable object within the group `SESG_SEARCH_CRM`:
 - ServiceRequest
 - Customer
 - Contract

For more information on permission sets used in building security context, see *Securing Searchable Objects Using Security Grants*, page 20-67.

Administering Crawls in Oracle SES

Crawling schedules define the frequency at which the index is updated with information about each source. Once searchable objects are deployed to Oracle SES, crawling schedules are automatically created along with the data sources in Oracle SES and visible in the Oracle E-Business Suite. However, these automatically created crawling schedules have the crawling frequency type set to the default value 'Manual Launch' which requires you to manually start the initial crawl. Otherwise, these schedules will never be started automatically.

Note: The initial crawl refers to the first time a searchable object is crawled. Since it usually involves a large set of data, it is highly

recommended that an initial crawling job should be scheduled by a low bandwidth job in non-peak hours.

To have fast performance on initial crawl, Oracle E-Business Suite Secure Enterprise Search uses the AD Parallel Update package to help split the large data set into smaller work units, and crawl the units in parallel by using the multi-thread crawling mechanism provided by Oracle SES.

If you want the source or index updated more frequently after the initial crawl is completed, you can update the crawling frequency for a schedule in the Edit Schedule page through the Oracle SES administrative UI.

Setting Crawling Frequency

Edit Schedule

Schedule Name

Assignment

To assign a source, select one or more available sources and click >>. After a source has been assigned to a group, it cannot be assigned to any other group. To undo assignments of a source, select one or more scheduled sources and click <<.

Get Available Sources for Type:

----- Available Sources -----	>>	<<	----- Assigned Sources ----- [Oracle E-Business Suite R12] oracle.apps.cs.csz.incident.server.CSSecureVO
-------------------------------	----	----	---

Note: * indicates source with "Delete password after crawl" option set

Update Crawler Recrawl Policy

When the crawler retrieves a Web, file, or table source document, it checks to see if that document has changed. By default, if the document has not changed, then the crawler does not process it. This significantly speeds up the crawling process. However, in certain situations, it might be desirable to force the crawler to reprocess all documents.

- Process Documents That Have Changed
- Process All Documents

Update Crawling Mode

This section lets you update the crawling mode.

- Automatically Accept All URLs for Indexing
- Examine URLs Before Indexing
- Index Only

Frequency

Select a frequency type. Be sure to specify all required data for the option. With "Manual launch", the schedule is never automatically started. You must manually start the schedule.

Frequency Type

- hourly
- daily
- weekly
- monthly
- Manual launch

To set crawling frequency in Oracle SES administrative page from Oracle E-Business Suite Secure Enterprise Search:

1. Log on to Oracle E-Business Suite with the Applications Search Administrator

responsibility.

2. Select the Configuration tab and click the **SES Admin Login** link from the Tasks region. This opens the Oracle SES login page.
3. Enter the user name and password you defined for an administrator in order to access an Oracle SES instance.
4. In Oracle SES, select the Home tab and Schedules subtab to access the Crawler Schedules page.
5. Select a schedule name and click the **Edit** icon to see the Edit Schedules page.
6. The selected schedule name is populated automatically in the Schedule Name field. You can select another schedule to update it if you want.
7. Leave the Assignment and Update Crawler Recrawl Policy regions unchanged with the default values.
8. In the Update Crawling Mode region, leave the **Automatically Accept All URLs for Indexing** radio button selected. This selection crawls and indexes all URLs in the source. It also extracts and indexes any links found in those URLs. If the URL has been crawled before, then it will be reindexed only if it has changed.
9. In the Frequency region, change the frequency type from the default 'Manual Launch' to daily, hourly, weekly, or monthly. Click **Update Frequency**.
10. Click **Finish** to save your changes.

Starting an Initial Crawl

ORACLE Secure Enterprise Search

Home Search Global Settings

General Sources Schedules Statistics

Crawler Schedules

Start Stop Previous 1-20 of 54 Next 20 Create

Select	Schedule Name	Status	Sources	Type	Log File	Last Crawled	Next Crawl	Edit	Delete
<input type="radio"/>	Mailing list Schedule	Failed	All mailing list sources	Mailing list		Aug 1, 2007 1:02:07 PM	Aug 1, 2007 2:00:00 PM		
<input checked="" type="radio"/>	oracle.apps.cs.csz.incident.server.CSSecureVO	Scheduled	oracle.apps.cs.csz.incident.server.CSSecureVO	Oracle E-Business Suite R12					
<input type="radio"/>	oracle.apps.cs.csz.incident.server.CsFormVO	Scheduled	oracle.apps.cs.csz.incident.server.CsFormVO	Oracle E-Business Suite R12					
<input type="radio"/>	oracle.apps.cs.csz.incident.server.CsSimpleVO	Scheduled	oracle.apps.cs.csz.incident.server.CsSimpleVO	Oracle E-Business Suite R12		Aug 31, 2007 2:56:26 AM			
<input type="radio"/>	oracle.apps.cs.csz.incident.server.IncidentSearchVO	Scheduled	oracle.apps.cs.csz.incident.server.IncidentSearchVO	Oracle E-Business Suite R12					
<input type="radio"/>	oracle.apps.cs.csz.incident.server.SRSearchVO	Scheduled	oracle.apps.cs.csz.incident.server.SRSearchVO	Oracle E-Business Suite R12		Aug 7, 2007 6:56:17 AM			

To start, stop, or delete a crawl in Oracle SES:

1. Log on to Oracle E-Business Suite with the Applications Search Administrator responsibility.
2. Select the Configuration tab and click the **SES Admin Login** link from the Tasks region. This opens the Oracle SES login page.
3. Enter the user name and password you defined for an administrator in order to access an Oracle SES instance.
4. In Oracle SES, select the Home tab and Schedules subtab to access the Crawler Schedules page.
5. Select a schedule name that you want to start the initial crawl and click **Start**. If you want to stop an existing crawl, select the schedule name and click **Stop** or click **Delete** to delete a schedule.
6. To update a schedule, select a schedule name and click **Edit**. See: To set crawling frequency in the Oracle SES administrative page from Oracle E-Business Suite Secure Enterprise Search, page 20-74.
7. To view a schedule status, click the link in the Status column, such as scheduled, disabled, launching, or failed, to see the schedule details.
8. Click the **Log File** icon to see detailed crawler settings and status.
9. Click **Create** to manually create a new schedule.

For more information on managing crawling schedules in Oracle SES, see the *Oracle Secure Enterprise Search Administrator's Guide* for details.

Testing Oracle E-Business Suite Secure Enterprise Search Setups

Use the following sections to validate whether you have successfully set up the Oracle E-Business Suite Secure Enterprise Search:

- Validate General Setups, page 20-77
- Test Deployment, page 20-77
- Test Schedules, page 20-78
- Test Searches, page 20-78

Validating General Setups

Use the following steps to validate general setups in Oracle E-Business Suite Secure Enterprise Search:

1. Test whether you have set the FND: Search Enabling Flag profile value to Yes. If it is not set to Yes, crawling should be disabled.
2. Assign the FND Search Crawler (SES_SEARCH_CRAWLER) responsibility and Application Search Administrator responsibility to a system administrator or search administrator. This administrator must be a valid FND user used as a proxy user for query.
3. Ensure you have set the correct value for the proxy parameters. To verify, log on to Oracle E-Business Suite with the Application Search Administrator responsibility, and select Configuration tab to view your setup parameters.

Use the **Update** check box to reset SES admin proxy and query proxy. For example, set SES admin proxy with user name `eqsys` and the password specified during installation; query proxy with user name `sysadmin` with appropriate password. The query user name must be a valid FND user with FND Search Crawler responsibility.

Important: Once you change the query proxy parameters, the SES instance needs to be restarted to reflect the changes.

Testing Deployment

Use the following steps to test whether you can deploy an object:

1. Log on to Oracle E-Business Suite with the Application Search Administrator responsibility.

2. Select the Searchable Objects tab and search for the object that you want to deploy.
3. Select the object to be deployed and click the **Deploy** icon.

Testing Schedules

Once searchable objects are deployed to the Oracle SES instance, you should be able to find their corresponding schedules automatically created in Oracle SES. Use the Oracle SES instance to start the crawling schedules.

Use the following steps to test crawling schedules whether they work properly:

1. Log on to the Oracle SES administrative page through the Configuration tab in the Application Search Administration page.
2. Select the Home tab and Schedules subtab. Refresh the page and you should be able to see the schedule for object you just deployed.
3. Select the schedule and click **Start** to observe the schedule status change for the selected schedule. Refresh the page if necessary to view the status updates.

Testing Searches

Once the setup tasks are completed, application users with appropriate privileges should be able to perform searches within the Oracle E-Business Suite.

Use the following steps to perform searches:

1. From the home page of the Oracle E-Business Suite, select a searchable group from the search drop-down list.
2. Enter a keyword in the text field, such as 'oracle' and click **Go**.

You should be able to find the search results populated in the results region.

Additional Administrative Tasks

In addition to setting up necessary tasks for Oracle E-Business Suite Secure Enterprise Search to ensure its seamless integration with Oracle SES, and performing administrative tasks to secure and deploy searchable objects, the system administrator and search administrator also need to perform the following tasks to proactively manage crawling schedules and optimize indexes:

- Managing Crawling Schedules, page 20-79
- Optimizing Indexes, page 20-81

Managing Crawling Schedules

Once searchable objects are deployed, crawling schedules are automatically created along with data sources in Oracle SES. After an initial crawl is completed, subsequent incremental crawls are scheduled and can be executed automatically triggered by business events, date changes, or crawling frequency, as well as other necessary manual crawls.

When a crawling job starts, each crawler retrieves business objects of a given type and then pushes the retrieved objects to be indexed by Oracle SES indexers. Finally, these objects with indexes are stored in the Oracle SES index store for user queries.

For example, a searchable object, such as a purchase order, may have source data from a number of tables (views), such as product description, employee e-mail address, and so on. When these fields change, the last updated date for the purchase order is also updated. In this way, when a scheduled crawl is performed, the purchase order gets reindexed and stored in the Oracle SES index store.

Oracle E-Business Suite Secure Enterprise Search allows the administrator to proactively manage the crawling schedules in the following ways:

- View the latest crawling status, last crawled, and next crawling schedule for a given schedule
- Stop a manual crawling job
- Recrawl all the data for a selected schedule
- Create an incrementally crawling job for a selected schedule once the initial crawling is completed

- The initial crawl refers to the first time a searchable object is crawled. Since it usually involves a large set of data, it is highly recommended that an initial crawling job should be scheduled by a low bandwidth job in non-peak hours.

To have fast performance on initial crawl, Oracle E-Business Suite Secure Enterprise Search uses the AD Parallel Update package to help split the large data set into smaller work units, and crawl the units in parallel by using the multi-thread crawling mechanism provided by Oracle SES.

- The incremental crawl refers to crawling the data to the original data source after the initial crawl.

Managing Crawling Schedules

ORACLE Application Search Administrator

Diagnostics Home Logout Preferences

Application Search Administration

Schedules Configuration Searchable Objects

Select Object: Incremental Crawl Force Crawl Stop Crawl |

Previous 1-10 of 31 Next 10

Select Schedule Name	Status	Sources	Last Crawled	Next Crawl
<input type="radio"/> Mailing list Schedule	failed	Mailing list Source	01-Aug-2007 18:02:07	01-Aug-2007 19:00:00
<input type="radio"/> oracle.apps.ota.learner.server.SimpleSearchVO	scheduled	oracle.apps.ota.learner.server.SimpleSearchVO	10-Jul-2007 06:49:59	Manual
<input type="radio"/> oracle.apps.ota.learner.search.server.SrchCourseVO	failed	oracle.apps.ota.learner.search.server.SrchCourseVO	03-Aug-2007 09:02:19	Manual
<input type="radio"/> oracle.apps.fnd.search.demo.FNDApplicationVO	scheduled	oracle.apps.fnd.search.demo.FNDApplicationVO	25-Jul-2007 15:07:18	Manual
<input type="radio"/> oracle.apps.ego.item.eu.server.EgoItemSesVO	scheduled	oracle.apps.ego.item.eu.server.EgoItemSesVO	01-Aug-2007 08:54:27	Manual
<input type="radio"/> oracle.apps.cs.csz.incident.server.CSSecureVO	failed	oracle.apps.cs.csz.incident.server.CSSecureVO	31-Jul-2007 05:10:14	Manual
<input type="radio"/> oracle.apps.fii.mtrtps.server.MetricsSesVO	scheduled	oracle.apps.fii.mtrtps.server.MetricsSesVO	17-Jul-2007 10:18:32	Manual
<input type="radio"/> oracle.apps.eam.asset.server.AppSearchVO	scheduled	oracle.apps.eam.asset.server.AppSearchVO	31-Jul-2007 09:44:59	Manual
<input type="radio"/> oracle.apps.po.contracts.server.PODemoVO	scheduled	oracle.apps.po.contracts.server.PODemoVO	06-Jul-2007 09:33:54	Manual
<input type="radio"/> oracle.apps.cs.csz.incident.server.IncidentSearchVO	scheduled	oracle.apps.cs.csz.incident.server.IncidentSearchVO	13-Jul-2007 19:12:43	Manual

Select Object: Incremental Crawl Force Crawl Stop Crawl |

Previous 1-10 of 31 Next 10

Diagnostics Home Logout Preferences

About this Page Privacy Statement Copyright (c) 2006, Oracle. All rights reserved.

To manage crawling schedules:

1. Log on to Oracle E-Business Suite with the Application Search Administrator responsibility and select the Application Search Administration link from the Navigator window.
2. From the Application Search Administration window, select the Schedules tab.
3. From the Schedules page, you can view the crawling details for a given schedule including schedule name, crawling status, source, last crawled, and next crawling schedule.
4. To view the latest schedule details, click the **Refresh Crawler Schedules** icon to get the schedule refreshed.
5. To create incremental crawling schedules, select a schedule name by clicking the **Select** radio button and click **Incremental Crawl** to have the next crawling schedule created in the Next Crawl field.

Incremental crawling can be raised by a business event or a date change to a searchable object since the last time it was crawled.
6. To stop an existing crawling job, after selecting a schedule name, click **Stop Crawl** to stop the crawling job for the selected schedule.

7. To recrawl all the data for a selected schedule, click **Force Crawl**.

Optimizing Indexes

Crawlers maintain active indexes of all documents crawled over all sources. To reduce fragmentation from crawls and increase the speed of searches, the administrator needs to create schedules for optimizing indexes through the Oracle SES administrative pages.

Oracle E-Business Suite Secure Enterprise Search also facilitates the index optimization performed in Oracle SES through a request.

Optimizing Indexes

The screenshot shows the Oracle Application Search Administrator interface. The top navigation bar includes 'Diagnostics', 'Home', 'Logout', and 'Preferences'. The main content area is titled 'Application Search Administration' and has three tabs: 'Schedules', 'Configuration', and 'Searchable Objects'. The 'Configuration' tab is active, showing three sections: 'Parameters', 'Admin Proxy', and 'Query Proxy'. The 'Parameters' section includes a 'SES End Point' field with a value of 'http://ap6009rems.us.oracle.com:7777' and an 'Update' button. The 'Admin Proxy' section includes 'User Name' (edqsys), 'Password' (masked), and 'Time Out' (1200 Sec) fields, with an 'Update' button. The 'Query Proxy' section includes 'User Name' (sysadmin), 'Password' (masked), and 'Time Out' (1200 Sec) fields, with an 'Update' button. On the right side, the 'Tasks' section is titled 'Optimize Indexes' and contains a description, a note about optimization during low usage, and an 'Optimize Index' button. Below this is an 'Additional Tasks' section with a link to 'SES Admin Login'.

To optimize indexes:

1. Log on to Oracle E-Business Suite with the Application Search Administrator responsibility and select the Application Search Administration link from the Navigator window.
2. From the Application Search Administration window, select the Configuration tab.
3. Click **Optimize Index** in the Optimize Indexes section. This raises an optimization request to Oracle SES and the indexes get optimized.

Important: In order to have minimal disruption to users, it is highly recommended that the index optimization should be done during hours of low usage.

For more information on optimizing indexes in Oracle SES, see the *Oracle Secure Enterprise Search Administrator's Guide* for details.

Error Messages

The following is a list of seeded error messages that Oracle E-Business Suite Secure Enterprise Search uses to notify or alert users when violations occur in interacting with the Oracle SES engine or during query:

Error Message Code	Description
FND_SEARCH_SECURITY	<p>This message occurs when security rules are violated by a query. The query module will terminate the process and throw security exception along with this message.</p> <p>Parameters in this message might include current FND user name. User-specified filters in a secured attribute is a security error. For example, you can enter keyword "oracle" to query. However, if you query on "EMP_ID:dlam content:oracle", an error message is returned because 'EMP_ID' is a secured attribute.</p>
FND_SEARCH_TOO_MANY_ENTRIES	<p>This message occurs when the query engine is to perform a post-query row-by-row process and there were too many rows. The query engine might perform some heuristic actions or throw an exception to the API user.</p>
FND_SEARCH_SYNTAX_ERROR	<p>This message occurs when the query syntax does not conform with Oracle SES. The query engine might rewrite the query.</p>
FND_SEARCH_SES_ERROR	<p>This message relays any potential error originating from Oracle SES when Oracle E-Business Suite Secure Enterprise Search interacts with the Oracle SES engine.</p>

The following table lists the error message type for the types of errors that occur during

the integration with Oracle SES engine or query:

Error Message Type	Description
FND_SEARCH_0001	Indicates search engine general errors.
FND_SEARCH_0002	Indicates security errors.
FND_SEARCH_0003	Indicates crawl time errors.
FND_SEARCH_0004	Indicates metadata errors.
FND_SEARCH_0005	Indicates query errors.

Administering Process Navigation

Overview of Process Navigation

A "process" is a series of actions taken to achieve a specific result. The Process Navigator utilizes Oracle Workflow to depict each of your business processes with a workflow diagram. A process diagram contains an icon for each step in the process; each icon acts as a visual cue and as an access point for the actual form associated with each step. You can navigate to any form involved in the process simply by clicking on the appropriate icon.

What is Oracle Workflow?

Oracle Workflow allows you to define business processes using a drag-and-drop designer. You can route relevant information to decision makers, automate processes, deliver electronic notifications to users in a given workflow, and monitor your processes as they are implemented. You can display any workflow diagram as a process in the Process Navigator. For more information, see the Oracle Workflow documentation.

What are Seeded Processes?

A seeded process is one that is delivered to you ready to use. Oracle E-Business Suite includes several seeded business processes which you can use as they are.

Modifying Your Menu

Before you begin, you should be aware that simply referencing a form from a process does not provide the required permissions for the responsibility to access the forms in the process. Form Functions for each form referenced from a process must be added to the Function Security Menu for the responsibility. If the Form Function is not accessible, the user will receive an error when attempting to access the form from the process in the Process Navigator.

Creating Process Navigator Processes

You must use Oracle Workflow Builder to create or customize any of the processes that are displayed in the Process Navigator. These instructions describe how to create new processes for the Process Navigator.

The following table lists the terms/components of a Process Navigator process and the corresponding components in Oracle Workflow Builder that define them.

Process Navigator Component	Description	Controlling Oracle Workflow Builder Component(s)
Process	The diagram that appears in the Process Navigator.	Process activity and process diagram
Process description	A description of the displayed process.	Process activity
Step	An icon in the process, which takes you directly to an Oracle E-Business Suite form when you double-click on it.	Notification activity
Step description	A description of the selected process step.	Message
Form associated with a step.	The Oracle E-Business Suite form that appears when you double click on a step in a Process Navigator process.	Form-type Message attribute

Note: The following procedures do not address most of the functionality of Oracle Workflow Builder, but are tailored to creating processes for the Process Navigator. The Oracle Workflow Builder is a tool used to design workflow processes. Workflow processes can range from routing documents through an approval process to setting up your Oracle E-Business Suite. See the Oracle Workflow documentation for more information.

Creating Process Navigator Processes

To create a new process for the Process Navigator, you must first create the necessary components in Oracle Workflow Builder. The components you create make up the

process definition, which is then saved to the database or to a flat file. The Process Navigator then reads the process definition from the database to display the process and its information and provide you access to the related Oracle E-Business Suite forms.

Creating a New Process Navigator Process

Note: For more information on creating a process, see the *Oracle Workflow Developer's Guide*.

1. Open Oracle Workflow Builder.
2. Create an item type. An item type is a repository that will contain all the components associated with the process you wish to build.
3. Create an Item Attribute of type role, whose internal name is USER_NAME.

Note: Enter a new display name for the message using the format <Verb><Form Title>. If the form title already contains a verb, then simply use the form title as the display name. If the form title does not contain a verb, then consider using one of the following verbs:

Define / Assign / Run / Load / Convert / Open / Set /
Generate / Review

4. Create a message to describe the task that is to be accomplished by a Process Navigator process step.
5. Create a form-type for the message. The seeded processes generally assign these message attributes an internal name of Open Form, but this is not required.
6. Create a notification activity to represent a Process Navigator process step.
7. Create a process activity to represent a Process Navigator Process.

Note: Enter a display name for your process. This name appears in the Process Navigator's process list. The naming convention for the process should be a functional name followed by the word "Process."

Enter a description for your process. The description appears when a user selects a process in the Process Navigator. The description is limited to 240 characters.

8. Draw the Process Diagram. Once you create a process activity, you can draw the process diagram that is associated with it. The process diagram is what appears

when you display a process in the Process Navigator.

Note: The Performer type of the Notification Activity you include in a process diagram for the Process Navigator must be set to the item attribute USER_NAME.

9. Save your changes. When you save your work to a database, you actually save everything in the current data store that has been modified. When you save your work to a flat file, you actually save everything in the current data store to the file.

Note: It is highly recommended that for new processes created for the Process Navigator that you always save a copy of your workflow process definition as a flat file and check that file into a source control system to maintain a working version of your process definition. Then when you want to update your definition in the database, you can pull up the flat file and save it directly to the database. Avoid using the process definition stored in your database as your source controlled version, as others with access to the database can update the definition.

10. Enable access to your process.

Enable access to your process

Before a process may be accessed in the Navigator you must complete the following two steps. Create a new function for your process in the Form Functions window, and add your process to a responsibility by adding the function you just created, to the responsibility's top menu in the Menus window.

Create a function for your process

Use the following procedure to create a function:

1. As the System Administrator navigate to the Form Functions window (Application->Function).

2. Enter a Function Name for your process using the format:

`<app>_<processname>`

Where `<app>` can be any application short name and `<processname>` is the internal name you entered when you created your process activity.

3. Enter a User Function Name. The name you enter here appears in the Navigator.
4. Enter "PROCESS" as your function type.

5. In the tabbed region 'Form' use the following format to enter a value in the Parameters field:

`<itemtype>:<processname>`

6. Save your work. No other fields are required to create your process function.

Add your function to a menu

In order for a user to access a process in the Navigator, the process must be added to a menu referenced by the user's responsibility. To determine the menu referenced by a particular responsibility use the Responsibilities window (Security->Responsibility->Define).

1. As the System Administrator navigate to the Menus window (Application->Menu).
2. Use the Find window to access the desired menu.
3. In a new row use the LOV to select the function you created for your process in the Functions field. You may optionally enter a description for the function. DO NOT enter any other fields. The Sequence field is automatically populated and the Navigator Prompt and Submenu fields must remain empty.
4. Save your work.

Access the Seeded Processes from the Database

To access the seeded processes, use the following procedure:

1. Run the Oracle Workflow Builder from your client PC.
2. Select Open from the File menu.
3. Choose Database.
4. In the User field, enter the user name for the APPS schema.
5. In the Password field, enter the password for the APPS schema.
6. In the Connect field, enter the connect string or alias for your database as shown in the `<ORACLE_HOME>\network\admin\tnsnames.ora` file on your client PC:
7. In the Show Item Types window, select the item type(s) associated with the seeded processes you wish to view. To select more than one item type, hold down your control key as you select the item types. Choose Show, and then choose OK.

Find the Form Function Name

Use the following procedure to find the form function name:

1. Log into Oracle E-Business Suite and navigate to the form of interest.
2. Choose About Oracle Applications... from the Help menu. Scroll down to Form Information and make note of the form name.
3. Now log into Oracle E-Business Suite using the Implementation System Administration responsibility and navigate to /Application/Form. Within the Form window, query for the form name you just made a note of in the Form field.
4. Make note of the value in the User Form Name field once your query completes.
5. Close the Form window and navigate to /Application/Function. Within the Function window, query for the User Form Name value that you just made a note of in the Form field.
6. The value that is returned in the Function field is the form function name that you need to associate a Process Navigator process step to a form.

Index

Symbols

%APPL_TOP%\envshell.cmd, 7-53
<language>_<filename>_ldt.log, 2-25
multi-node systems
 distributing tasks across, 7-21

A

abandon
 AD utilities command line argument, 7-44
abandoned node
 in patching, 3-18
abort
 command
 stopping a utility, 7-55
ACKNOWLEDGE_QUIT
 defaults file option, 7-60
Action Detail report, 5-14
Action Summary report, 3-57, 5-13
AD_DEFERRED_JOBS
 table, 7-55
AD_FILES table
 reviewing for customized files, 2-29
ad_monitor, 3-63
AD_SNAPSHOT_BUGFIXES table
 role in maintaining snapshot information, 7-72
AD_SNAPSHOT_FILES table
 role in maintaining snapshot information, 7-72
AD_SNAPSHOTS table
 role in maintaining snapshot information, 7-72
adadmin, 7-30
 relinking, 7-11
AD Administration, 7-30, 7-38, 7-44, 7-55
 command, 7-66
 copying files to destinations, 7-70
 exiting, 7-66
 Generate Applications Files menu, 7-62
 generating product files, 7-2
 interface, 7-62
 log files, 7-39
 main menu, 7-62
 overview, 7-61
 relinking Oracle E-Business Suite executable programs, 7-69
 running interactively, 7-65
 running non-interactively, 7-62
 running non-interactively, 7-62
 task-specific prompts, 7-62
 Timing Details report, 5-21
 Timing Information, 8-1
 using with defaults file, 7-63
adadminsvctl.cmd, 7-24
adadminsvctl.sh, 7-24
adalnctl.cmd, 7-24
adalnctl.sh, 7-24
adapcctl.cmd, 7-23
adapcctl.sh, 7-23
adaprstctl.cmd, 3-63
AD Check Digest, 7-30
adchkdig, 7-30
adcmctl.cmd, 7-23, 7-25
adcmctl.sh, 7-23, 7-24
adconfig.txt, 7-40
AD Configuration, 7-30
AD Configuration report
 and AD Splicer, 7-5

- AD Controller, 7-30, 7-38, 7-39, 7-45, 7-55, 9-1
 - in parallel processing, 7-55
 - reviewing worker status, 9-7
 - shutting down a manager, 9-9
 - using restart option, 9-7
 - using to monitor workers, 9-1
- adctrl, 7-30, 7-55, 9-1
 - relinking, 7-11
- adctrl.log
 - AD Controller log file, 7-55
- addbctl.cmd, 7-29
- addbctl.sh, 7-28
- addlnctl.cmd, 7-28
- addlnctl.sh, 7-27
- adhrept.sql, 3-60
- AD File Character Set Converter, 7-31
- AD File Identification, 7-30, 8-3
- adformsrvctl.cmd, 7-23
- adformsrvctl.sh, 7-23
- adgennls.pl
 - See* Translation Synchronization Patch utility
- adident, 7-30, 8-3
- AD Job Timing Report, 7-31, 8-1
- adlibin.log, 2-25
- adlibout.log, 2-25
- admanagedsrvctl.cmd, 7-23
- admanagedsrvctl.sh, 7-23
- AD Merge Patch, 7-31, 7-54
 - about, 2-44
 - admrgpch, 2-46
 - and NLS patches, 7-13
 - destination directory, 2-44
 - directory structure, 2-44
 - merging patches with, 2-44
 - merging patch zip files, 2-45
 - source directory, 2-44
 - superseded by adop, 1-8
 - using, 3-37
- Administer Folders, 6-23
- Administering Oracle E-Business Suite Secure Enterprise Search
 - Administering Crawl, 20-72
 - Administering Searchable Objects, 20-67
 - Administrative Tasks, 20-26
 - Administrative Tasks for SES Integration, 20-31
 - Architecture, 20-4
 - Assigning Responsibility, 20-30
 - Crawl Time, 20-7
 - Deploying Searchable Objects, 20-68
 - Design Time, 20-6
 - Enabling Searches in EBS, 20-28
 - Error Messages, 20-82
 - Overview, 20-1
 - Performing Administrative and Setup Tasks, 20-25
 - Personalization Setup Steps, 20-30
 - Query Time, 20-13
 - Search Administrator, 20-25
 - Search Setups in EBS, 20-28
 - Securing Searchable Objects, 20-38, 20-67
 - Setting Language Preferences, 20-28
 - Setting Profile Options, 20-28
 - Setup Overview, 20-26
 - Setups in SES for Oracle E-Business Integration, 20-36
 - Terms and Definition, 20-2
 - Testing Setups, 20-77
- Administering Searchable Objects
 - Administering Crawl, 20-72
 - Deploying Searchable Objects, 20-68
 - Securing Searchable Objects, 20-67
- Administering Searches
 - Configuring Proxy Parameters, 20-34
 - Installing, 20-33
- Administering Oracle E-Business Suite Secure Enterprise Search
 - Additional Administrative Tasks, 20-78
 - Managing Crawling Schedules, 20-79
 - Optimizing Indexes, 20-81
- admrgpch, 1-8, 2-46, 7-31
- admsi.pl, 1-8, 2-2, 3-5, 7-32
 - customized installation instructions, 3-35
- adncnv, 7-31
- adnodemgrctl.cmd, 7-24
- adnodemgrctl.sh, 7-24
- AD online patching
 - See* adop
- adop, 1-8, 7-31, 7-38, 7-44, 7-54
 - arguments
 - apply, 2-34
 - driver, 2-35
 - patchtop, 2-35
 - preinstall, 2-35

- uploadph, 2-36
- command line arguments, 2-33, 2-41
- description of, 1-8
- how to run, 2-41
- in online patching cycle, 3-6
- introduction, 1-8, 2-2
- log files, 3-36
- modes, 2-31
- non-interactive patching, 3-38
- option
 - parallel, 2-40
 - phtofile, 2-41
- options
 - about, 2-36
 - actiondetails, 2-36
 - autoconfig, 2-37
 - checkfile, 2-37
 - compiledb, 2-38
 - compilejsp, 2-38
 - copyportion, 2-38
 - databaseportion, 2-39
 - forceapply, 2-39
 - generateportion, 2-39
 - hotpatch, 2-40
 - integrity, 2-40
 - validate, 2-41
- phase order, 2-29
- prompts, 2-26
- responding to prompts, 3-35
- restarting, 2-43
- running
 - in pre-install mode, 2-32
 - in test mode, 2-33
 - non-interactively, 2-29
 - specifying parallel workers, 2-28
 - Timing Details report, 5-18
 - use in interactive patching, 3-33
 - validating passwords, 2-41
- adop.log
 - log files, 2-23
- adop arguments, 2-34
- adopmnctl.cmd, 7-23
- adopmnctl.sh, 7-23
- adop remote invocation
 - in multi-node environments, 3-17
 - requirement for ssh, 3-3
- adovars.cmd, 7-41
- adovars.cmd file, 7-70
- adovars.env, 7-41
- adovars.env file, 7-70
- adpchlst.sql, 3-61
- adphrept.sql, 3-60
- AD Relink, 7-10, 7-31
 - overview, 7-61
- adrelink.log, 2-25
- adrelink.sh, 7-31
- AD reporting utilities
 - categories, 8-3
- adsplce, 7-31
- AD Splicer, 7-31
 - and off-cycle products, 7-4
- AD Splicer
 - control files, 7-4
- adstpall.cmd, 7-24, 7-27
- adstpall.sh, 7-26
- adstpall.sh , 7-24
- adstrtal.cmd, 7-24
- adstrtal.sh, 7-24
- adstrtl.cmd, 7-26
- adstrtl.sh, 7-25
- adt<session_id>.lst, 5-16
- adtimrpt.sql, 7-31
- adutconf.sql, 7-30
 - AD Configuration script, 8-2
- advrfapp.sql, 7-15
- adworker
 - worker program, 7-55
- adworkxxx.log, 2-25
- ADXCKPIN.sql
 - SGA monitoring script, 7-19
- ADXGNPIN.sql, 3-36
 - pinning script, 7-18
- ADXGNPNS.sql, 3-36
 - pinning script, 7-18
- AFPASSWD utility, 6-4
- Agent listeners, 15-64
- Agents, 15-110
- APPL_ CONFIG
 - environment variable, 7-53
- APPL_TOP name, 2-27
- applcust.txt file, 4-23
- Application
 - registering, 6-18
- Application basepath, 6-18

- Application environment variable, 6-18
- Applications DBA utilities
 - definition, 7-29
- Applied Patch Check field, 5-32
- Applied Patches
 - about, 3-54
 - Action Detail report, 5-14
 - Action Summary report, 3-57, 5-13
 - Advanced Search page, 5-5
 - Bug Fixes report, 3-57, 5-12
 - description of, 1-8
 - Files Copied report, 3-56, 5-12
 - Oracle Applications Manager utility, 7-33
 - page navigation, 5-2
 - Patch Details report, 3-55, 5-9
 - reporting tool, 5-1
 - Simple Search page, 5-3
 - using, 3-54
- applmgr
 - account, 7-52
- applora.txt, 7-41
- applorau.txt, 7-42
- applprod.txt, 7-42
- applsyst
 - ORACLE ID, 6-15
- APPLSYSPUB
 - Oracle E-Business Suite public schema, 7-77
- appltterr.txt, 7-42
- apply=no, 3-37
- APPS accounts
 - password, 6-15
- APPSCONTEXT_NAME>.env, 7-53
- APPS schema
 - current invalid objects, 2-28
- Assign default folders, 6-23
- AutoConfig
 - Oracle Applications Manager utility, 7-33
- AutoConfig test mode, 3-37
- AutoPatch, 7-29, 7-31, 7-55
 - modes, 5-2
 - role in maintaining snapshot information, 7-72
 - Timing Information, 8-1

B

- Background engines, 15-86
- backup_mode

- optional AD Relink command line argument, 7-10
- BACKUP_SCHEMA_STATS procedure (CBO), 19-10
- BACKUP_TABLE_STATS (CBO), 19-9
- backup files
 - deleting after patching, 2-31
- binary forms files
 - generating with AD Administration, 7-68
- binary message files
 - generating with AD Administration, 7-67
- binary report files
 - generating with AD Administration, 7-68
- Bug Fixes report, 3-57, 5-12
- bug fix patch
 - definition of, 1-2

C

- Caching Framework, 17-1
- Case-sensitive database passwords, 6-12
- Changing passwords with
 - AFPASSWD utility, 6-4
 - FNDCPASS utility, 6-4
- CHECK_HISTOGRAM_COLS procedure (CBO), 19-22
- child processes
 - of worker processes
 - terminating, 9-7
- Client System Analyzer (OAM), 14-39
- codelevels
 - Codelevel Introduced report, 5-11
 - current and new, 4-3
 - definition of, 1-4
- Codelevels Summary page, 4-19
- codelines, 3-44, 4-18, 5-11
 - definition of, 1-3
- compiling invalid objects
 - adadmin task, 7-17
- Concurrent Manager Recovery (OAM), 14-41
- concurrent programs, 4-2
- Configuration
 - logging, 12-5
- configuration and environment files
 - generated by AutoConfig, 7-40
- Configuration Change Detector
 - in file system synchronization, 7-1

- consolidated upgrade patch
 - definition of, 1-3
- Consumer groups
 - Resource consumer groups, 6-3
- CONTROL_CODE
 - column in FND_INSTALL_PROCESSES table, 7-55
- Convert Character Set
 - admanifest_excp.lst file, 7-71
 - admanifest_lossy.lst file , 7-71
 - admanifest.lst file , 7-71
 - lossy conversions, 7-70
 - running, 7-71
 - scanning a CUSTOM directory for exceptions, 7-71
- copy portion
 - unified driver, 2-27
- CP Signature Wizard (OAM), 14-45
- CREATE_STAT_TABLE Procedure (CBO), 19-8
- current working directory
 - and log file storage, 7-55
- custom filters for Patch Wizard, 4-10
- customized files, 4-24
 - in AD_FILES table, 2-29
- customized installation instructions, 3-35

D

- database portion
 - unified driver, 2-28
- Database Resource Manager, 6-3
- Debugging, 12-5
- Debug Workbench, 14-38
- Default folders, 6-23
- defaultsfile
 - AD utilities command line argument, 7-45
- defaults file
 - for non-standard patches, 3-40
 - use in non-interactive maintenance, 7-20
 - use in non-interactive maintenance tasks, 7-19
 - using, 2-20
 - using with AD Controller, 7-60
- deferred jobs
 - use in recovery from job failure, 7-55
- Diagnostics in Oracle Applications Manager, 14-37
- Diagnostics menu, 13-1

- DISABLE_SCHEMA_MONITORING procedure (CBO), 19-12
- DISPLAY
 - setting for PAA, 3-5
- Distributed AD
 - definition, 7-21, 7-55
- driver
 - unified, 2-27
- driver file actions, 3-55
- DUAL table, 7-78

E

- emergency patches
 - applying, 3-38
- ENABLE_SCHEMA_MONITORING procedure (CBO), 19-12
- Environment variable, 6-18
- error messages, 2-30
- Exception Report, 5-27

F

- File History
 - about, 5-6
 - Advanced Search page, 5-8
 - description of, 1-9
 - Oracle Applications Manager utility, 7-34
 - Simple Search page, 5-7
- Files Copied report, 3-56, 5-12
- filters for Patch Wizard
 - custom filters, 4-10
 - pre-seeded, 4-10
- flexfields
 - compiling, 7-76
- FND_INSTALL_PROCESSES
 - table, 7-55
 - dropping, 7-55
- FND_STATS package, 19-8
- FNDCPASS utility, 6-4
- findenv.env, 7-42
- FNDLOAD, 7-55
- FNDWFAASTATCC, 15-2
- FNDWFMLRSTATCC, 15-2
- FNDWFWITSTATCC, 15-2
- Folder Administration, 6-23
- Folder Set, 6-23
- Forms

- Applications, 6-18
 - Register ORACLE IDs, 6-15
- Functional Administrator, 17-1
- Functional Developer, 17-1

G

- GATHER_ALL_COLUMN_STATS procedure (CBO), 19-20
- GATHER_COLUMN_STATS procedure (CBO), 19-19
- GATHER_INDEX_STATS procedure (CBO), 19-16
- GATHER_SCHEMA_STATS procedure (CBO), 19-13
- GATHER_TABLE_STATS procedure (CBO), 19-17
- Gather Schema Statistics
 - usage, 7-14
- generate portion
 - unified driver, 2-28
- Generic Collection Service and Forms Monitoring Wizard (OAM), 14-45

H

- help
 - obtaining in AD command line utilities, 7-34
 - obtaining in Oracle Applications Manager
 - See OAM Help
- help
 - AD utilities command line argument, 7-45
- hide patch, 4-17
- hidepw
 - AD utilities command line argument, 7-50

I

- INFORM_FAILURE
 - defaults file option, 7-60
- informational messages, 2-30
- In-Progress Timing Details page, 5-22
- input_file
 - use in non-interactive patching, 2-29
- input file
 - use in non-interactive patching, 2-22
- interactive
 - AD utilities command line argument, 7-45

- interactive patching
 - definition of, 3-33, 3-38
- In Use flag, 4-9
- Invokerxd5 s Rights processing
 - removed in Release 12, 2-28

J

- JAR files, 7-2
- Java agent listeners, 15-72
- Job Timing report, 5-23

L

- License Manager
 - adding languages with, 7-13
 - Oracle Applications Manager utility, 7-34
- localizations
 - definition and types, 16-9
- localworkers
 - AD utilities command line argument, 7-46
- logfile
 - AD utilities command line argument, 7-46
- log files
 - <language>_<filename>_ldt.log, 2-25
 - adlibin.log, 2-25
 - adlibout.log, 2-25
 - adrelink.log, 2-25
 - adworkxxx.log, 2-25
 - used by AD utilities, 7-39
 - viewing details, 5-30
 - viewing from Timing Reports main page, 5-18
 - viewing from View Log Files page, 5-28
- Logging
 - configuration, 12-1, 12-5
 - C environment variables, 12-1
 - Java system properties, 12-1
 - middle-tier properties, 12-1
 - profile options, 12-1
 - disabling, 12-5
 - in high volume scenarios, 12-5
 - logging to screen, 12-5
 - messages
 - purging, 12-5, 12-8
 - viewing, 12-8
 - overview, 12-1
 - specific user errors, 12-5
 - startup behavior, 12-1

- logging
 - AD utilities command line argument, 7-50
- Logging to Screen, 12-1
 - with CRM Technology Foundation, 12-1
 - with Oracle Application Framework, 12-1
- Logs, 12-10

M

- Maintain Current View Snapshot Information
 - options, 7-74
- maintaining snapshot information
 - APPL_TOP snapshots
 - definition, 7-72
 - current view snapshots
 - definition, 7-72
 - global snapshots
 - definition, 7-72
 - named view snapshots
 - definition, 7-72
- Maintain Snapshot
 - tasks, 7-6
- Maintain Snapshot Information menu
 - options, 7-72
- manual steps for patching, 4-17
- menu_option
 - AD utilities command line argument, 7-47
- merged patches
 - creating, 4-8
 - naming, 2-45
- merge option defaults, 4-8
- merge options, 4-14
- messages
 - error, 2-30
 - informational, 2-30
 - successful completion, 2-31
- modifiers
 - AD utility command line, 7-43
- MRC schema
 - not required, 2-28
- My Oracle Support credentials, 3-43, 4-1

N

- Network latency
 - testing, 6-21
- Network Test window, 6-21
- NLS

- merging patches, 3-42
- multiple patches, 3-43
- prompts for adop, 2-29
- single patch, 3-42
- non-interactive patching
 - abandoning, 3-40
 - definition of, 3-33
 - non-standard patches, 3-40
 - preparing for, 2-29
 - restarting in adop, 3-40
- non-interactive processing, 7-38
- non-standard patch
 - applying non-interactively, 3-40
 - definition of, 3-40
- Notification mailers, 15-19

O

- OAM Help, 7-35
- oamreport.csv, 4-23
- OAM Site Map
 - obtaining help in, 7-36
- OAM Timing Reports feature, 8-1
- off-cycle products, 7-4
- Online Patching Diagnostic Reports, 3-32
- optimization
 - and SQL statements, 7-14
- options=validate, 3-37
- Oracle Application Framework, 17-1
- Oracle Application Object Library, 7-76
- Oracle Applications Manager
 - Applications Dashboard, 14-1
 - Client System Analyzer, 14-39
 - Debug Workbench, 14-38
 - Log, 14-47
 - Monitoring, 14-12
 - Running Web-based patching utilities from, 1-8
 - Support Cart, 14-46
 - System alerts, metrics, and logs, 14-27
 - troubleshooting wizards, 14-40
- Oracle Applications Manager (OAM), 7-33
 - Applied Patches, 5-1
 - File History, 5-6
 - Patch Impact Analysis, 4-19
 - Patch Wizard, 4-3
 - Register Flagged Files, 4-22

- Timing Reports, 5-15
- URL for, 4-4
- OracleConcMgr<CONTEXTNAME>
 - Windows Concurrent manager service, 7-54
- Oracle Database Patchset Updates
 - applying, 1-7
- Oracle E-Business Suite Secure Enterprise Search
 - Administering Oracle E-Business Suite Secure Enterprise Search, 20-1
- ORACLE ID
 - applsys - password warning, 6-15
 - explained, 6-1
 - Oracle username, 6-15
 - registering, 6-1, 6-15
 - requirement for database access, 6-1
- Oracle Patch Application Assistant
 - uses, 3-1
- Oracle Patch Application Assistant
 - use with readme files, 1-2
- ORACLE schemas, 6-1
- OracleService<SID>
 - Windows Database service, 7-54
- ORACLE usernames, 6-1
- Oracle Workflow
 - managing from Oracle Applications Manager, 7-33
 - system status, 15-3
- Oracle Workflow Manager, 15-1

P

- PAA
 - See Patch Application Assistant, 3-5
- packages
 - pinning, 7-18
- parallel_index_threshold
 - AD utilities argument
 - choosing values, 7-55
- parallel_index_threshold
 - AD utilities command line argument, 7-47
- parallel processing
 - uses, 7-55
- parallel workers
 - specifying in adop, 2-28
- partial snapshot, 7-7
- passwords
 - validating for patching, 3-37
- patch
 - copy portion of a unified driver, 2-27
 - database portion of a unified driver, 2-28
 - file structure, 1-1
 - generate portion of a unified driver, 2-28
 - information files, 5-1
 - storing history, 5-1
 - top-level directory files, 1-1
 - unified driver, 2-27
- Patch Application Assistant, 7-32
 - about, 1-8, 2-1
 - completed steps, 3-5
 - setting DISPLAY variable, 3-5
 - static README.html file, 3-5
 - using, 3-5
 - with adop, 2-42, 3-34
- Patch Details report, 3-55, 5-9
- patch directory, specifying, 2-26
- patch driver file
 - about, 2-27
- patches, 3-10
 - determining if applied, 3-54
 - directions for applying, 3-33
 - emergency, 3-38
 - merging
 - destination directory, 2-44
 - source directory, 2-44
 - monitor the progress of, 3-63
 - non-standard, 3-40
 - Patch Impact Analysis report, 3-37
 - prerequisite, 3-33
 - readme files, 3-34
 - recommended, 3-44
 - testing, 3-37
 - tracking prerequisites, dependencies, and compatibilities, 1-4
 - unzipping, 3-34
 - using Patch Wizard, 3-37
- patch formats
 - bug fix, 1-2
 - consolidated upgrade patch, 1-3
 - pre-upgrade, 1-2
 - product release update pack, 1-2
 - release update pack, 1-2
- patch history
 - database, 3-54
 - searching for, 5-3

- patch history XML reports
 - adfhrept.sql, 3-60
 - adpchlst.sql, 3-61
 - adphrept.sql, 3-60
- Patch Impact Analysis report, 3-37, 4-17
- patch information bundle
 - about, 4-2
- patching
 - about, 3-33
 - applying part of a unified driver, 3-36
 - customization issues, 3-35
 - interactive, 3-33
 - NLS systems, 3-42, 3-43
 - non-interactive, 3-33, 3-38
 - password validation, 3-37
 - patch top directory, 3-34
 - pinning SGA packages, 3-36
 - reviewing log files, 3-36
 - supplying driver name, 3-35
 - supplying unified driver name, 3-36
 - test mode, 3-37
- Patch Recommendation Requests section, 5-32
- Patch Summary report, 3-59
- patchtop
 - definition, 2-26
- Patch Wizard
 - accessing, 4-4
 - and snapshots, 7-6
 - concurrent programs, 4-2
 - description of, 1-9
 - details icons, 4-4
 - downloading patches
 - display option defaults, 4-9
 - language and platform details, 4-9
 - merge option defaults, 4-8
 - products in-use in your system, 4-9
 - filters
 - creating, 4-11
 - custom filters, 4-10
 - defining, 4-10
 - pre-seeded filters, 4-10
 - generating patch recommendations, 4-12
 - main page, 4-3
 - Main page, 4-4
 - Oracle Applications Manager utility, 7-34
 - Patch Impact Analysis, 4-19
 - patch information bundle, 4-2
 - Preferences page, 4-6
 - Recommended Patches Results page, 4-15
 - Recommend Patches page, 4-11
 - role in maintaining snapshots, 7-72
 - scheduling requests, 4-12
 - setting up, 4-6
 - task icons, 4-3
 - using, 3-37, 3-50
- Personalization
 - Oracle Application Framework, 17-1
- Phase Information report, 5-25
- phase order
 - in adop, 2-29
- phases
 - of online patching, 2-22
 - of online patching cycle, 2-22
- PL/SQL agent listeners, 15-64
- Predefined alerts
 - action sets - definition of, 18-3
 - alert action - definition of, 18-2
 - alert - definition of, 18-2
 - customizing, 18-4
 - DBA alerts, 18-7
 - event alert - definition of, 18-2
 - explained, 18-2, 18-3
 - overview of Oracle Alert, 18-1
 - periodic alert - definition of, 18-2
 - precoded custom alerts, 18-6
 - purging alerts, 18-10
 - using, 18-3
- Preferences page, 4-6
- Preinstall_Codelevel_AD.txt
 - adop file, 2-32
- Preinstall_Codelevel_MP.txt
 - adop file, 2-32
- pre-install mode
 - arguments for, 2-35
 - information not stored when using, 3-54
 - running adop in, 2-32
- prerequisite patches
 - when to apply, 3-33
- pre-seeded filters, 4-10
- pre-upgrade patch
 - definition of, 1-2
- pre-upgrade patches, 4-10
 - disk space for, 2-31
- primary node

- in multi-node environment, 3-17
- printdebug
 - AD utilities command line argument, 7-48
- private/public key pair
 - in ssh, 3-3
- Process Navigator
 - overview, 21-1
- product family RUPs, 1-2, 4-10
 - disk space for, 2-31
- product JAR files
 - generating with AD Administration, 7-68
- product release update pack
 - definition of, 1-2
- PURGE_STAT_HISTORY procedure (CBO), 19-22
- Purge Debug Log and System Alerts concurrent program, 12-8
- Purging
 - Oracle Workflow data, 15-90
- Purging (OAM), 14-48

Q

- Queue propagation, 15-114

R

- Rapid Install, 7-32
 - role in maintaining snapshot information, 7-72
- rapidwiz, 7-32
- README.html file
 - using with PAA, 2-42
- readme.txt
 - function, 1-2
- readme files, 3-34, 4-2, 4-20
 - in top-level directory, 1-1
 - post patching steps in, 2-30
 - reviewing manual steps, 2-42
- Recommended Patches Results page, 4-15
- Recommend Patches page, 4-11
- recreate grants and synonyms
 - adadmin task, 7-17
- Register
 - application, 6-18
- Register Flagged Files
 - about, 4-22
 - accessing, 4-22
 - adding, 4-24
 - description of, 1-9
 - exporting, 4-23
 - importing, 4-23, 4-25
 - Oracle Applications Manager utility, 7-34
 - registering customized files, 3-35
- release update pack, 7-3
- release update packs (RUPs), 4-10
 - definition of, 1-5
 - disk space for, 2-31
- relinkenv.cmd, 7-10
- reports
 - Action Detail, 5-14
 - Action Summary, 5-13
 - AD Administration Timing Details, 5-21
 - adop Timing Details, 5-18
 - Bug Fixes, 5-12
 - Exception, 5-27
 - Files Copied, 5-12
 - Job Timing, 5-23
 - Patch Details, 3-55, 5-9
 - Patch Impact Analysis, 3-37, 4-17
 - Patch Summary, 3-59
 - Phase Information, 5-25
- request set ID, 4-13
- Resource consumer groups, 6-3
- restart
 - utility, 7-55
- restart
 - AD utilities command line argument, 7-48
- RESTART_JOB
 - defaults file option, 7-60
- restart=yes, 3-40
- restart file
 - worker, 7-55
- restart files, 7-39
- RESTORE_COLUMN_STATS procedure (CBO), 19-11
- RESTORE_SCHEMA_STATS procedure (CBO), 19-10
- RESTORE_TABLE_STATS procedure (CBO), 19-11
- restricted mode
 - URL, 3-63
 - using, 3-63
- RUP
 - suite release update pack, 1-2

S

- Search Security Plug-ins
 - ACL Security, 20-49
 - How Security Plug-ins Works, 20-45
 - Other Considerations, 20-61
 - Query Rewrite Security, 20-51
 - Security Logic and General Plug-ins, 20-47
 - Security Models with Search Plug-ins, 20-55
- secondary nodes
 - in multi-node environment, 3-17
- Securing Searchable Objects
 - Creating Grants, 20-40
 - RBAC Security, 20-39
 - Search Security Plug-ins, 20-43
 - UAC, 20-63
- Security Patch Updates
 - applying, 1-7
- Service components, 15-6
- Service Infrastructure (OAM), 14-43
- SGA
 - pinning packages after patching, 3-36
- shared pool
 - monitoring, 7-19
- SHOW_STATUS
 - defaults file option, 7-60
- SHUTDOWN_WORKER
 - defaults file option, 7-60
- snapshot
 - definition, 7-6
- Software Updates
 - about, 5-31
 - accessing, 3-54, 5-31
 - Applied Patch Check, 5-32
 - description of, 1-9
 - home page, 5-31
 - Maintenance Activities, 5-32
 - Oracle Applications Manager utility, 7-34
 - Patch Recommendation Requests, 5-32
 - Related Links, 5-33
- SQL*Plus
 - validating APPS schema from, 7-16
- ssh
 - requirement for in multi-node application tiers, 3-3
- staging directory, 4-3

- START_WORKER
 - defaults file option, 7-60
- Statistics
 - gathering for Oracle Workflow, 15-2
- STATUS
 - column in FND_INSTALL_PROCESSES table, 7-55
 - successful completion message, 2-31
- Support Cart (OAM), 14-46
- Supporting Security Models with Search Plug-ins
 - Business Group, 20-57, 20-60
 - Legal Entity, 20-59
 - Organization, 20-59
- System Global Area, 7-18

T

- Technology Inventory Utility, 10-1
- Testing Setups
 - General setups, 20-77
 - Testing Deployment, 20-77
 - Testing Schedules, 20-78
 - Testing Searches, 20-78
- test mode
 - arguments for, 2-34
 - information not stored when using, 3-54
 - running adop in, 2-33
 - using, 3-37
- Timing Reports
 - about, 5-15
 - accessing, 5-17
 - AD Administration Timing Details report, 5-21
 - adop Timing Details report, 5-18
 - description of, 1-9
 - Exception report, 5-27
 - Job Timing report, 5-23
 - Main Page, 5-17
 - Oracle Applications Manager utility, 7-34
 - Phase Information report, 5-25
 - restricted mode, 3-63
- trace
 - AD utilities command line argument, 7-51
- translations
 - searching for patches, 3-58
- Translation Synchronization Patch, 7-14
- Translation Synchronization Patch utility, 7-13, 7-

13, 7-14

TWO_TASK

environment variable, 7-53

U

unified driver

about, 2-27

applying part of, 3-36

copy portion, 2-27

database portion, 2-28

generate portion, 2-28

in top-level directory, 1-1

unified driver file

description of, 1-2

updates

see patching, 3-33

V

validate APPS schema

adadmin task, 7-16

VERIFY_STATS procedure (CBO), 19-23

View Log Details page, 5-30

View Log Files page, 5-28

W

wait_on_failed_job

AD utilities command line argument, 7-49

Web services outbound components, 15-80

Windows Task Manager, 9-1

worker log files, 9-1

workers

affected by restart, 2-43

in deferred jobs, 2-31

number of, 2-28

recording actions of, 2-25

workers

AD utilities command line argument, 7-49

Workflow Agent Activity Statistics Concurrent Program, 15-2

Workflow control queue, 15-96

Workflow Mailer Statistics Concurrent Program, 15-2

Workflow Notification Mailer

integration with Oracle Alert, 18-4

Workflow Work Items Statistics Concurrent

Program, 15-2

Work items

active, 15-98

completed, 15-92

deferred, 15-101

errored, 15-107

suspended, 15-104